

Administrator's Guide

ArcSight Logger™ v4.5 GA

July 6, 2010



Administrator's Guide ArcSight Logger™ v4.5 GA

July 6, 2010

Copyright © 2006-2010 ArcSight, Inc. All rights reserved.

ArcSight, the ArcSight logo, ArcSight TRM, ArcSight NCM, ArcSight Enterprise Security Alliance, ArcSight Enterprise Security Alliance logo, ArcSight Interactive Discovery, ArcSight Pattern Discovery, ArcSight Logger, FlexConnector, SmartConnector, SmartStorage and CounterACT are trademarks of ArcSight, Inc. All other brands, products and company names used herein may be trademarks of their respective owners.

Follow this link to see a complete statement of ArcSight's copyrights, trademarks, and acknowledgements:
<http://www.arcsight.com/company/copyright/>

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

This document is ArcSight Confidential.

Revision History

Date	Product Version	Description
07/06/10	v4.5 GA	v4.5 GA.
01/22/10	v4.0 SP1	v4.0 SP1 release.
11/03/09	v4.0 GA	v4.0 GA release.
07/22/09	v3.0 SP1	Service pack 1 for version 3.0.
01/07/09	v3.0 Patch 1	v3.0 Patch 1 version.
10/16/08	v3.0	v3.0 GA version.
06/30/08	v2.5.1	v2.5 SP1 release.
05/29/08	v2.5	v2.5 GA release.
11/19/07	v2.0 Patch 1	v2.0 guide; added Report Category Filter section.
10/19/07	v2.0	v2.0 guide.

Document template version: 1.0.2.9

ArcSight Customer Support

Phone	1-866-535-3285 (North America) +44 (0)870 141 7487 (EMEA)
E-mail	support@arcsight.com
Support Web Site	https://support.arcsight.com
Protect 724 Community	https://protect724.arcsight.com

Contents

About the Online Help	xv
Chapter 1: Overview	1
Introduction	1
Logger Features	3
Storage Configuration	3
Receiver Configuration	3
Analyzing Events	4
Grouping Events	5
Exporting	5
Forwarder Configuration	5
User Management	6
Other Setup and Maintenance	6
Deployment Scenarios	6
What's New in Version 4.5 GA	9
Chapter 2: Installation and Initialization	15
Section 1: Installing the Logger Appliance	15
Installing the Logger Appliance	15
Setting an IP Address on the Appliance	16
Initializing the Logger Appliance	19
Configure DRAC for Remotely Accessing a Logger Appliance	19
Section 2: Installing the Software Version of Logger	19
How Licensing Works on the Software Version of Logger	19
Supported Platforms and Browsers	21
Installing the Software Version of Logger for the First Time	22
Initializing the Software Version of Logger	22
Applying a New License on the Software Version of Logger	22
Uninstalling the Software Version of Logger	22
Initialization for all Loggers	22
Planning	22
Storage Strategy	22
Retention Policy	23
Peer Loggers	23

Initialization Sequence (for all Loggers)	24
1 License	24
2 SAN	25
3 Storage Volume	25
3 Storage Groups	26
4 Time Settings	26
5 Index Fields and Full-text Indexing	27
6 Reboot	28
7 Receivers	28
8 Devices	28
9 Device Groups	28
10 Storage Rules	28
Installing SmartConnectors to Send Events to Logger	29
SmartMessage	30
Forwarding Logger Events to an ESM Manager	30
Configuring SmartConnectors to Send Events to Both Logger and an ESM Manager	31
Configuring SmartConnectors for Failover Destinations	31
Sending Events from ArcSight ESM to Logger	31
Chapter 3: Using the User Interface	35
Logger User Interface	35
Browser Requirements	35
Navigating the User Interface	36
Help	36
Options	37
Logout	37
Monitor	37
Platform	39
Network	39
Logger	40
Receivers	40
Forwarders	40
Storage	40
Chapter 4: Searching and Analyzing Events	43
The Need to Search Events	43
The Process of Searching Events	43
Elements of a Search Query	44
Query Expression	44
Syntax Reference for Query Expression	60
Using the Search Builder Tool	64
Accessing Search Builder	65
Nested Conditions	67

Alternate Views for Query Building in Search Builder	67
Search Analyzer	68
Regex Helper Tool	70
Searching for Events on Logger	72
Advanced Search Options	73
Searching Peer Loggers (Distributed Search)	73
Understanding the Search Results Display	74
Multi-line Data Display	76
Auto Updating Search Results	77
Exporting Search Results	77
Scheduling an Export Operation	80
Indexing	80
How indexing works	80
Full-text Indexing (Keyword Indexing)	81
Field-based Indexing	81
Saving Queries (Saved Filters and Searches)	84
Saving a Query	85
Using a Saved Filter or a Saved Search	85
System Filters/Predefined Filters	86
Using a System Filter	88
Monitoring System Health	88
System Health Events	88
Alerts	89
Viewing Alerts	90
Receiving Alerts for Events	90
Base Event Fields	91
Go, Export, and Auto Update Options	91
Chapter 5: Reporting	93
Navigating to Reports	93
Report Groups	94
Foundation Reports	95
SANS Top 5 Reports	95
Network Monitoring Reports	96
Intrusion Monitoring Reports	96
Configuration Monitoring Reports	97
Solution Reports	97
Device Monitoring Reports	97
Anti-Virus Reports	98
Cross Device Reports	98
Database Reports	98
Firewall Reports	98
Identity Management Reports	98

IDS-IPS Reports	98
Network Reports	98
Operating System Reports	98
VPN Reports	98
User Reports	98
Reports Home Page	99
Using the Dashboard	100
Viewing the Dashboard	100
Designing Dashboards	101
What items can a dashboard include?	102
Quick Start - Creating a New Dashboard	102
Add an Empty Dashboard	103
Creating Widgets	105
Placing Dashboard Items on the Layout	105
Placing a Report on a Dashboard	105
Placing a Use Case on a Dashboard	108
Placing an External Link on a Dashboard	110
Swapping Items on Widgets	111
Setting Dashboard Preferences	112
Working with Available Dashboards	112
Selecting a Dashboard View	112
Modifying or Removing Existing Dashboards	113
Running, Viewing, and Publishing Reports	113
Best Practices	114
Finding Reports	114
Task Options on Available Reports	114
Running and Viewing Reports	116
About the Pagination of Reports	116
Quick Run / Run In Background Report Parameters	117
Run Report Parameters	119
Report File Formats	120
Publishing Reports	121
Report Delivery Options	122
Refreshing a Report	122
E-mailing a Report	122
Exporting and Saving a Report	123
Viewing the Output of a Published Report	124
Designing Reports	124
Opening the Report Designer	125
Creating New Reports	125
Quick Start: Base a New Report on an Existing One	125
Designing New Reports	128
Select Filter Criteria	130

Select Grouping	132
Select Totals	134
Sort Order	134
Highlighting	135
Create Matrix	135
Create Chart	136
Editing a Report	138
Adhoc Report Designer	139
Setting Access Rights on Reports	140
Setting up Queries	140
How Search and Report Queries Differ	141
Overview of Query Design Elements	141
Creating a Copy of an Existing Query	142
Designing a New SQL Query	142
Modifying a Query Object	155
Deleting a Query Object	156
Defining SQL in the Editor	156
Working with Parameters	163
Creating New Parameters	164
Modifying a Parameter	169
Deleting a Parameter	169
Configuring Parameter Value Groups	170
Applying Report Template Styles	172
Defining a New Template	173
Scheduling Reports	173
Viewing and Editing Scheduled Reports	174
Scheduling a Report	174
Deploying a Report Package	177
Report Server Administration	178
Using Report Category Filters	180
Backup and Restore of Report Content	180
Chapter 6: Configuration	181
Configuration	181
Devices	181
Devices	182
Device Groups	183
Event Archives	184
Event Archives	185
Scheduled Event Archive	186
Archive Storage Settings	186
Storage	187
Storage Groups	187

Storage Rules	189
Storage Volume	191
Event Input/Output	193
Receivers	193
Forwarders	199
ESM Destinations	203
Forwarding Log File Events to ESM	207
Alerts	207
Configuring and Managing Real Time Alerts	210
Creating and Managing Saved Search Alerts	212
SNMP Destinations	218
Syslog Destinations	219
ESM Destinations	220
Export	221
Scheduled Tasks	221
Scheduled Tasks	221
Currently Running Tasks	222
Finished Tasks	222
Filters	222
Filters	222
Search Group Filters	223
Export	225
Saved Searches	225
Saved Searches	225
Scheduled Saved Search	226
Saved Search Files	229
Search Optimization	229
Add Search Indexes	229
Tuning Advanced Search Options	230
Deleting Custom Field Sets	231
Peer Loggers	231
Guidelines	232
Configuration Backup and Restore	235
Running a Configuration Backup (Ad-hoc or Scheduled)	236
Restoring from a Configuration Backup	237
Editing Configuration Backup Settings	237
System Maintenance	238
Database Defragmentation	239
Storage Volume Size Increase	244
License Information	245
Retrieve Logs	246
Exporting and Importing Content	247
Guidelines for Exporting and Importing	247

Exporting Content	248
Importing Content	249
Chapter 7: System Admin	251
Section 1: Logger Appliance System Administration	251
Reboot	252
DNS Settings	253
Hosts	253
Network	254
Time/NTP	256
SMTP Settings	258
Static Routes	259
License & Update	259
Process Status	260
Support Login	261
Logs - Audit and Error	262
Logs - Audit Forwarding	262
Storage	263
CIFS Settings	263
Network File System (NFS) Settings	265
SAN	267
Restoring a SAN	269
RAID Controller	270
Security	270
SSL Server Certificate	271
SSL Client Authentication (CAC Authentication)	273
FIPS 140-2	275
Users/Groups	279
Authentication Settings	279
Groups	284
Users	289
Change Password	290
Section 2: Software Version Logger Administration	291
Network - SMTP Settings	291
Process Status	292
Logs - Audit and Error	293
Logs - Audit Forwarding	293
Users/Groups - Groups	294
User Groups	294
Users	299
Users/Groups - Change Password	300
Using a CA-signed Certificate on Software Version of Logger	301


Chapter 8: Managing Connectors on Connector Appliance	303
SmartConnector Overview	304
Navigating the Manage Connectors Tab	305
Locations	307
Viewing All Locations	307
Viewing Hosts, Containers, and Connectors in a Location	307
Adding a Location	308
Adding Locations and Hosts from a File	308
Editing a Location	309
Deleting a Location	309
Adding Hosts to a Location	309
Hosts	310
Viewing All Hosts	310
Viewing Containers and Connectors in a Host	310
Adding a Host	311
Scanning a Host	313
Deleting a Host	314
Moving a Host to a Different Location	314
Editing a Host	315
Upgrading a Host Remotely	315
Adding a Container to a Host	315
Containers	316
Viewing All Containers	316
Viewing Connectors in a Container	317
Adding a Container	317
Adding a Connector to a Container	317
Editing a Container	317
Deleting a Container	318
Updating Container Properties	318
Changing Container Credentials	319
Enabling and Disabling FIPS on a Container	320
Managing Certificates on a Container	321
Enabling or Disabling a Demo Certificate on a Container	321
Adding CA Certificates on a Container	322
Adding a CA Certs File on a Container	323
Removing CA Certificates from a Container	324
Viewing Certificates on a Container	325
Resolving Invalid Certificate Errors	327
Running a Command on a Container	327
Upgrading a Container to a Specific Connector Version	328
Viewing Container Logs	329
Deleting Container Logs	329
Running Logfu on a Container	329

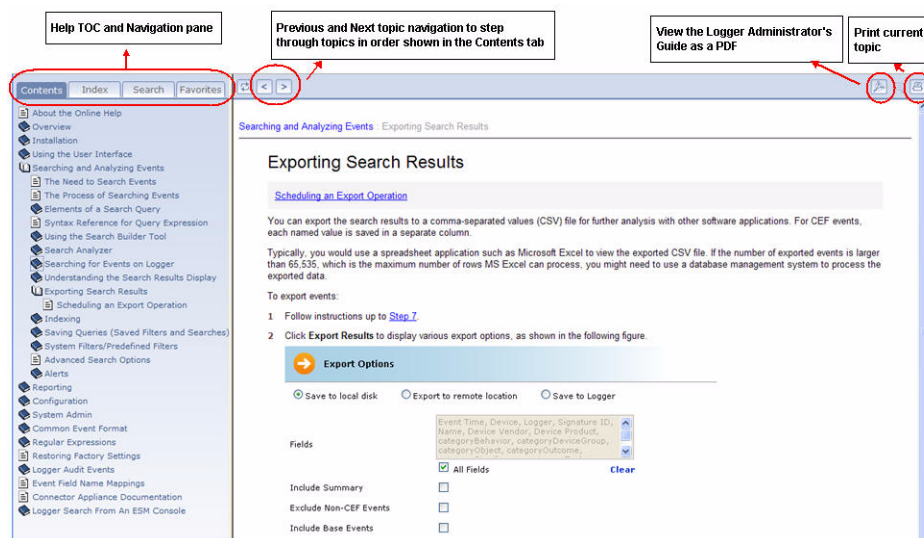
Connectors	331
Viewing all Connectors	331
Adding a Connector	331
Editing Connector Parameters	334
Updating Simple Parameters for a Specific Connector	334
Updating Table Parameters for a Specific Connector	336
Managing Connector Destinations	337
Adding a Primary Destination to a Specific Connector	337
Adding a Failover Destination to a Specific Connector	338
Adding a Primary or Failover Destination to Multiple Connectors	339
Removing Destinations	340
Editing Destination Parameters	341
Editing Destination Runtime Parameters	342
Managing Alternate Configurations	344
Sending a Command to a Destination	346
Removing a Connector	347
Sending a Command to a Connector	347
Running Logfu on a Connector	348
Configuration Suggestions for SmartConnector Types	349
Deploying FlexConnectors	350
Configuring the Check Point OPSEC NG Connector	350
Adding the MS SQL Server JDBC Driver	352
Troubleshooting Connector Communication Issues	353
Chapter 9: Managing Repositories in Connector Appliance	355
Overview	356
Logs Repository	358
Uploading a File to the Logs Repository	358
CA Certs Repository	359
Uploading CA Certificates to the Repository	360
Removing CA Certificates from the Repository	360
AUP Repository	361
About the AUP Upgrade Process	361
Uploading an AUP Upgrade File to the Repository	361
Removing a Connector Upgrade from the Repository	362
Content AUP Repository	363
Applying a New Content AUP	363
Applying an Older Content AUP	364
Emergency Restore	364
User-Defined Repositories	365
Creating a User-Defined Repository	365
Retrieving Container Files	367
Uploading Files to a Repository	367

Deleting a Repository	367
Updating Repository Settings	368
Managing Files in a Repository	369
Retrieving a File from the Repository	369
Uploading a File from the Repository	369
Pre-Defined Repositories	370
Cloning Container Configuration	371
Appendix A: Common Event Format	373
Common Exchange Format	373
Common Extension Dictionary	375
Appendix B: Regular Expressions	379
Regex Overview	379
Simple Regular Expressions	380
Metacharacters	380
Forbidden Characters	385
Things To Remember	386
Appendix C: Using the Rex Operator	387
Syntax of the rex Operator	387
Understanding the rex Operator Syntax	387
Ways to Create a rex Expression	388
Creating a rex Expression Manually	388
Samples of rex Expressions	389
Appendix D: Restoring Factory Settings	393
Appendix E: Logger Audit Events	401
Types of Audit Events	401
Information in an Audit Event	401
Platform Events	402
Logger Application Events	408
Appendix F: Event Field Name Mappings	421
Appendix G: Connector Appliance Documentation	427
Appendix H: Destination Runtime Parameters	429
Appendix I: Logger Search From An ESM Console	437
Understanding the Integrated Search Functionality	437
Prerequisites	438
Setup and Configuration	438
ESM	438




Logger	438
Supported Search Options	439
Guidelines	439
Searching on Logger From ESM Console	439
Index	441

About the Online Help

Logger online Help is available through the ArcSight Logger's web user interface (UI). Click the  icon in the top right-hand corner of any Logger UI page to access the Help for that page.



The online Help includes the following features:


- Left panel Help navigation - Click a tab for Contents (TOC), Index, Search, or Favorites.
- Next, Previous topic navigation- Click the Previous button () to view the preceding topic in the history, or the Next button () to view the subsequent topic.
- Topic display window - Click a topic in the Contents, Index, Search hit list, or saved Favorites to view it in the display window.
- Access to the Administrator's Guide as an Adobe Acrobat PDF document. Click the PDF button () in the upper right of the Online Help toolbar to get PDF.

- Print capabilities - Click the Print () button to print a copy of the current topic.



Connector Appliance documentation is available as follows:

The Chapter 8, Managing Connectors on Connector Appliance, on page 303 and Chapter 9, Managing Repositories in Connector Appliance, on page 355 chapters in this guide.

Through the Help icon () on any user interface page, when you are in the Connector Appliance context. When you click this icon, a PDF of the Connector Appliance Administrator's Guide is displayed. All information in this guide except system administration is applicable to your product.

Through the ArcSight Customer Support site at <https://support.arcsight.com>.

Chapter 1

Overview

The following topics provide an overview of ArcSight Logger, including information on what's new in this release; storage, receiver, and forwarder configuration; working with events; user management; and setup and maintenance considerations.

["Introduction" on page 1](#)

["Logger Features" on page 3](#)

["Deployment Scenarios" on page 6](#)

["What's New in Version 4.5 GA" on page 9](#)

Introduction

ArcSight Logger is a log management solution that is optimized for extremely high event throughput, efficient long-term storage, and rapid data analysis. An event is a time-stamped text message, either a syslog message sent by a host or a line appended to a log file. Logger receives and stores events; supports search, retrieval, and reporting; and can optionally forward selected events.

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.

This chapter presents an overview of Logger's capabilities, with references to other parts of this document for more detail.

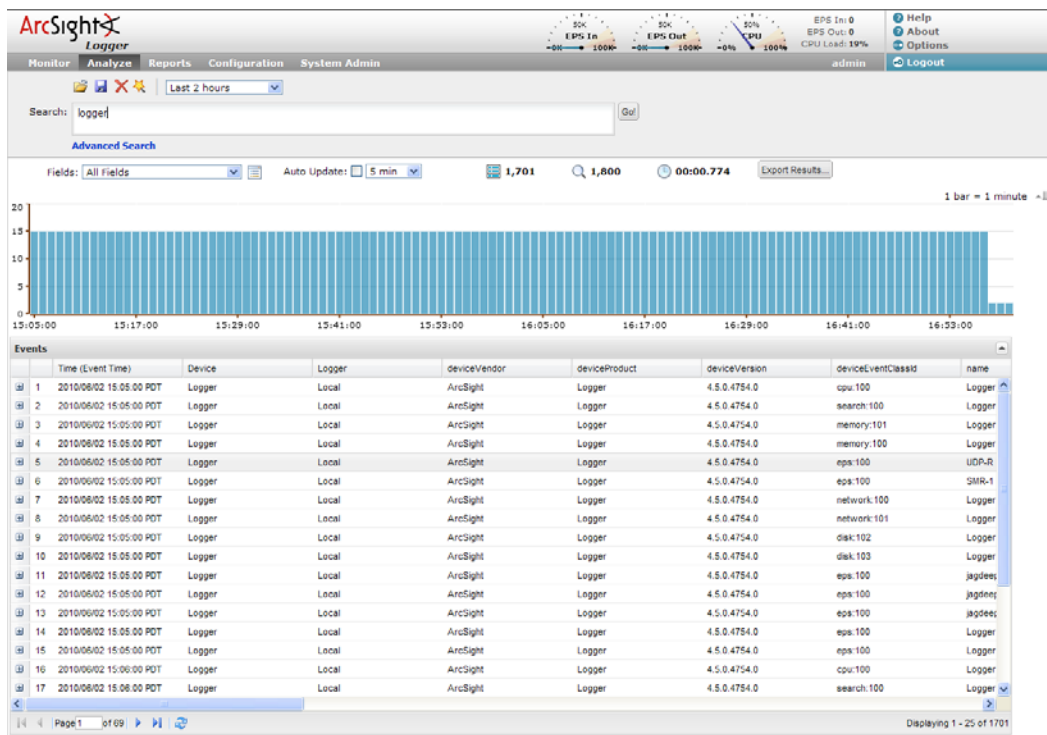


Figure 1-1 Logger web interface, Analyze tab

Logger stores time-stamped text messages, called events, at high sustained input rates. Logger compresses raw data, but can always retrieve unmodified data on demand, for forensics-quality litigation data.

Similar to ArcSight ESM, Logger leverages the ArcSight SmartConnector framework to collect events. Similar to ArcSight ESM, Logger can receive normalized CEF events from the SmartConnectors. The file-type Receivers configured on Logger only parse event time from an event.

Multiple Loggers can work together to scale up to support extremely high event volume. Loggers can be configured as a peer network, with search queries distributed across all peer Loggers.

Syslog is a loose standard (characterized, not defined, in RFC 3164) for event messages. 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. (See [Appendix A, Common Event Format, on page 373](#) for more information.)

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.

- [“Peer Loggers” on page 231](#)
- [“Common Event Format” on page 373](#)

Logger Features

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

Storage Configuration

Logger appliance includes onboard storage for events. Some Logger models include RAID 1 or RAID 5 storage systems. (See the *Appliance Specifications* document for more details. This document is available on the ArcSight Customer Support web site at <https://support.arcsight.com>.) 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, as shown in [Figure 1-2](#). Using a Network File System (NFS) as primary storage for events on a Logger appliance is not recommended.

On the **software version of Logger**, you need to have at least the minimum disk space mentioned in [“Supported Platforms and Browsers” on page 21](#) to store events. The disk space needs to be on the partition where the `/opt` directory exists. Specifically, most of this space should be available for `/opt/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 `/opt/data/logger` 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 can not 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. Storage Groups must be created when Logger is first configured. New Storage Groups cannot be added later, however, a Storage Group's size can be increased or decreased, and the retention policy defined for it can be changed.

- [“Planning” on page 22](#)
- [“Initialization Sequence \(for all Loggers\)” on page 24](#)
- [“Storage” on page 187](#)
- [“Storage” on page 263](#)

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 and the software version of Logger, a SAN.

Logger may also receive events from an ESM 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 193](#)
- [“Installing SmartConnectors to Send Events to Logger” on page 29](#)
- [“Sending Events from ArcSight ESM to Logger” on page 31](#)

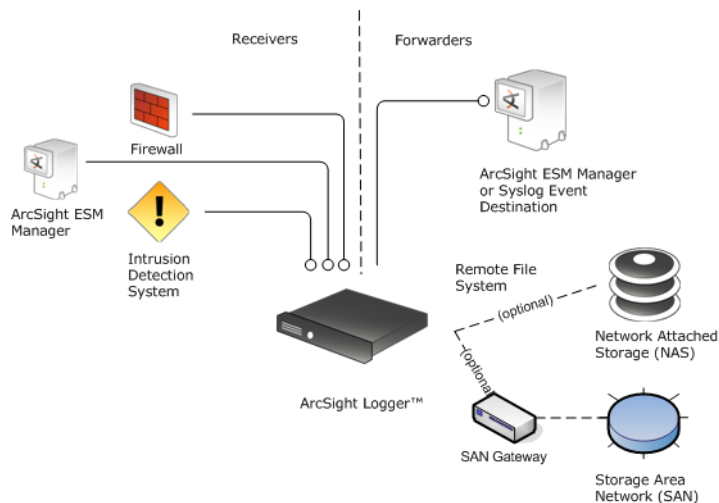


Figure 1-2 Logger appliance has multiple network interface cards (NICs) so that it can receive events on one subnet and forward events on a different subnet.

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 72](#)
- [“Saving Queries \(Saved Filters and Searches\)” on page 84](#)
- [“Filters” on page 222](#)
- [“Saved Searches” on page 225](#)
- [“Peer Loggers” on page 231](#)

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 181](#)
- [“Storage Rules” on page 189](#)
- [“Searching Peer Loggers \(Distributed Search\)” on page 73](#)

Exporting

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 version of the Logger** can be only exported locally to the Logger (to the `/opt/data/logger` directory) or to the browser from which you connect to the Logger. The `/opt/data/logger` directory can be mounted to an NFS, CIFS, or a SAN LUN.

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 and CEF events, 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 77](#)
- [“Impact of Daylight Savings Time Change on Logger Operations” on page 258](#)
- [“Scheduled Saved Search” on page 226](#)

Forwarder Configuration

Logger can send events (as they are received or past events) to other hosts using UDP or TCP, to an ArcSight Logger Streaming SmartConnector, or to an ArcSight ESM 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 ESM Manager using a syslog SmartConnector, but Logger can also send CEF events directly to a 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 ESM Manager, as shown in [Figure 1-3](#).

- ["Forwarders" on page 199](#)
- ["ESM Destinations" on page 203](#)

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" on page 289](#)
- ["Change Password" on page 290](#)
- ["User Groups" on page 284](#)
- ["Search Group Filters" on page 223](#)

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.

Logger **appliance** can be rebooted using controls in the browser user interface. To reboot the system on which the **software version of Logger** is installed, follow the instructions in the documentation for that system's operating system.

Various other system settings can be modified. Some require a system reboot for the changes to take effect.

- ["Configuration Backup and Restore" on page 235](#)
- ["Retrieve Logs" on page 246](#)
- ["Storage" on page 263](#)
- ["Reboot" on page 252](#)
- ["License & Update" on page 259](#)
- ["Network" on page 254](#)

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 ESM 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 ESM for real-time monitoring and correlation, as shown in [Figure 1-3](#). Logger can store the raw firewall data for compliance or service level agreement purposes.



In the following illustrations ArcSight Logger can be the Logger appliance or the software version of Logger that is installed on a supported platform of your choice.

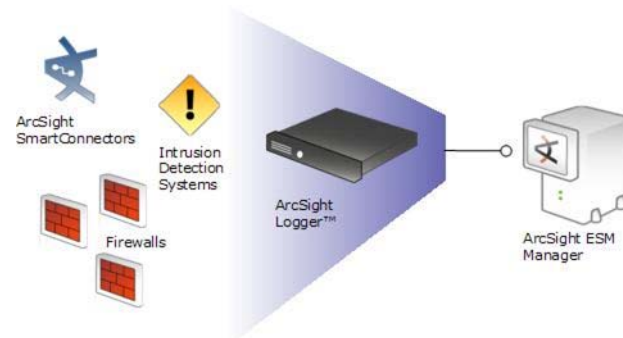


Figure 1-3 Logger can act as a funnel, forwarding selected events to ArcSight ESM.

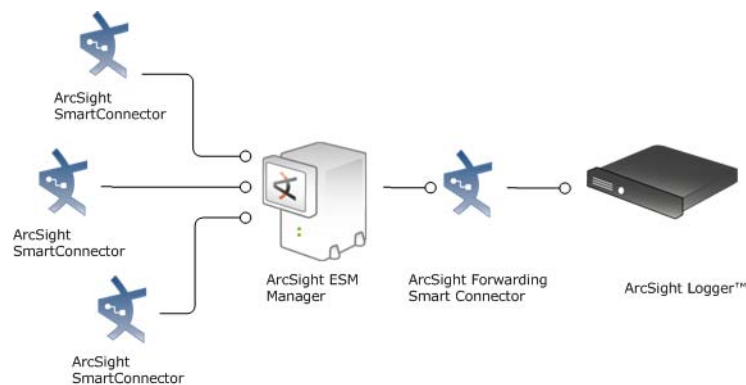


Figure 1-4 Logger can store events sent by ArcSight ESM.

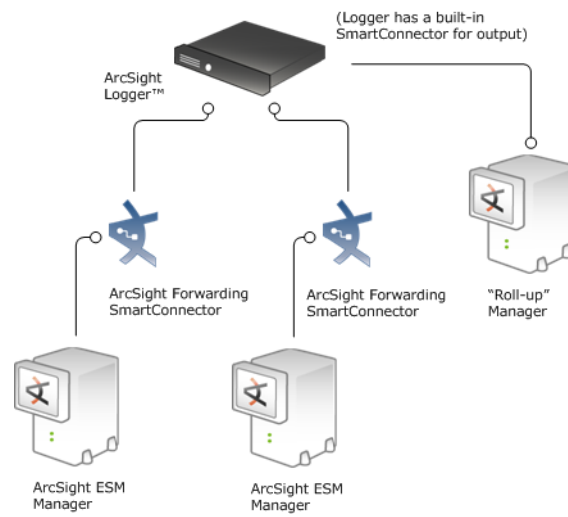


Figure 1-5 Logger can store and forward filtered events in a hierarchical ArcSight Manager deployment.

What's New in Version 4.5 GA

Logger v4.5 GA is the first Logger release that is available for two form factors: the Logger appliance and a supported platform of your choice. Although the features available in both form factors are identical, licensing requirements for both are unique and are described in this section. In addition, this section lists the new features and enhancements introduced in Logger v4.5 GA. Also refer to the Release Notes for this release for late-breaking information.

Software Version of Logger

This release introduces the first Logger in a software form factor. You can download and install this software version of Logger on a supported platform of your choice. You can also install the software on a VM image of a supported platform. You need a valid license file to install and use the software version of Logger.

The Logger software is available for download from the ArcSight Customer Support web site at <https://support.arcsight.com>. You can also obtain a valid license, which required to use the product, from ArcSight Customer Support. For details about supported platforms and installation instructions, see [“Section 2: Installing the Software Version of Logger” on page 19](#).

License Enforcement

Both form factors of Logger require valid license files. A license file enforces validity and limits on the time, maximum amount of data stored per day, and aggregated storage used on the system. On the software version of Logger, limits depend on the type of license you purchased from ArcSight.

For a detailed description of how licensing works on the software version of Logger, see [“How Licensing Works on the Software Version of Logger” on page 19](#).

Saved Search Alert

In addition to real-time alerts, Saved Search Alerts are now available on Logger. The Saved Search alerts are saved queries that run on a preconfigured schedule and send alerts when a specified number of matches occur within the specified threshold. Alerts can be sent to preconfigured e-mail, SNMP, ESM, or syslog destinations.

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.

For more information about Saved Search Alerts, see [“Alerts” on page 207](#).

Storage Volume Increase

You can extend the storage volume size you established during Logger initialization at any time. Once extended, the volume size cannot be reduced.

For more information, see [“Storage Volume Size Increase” on page 244](#).

Search Operators and Regex Helper

The significant usability and functionality improvements made in Logger v4.0 for searching events are taken a step further in this release. A flow-based search language that allows you to specify multiple search commands in a pipeline format, including regular expressions, has been introduced in this release. The language includes several operators that enable you to extract data of interest from matching queries, process it, and (optionally) create charts and reports from it.

The following pipeline operators have been introduced in this release:

- chart
- eval
- fields
- head
- rare
- regex*
- rex
- sort
- tail
- top
- where

* not a new operator in this release, but follows a new syntax.

For more information about search operators, see [“Search Operators” on page 47](#)

To ease the task of creating regular expressions for the `rex` operator, a Regex Helper tool is available in this release. The tool makes inserting rex expressions in a search query efficient and error free. The tool 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.

For more information about the Regex Helper tool, see [“Regex Helper Tool” on page 70](#).

For more information about searching and analyzing events, see [“Searching for Events on Logger” on page 72](#).

Histogram for Search Results

A histogram provides a graphical representation of the distribution of 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. Histogram enables you to randomly drill-down to events in a specific time period. For example, to investigate a spike in failed logins during a particular time, or a drop in the number of TCP connections to a server within a time period.

Histograms are automatically generated for search queries run through the Search page. Scheduled searches do not output a histogram.



Note

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 either 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 adjust the query to 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.

For detailed information about histograms, see [“Guidelines for Using the Histogram” on page 75](#).

System Content Update

This release includes system content applicable for IT operations and application development environments. A number of predefined filters for commonly searched event types are available. For example, Net-DHCP Lease Events, Net-Port Links Up and Down, bandwidth utilization, failed logins, Unix-Password Changes.

You can use these filters to quickly find events of interest without defining queries first. You can also use these filters as a starting point for creating customized filters for your environment.

For a complete list of available filters, see [“System Filters/Predefined Filters” on page 86](#).

Exporting Search Results

Search results can now be also exported in PDF form. The PDF format is 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 and CEF events, can be included in the exported report.

For more information, see [“Exporting Search Results” on page 77](#).

New Internal Event

The following new internal event for the amount of storage space used by a storage group has been added.

Device Event Category: /Monitor/StorageGroup/Space/Used

Device Event Class ID: storagegroup:100

Peer Search Update

Peer Loggers can run different versions and peers can be configured on different form factors. However, these are the only supported paths for running a search across peers:

- A search from a v4.0.x Logger to v4.5
- A search from v4.5 Logger to v4.0.x
- A search from v4.5 Logger to v4.5

Search operators (such as cef, chart, top, and so on) cannot be used for searches across peer Loggers.

For more information about peer searching, see [“Searching Peer Loggers \(Distributed Search\)” on page 73](#).

Number of Reports that can Run Concurrently

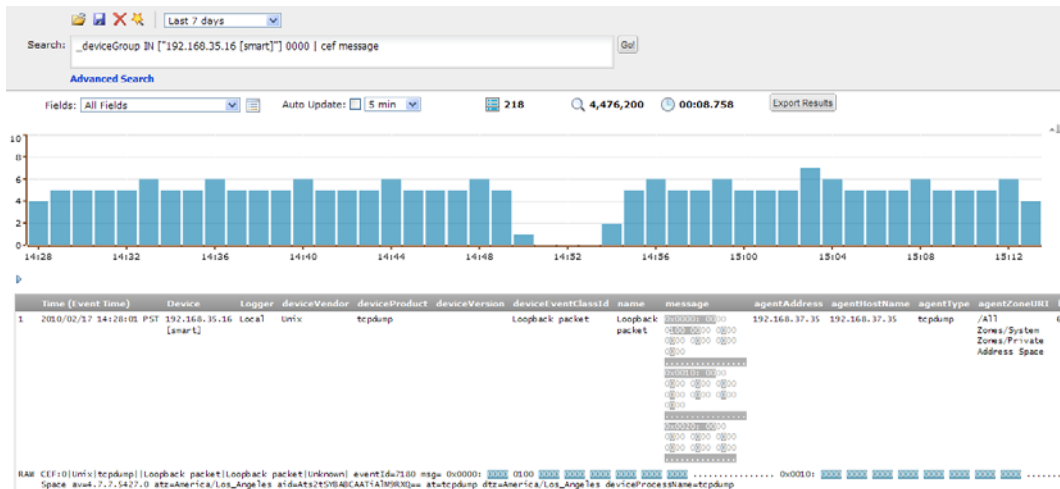
In earlier versions of Logger, you could run up to two report concurrently. Starting with this release, you can run up to five reports on a Logger.

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 0000 .....
```

Logger v4.x 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.



Enhancements to Search Results Display

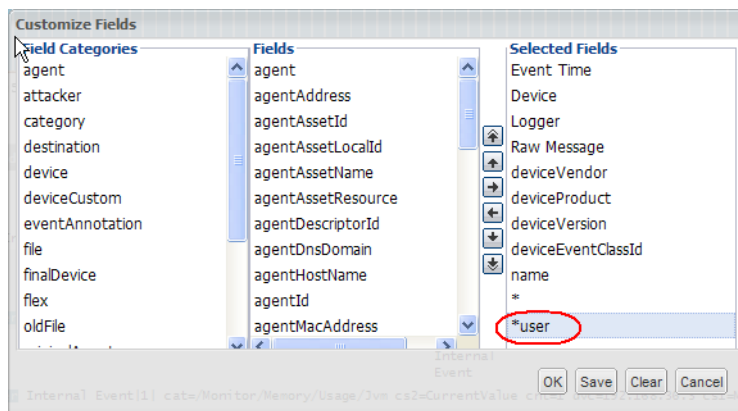
The following enhancements have been made to the search results display page:

- A well layed out, columnar, grid-like search results display that is easier to read.
- Ability to adjust the widths of displayed columns to suit your needs.
- Ability to display the raw form of a single event or all events.
- Ability to skip to the last page or to a specific page number in the search results display.

For more information about these enhancements, see [“Understanding the Search Results Display” on page 74](#).

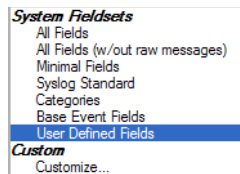
User-Defined Field Set

When you use a search operator that defines a new field, such as cef, rex, or eval, a new column for each field is added to the currently selected display. These newly defined fields are displayed by default. A new field, *user (shown below), in field sets controls the display of fields defined by search operators. When *user is included in the Selected Fields list of a custom field set, the newly defined fields are displayed.



A new field set, User-Defined Fields, is also available in this release that enables you to view only the newly defined fields.

The “User-Defined Fields” field set is available as a drop-down option from the “Fields:” menu on the page where search results are displayed.



For detailed information, see [“Field Set” on page 56](#).

Documentation for Logger

Documentation for Loggers in both form factors is available as online Help and a PDF formatted Administrator's Guide. Both forms of documentation are integrated in the product.

The Logger Administrator's Guide now contains information pertinent to both form factors of Logger. Whenever an option or a field is handled differently on the appliance than the software version of Logger, the document explains the action for both form factors.

Example: 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 version of the Logger** can be only exported locally to the Logger (to the `/opt/data/logger` directory) or to the browser from which you connect to the Logger.

Whenever there are significant differences between the two form factors, the information has been divided into sections, with one section containing information about the appliance and the other section containing information about the software form factor. For example, the System Administration chapter is divided in two sections—Section 1 for the options available on the appliance, and Section 2 for options available on the software version of Logger.

Chapter 2

Installation and Initialization

This chapter describes how to install and initialize a Logger appliance and the software version of Logger. The installation process is specific to Logger type, therefore, the installation instructions are provided in two sections:

- [“Section 1: Installing the Logger Appliance” on page 15](#)
- [“Section 2: Installing the Software Version of Logger” on page 19](#)

This chapter also includes the following information, which is applicable to both Logger types:

- [“Initialization for all Loggers” on page 22](#)
- [“Installing SmartConnectors to Send Events to Logger” on page 29](#)
- [“Sending Events from ArcSight ESM to Logger” on page 31](#)

Section 1: Installing the Logger Appliance

Installing the Logger appliance includes these steps:

- 1 Installing the device, as described in [“Installing the Logger Appliance” on page 15](#).
- 2 Setting an IP address, as described in [“Setting an IP Address on the Appliance” on page 16](#)
- 3 Initializing the appliance, as described in [“Initialization for all Loggers” on page 22](#).

Installing the Logger Appliance

To install the Logger appliance, follow the instructions in the rack installation instructions, included in the package.

ArcSight Logger Package Contents

Inspect the shipping container for signs of damage or missing items. Different appliance models contain some or all of the following:

- Logger appliance chassis (main system)
- Face plate (bezel) if not attached
- One or two power cables (North America)
- Slide rail/rack mount parts kit
- Packaging Checklist, *ArcSight Getting Started Guide* document, rack installation instructions for your platform.

If any items are missing, or there is physical damage, contact:

ArcSight Customer Support
1-866-535-3285 (North America)
+44 (0)870 141 7487 (EMEA)
E-mail: support@arcsight.com

Safety Precautions

There are a few safety concerns with any electrical appliance. Please review and observe all cautions described in the *Platform Installation Guide*, included with the appliance.

Do not remove the top cover of the Logger appliance. Opening the appliance will void the warranty, and there is generally no reason for opening the appliance, which carries the risk of electrostatic discharge or even electrocution.



Power supplies used in the Logger appliance may produce high voltages and energy hazards, which can cause bodily harm. Unless you are instructed otherwise by ArcSight, only trained service technicians are authorized to remove the covers and access any of the components inside the Logger appliance.

Do not operate Logger if the power cables are damaged, if liquids or foreign objects have entered the appliance, or if the appliance has been damaged by dropping or other physical shock, or if the device has been exposed to water.

Do not operate Logger in a wet environment. Do not modify power cables or plugs. Consult a licensed electrician or your power company if site modifications are necessary. Always follow national or local electrical wiring regulations.

When connecting or disconnecting power to hot-swappable power supplies, observe these guidelines:

- Install the power supply before connecting the power cable to the power supply.
- Unplug the power cable before removing the power supply.
- Disconnect power to Logger by unplugging **both** power cables from the power supplies.

Setting an IP Address on the Appliance

Before logging in to a Logger appliance for the first time, you need to configure at least one valid IP address. There are three ways to accomplish this:

- Attach a terminal to the serial port on Logger and use the Command Line Interface to change the default IP addresses; or
- Attach a monitor and keyboard to the rear panel connectors and use the Command Line Interface to change the default IP addresses; or
- Configure a host to be a subnet that matches the predefined Logger IP (192.168.35.*) and use a browser from that host to log in and change the default IP addresses.

Connecting to 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.

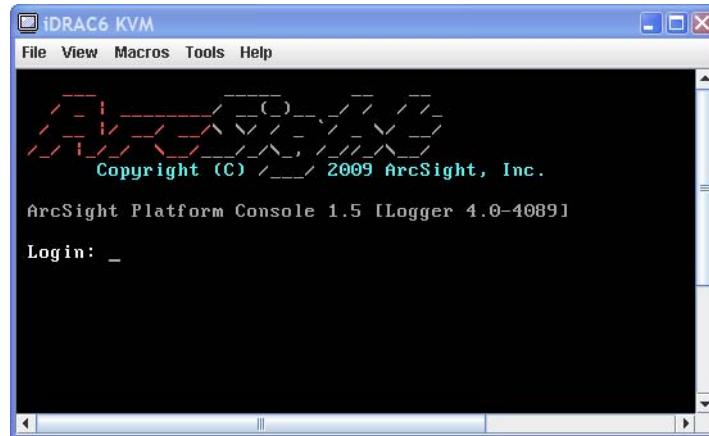


Figure 2-1 The ArcSight Platform Console sign-on. Note that the platform version will not match the current Logger application version.

To set an IP address using the CLI:

- 1 Connect a terminal to the serial port on Logger. Turn on Logger. Enter user name and password (initially, admin/password). CLI credentials are the same as for the web user interface. The terminal should display the ArcSight Platform Console prompt shown in Figure 2-1.

For security reasons, change the default password after the initialization is complete.

- 2 Enter `set password <pwd>` (replace <pwd> with your chosen password) to set the current user's password.
- 3 Enter `set ip eth0 192.168.35.35/<prefix>`, replacing the IP address with the IP address desired and <prefix> with the number of bits in the subnet mask. (For example, /24 = 255.255.255.0.)
- 4 Enter `set hostname <logger>`, replacing <logger> with the fully-qualified domain name (FQDN) of the desired host.
- 5 Enter `set dns <search_domain> <name_server>`, replacing the <search_domain> with your domain and <name_server> with the hostname or IP address of your nameserver.
- 6 Enter `set defaultgw 192.168.35.2`, replacing the IP address with your default gateway IP address.
- 7 The preceding changes take effect immediately. To confirm that the settings are correct for your environment, enter `show config`.

Using a Browser to Set an IP Address

- 1 Open a modern, Flash-enabled browser (Microsoft Internet Explorer 6.0, 7.0 or Firefox 1.5 or later). Specify Logger's default IP address, like this:

<https://192.168.35.35/>

- 2 At the login screen, enter `admin` for user name and `password` for password. Logger reminds you to set up a Storage Volume, but that step is described later, in

[“Initialization Sequence \(for all Loggers\)” on page 24](#). It is very important that you do not specify a Storage Volume or make other critical deployment decisions at this time.

- 3 Click the **System Admin** tab.
- 4 On the sub-menu, click **Platform** (under Settings).
- 5 Click the **Network** tab and enter the desired host name, default gateway, IP address(es) and other information. Click **Update Settings**.



It is important that the host name is resolvable by DNS and that it resolves to the Logger's IP address. Performance is significantly affected if DNS cannot resolve the host name.

- 6 Click the **Change Password** sub-menu (under User/Groups). Enter the old password ('password'), enter a new password and confirm it. Click **Set Password**.
- 7 On the sub-menu (under System Configuration), click **System Reboot**. Click **Start Reboot Now**. The setting changes take effect after Logger is rebooted.

Other CLI Commands

The following commands are available at the CLI prompt:

Command	Description
exit	Logout
halt	Stop and power down the Logger appliance
reboot	Reboot the Logger appliance
set defaultgw <IP> [nic]	Set the default gateway for one or all network interfaces
set dns <dn1> [, <dn2>][, <dn3>] ns1 [, ns2]	Set DNS name server(s). dn=search domain name, ns=nameserver
set hostname <host>	Set Logger's host name
set ip <nic> <IP> [/prefix] [netmask]	Set Logger's IP address for a specific network interface.
set password	Set the password the current user's account.
show admin	Show the default administrator user's name
show config	Show host name, IP address, DNS, and default gateway for this Logger
show defaultgw [nic]	Display the default gateway for all or the specified network interface
show dns	Display the DNS name servers currently configured
show hostname	Display the current hostname
show ip [nic]	Show the IP addresses of all or the specified network interface

Command	Description
enable support	Enable access by ArcSight Customer Support for one session. When Customer Support logs out, access is automatically disabled. Support access is disabled by default. Once this command is given, the disable support command can rescind it.

Initializing the Logger Appliance

See [“Initialization for all Loggers” on page 22](#).

Configure DRAC for Remotely Accessing a Logger Appliance

ArcSight recommends configuring Dell Remote Access Controller (DRAC) on a Logger appliance. Doing so ensures that you (and ArcSight Customer Support, with your permission and assistance) can remotely access your appliance's console and control its power. To enable DRAC, you need to configure a network interface dedicated for this purpose at the back side of the appliance. For details about configuring DRAC on your appliance, contact ArcSight Customer Support.

Section 2: Installing the Software Version of Logger

The software version of Logger is available for download from the ArcSight Customer Support web site at <https://support.arcsight.com>. You need to have a server with supported operating system and storage available to install the software Logger. A valid license is required to use the product, which is also available from ArcSight Customer Support.

This section describes how licensing works on software version of Loggers, installing such a Logger for the first time, updating an expired license, and uninstalling the Logger.

How Licensing Works on the Software Version of Logger

A license for the software version of Logger defines the validity and limits for the following:

- Time limit: The length of time for which the license is valid. For example, 60 days.
- Data limit: A per day limit on the amount of incoming data. For example, 20 GB per day. The sum of the size of the original events is used to determine this value.

Even if this limit is exceeded, the software version of Logger continues to collect and store events; therefore, no events are lost. However, if this limit is exceeded 6 times (that is, any 6 days) in a 30-day sliding window, you cannot search or run reports on the collected events until the 30-day sliding window contains 5 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 5 violations so far, you can search and report on the stored event data on January 21st. However, if there is another violation on January 30th, you cannot 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

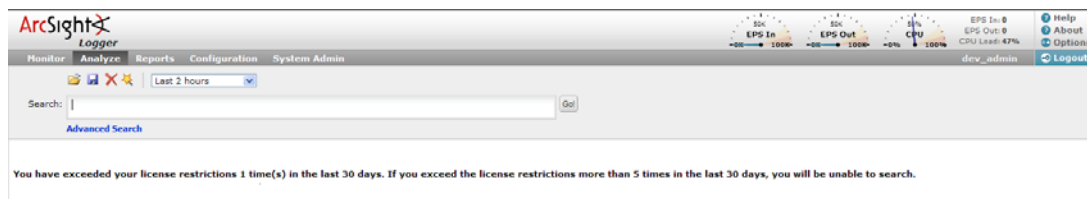
ability to search resumes on February 5th because the January 5th violation is now outside of the 30-day sliding window.



If you are using ArcSight connectors to send events to the software version of Logger, make sure you are running connector version 5.0.0.5560.0-0 on your connectors to ensure that event size is accurately accounted on the Logger.

- **Aggregated storage limit:** A limit on the aggregated storage—the sum of storage used to store incoming events and the storage consumed due to retention—used on the Logger. For example 80 GB.

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



You can also view the data limit violation information on the License Information page (**Configuration > License Information**). The License Information 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.

Date	Data Stored	Limit Exceeded
Mon Jan 04 00:00:00 PST 2010	0	false
Tue Jan 05 00:00:00 PST 2010	0	false
Wed Jan 06 00:00:00 PST 2010	0	false
Thu Jan 07 00:00:00 PST 2010	0	false
Fri Jan 08 00:00:00 PST 2010	0	false
Sat Jan 09 00:00:00 PST 2010	0	false
Sun Jan 10 00:00:00 PST 2010	0	false
Mon Jan 11 00:00:00 PST 2010	0	false
Tue Jan 12 00:00:00 PST 2010	0	false
Wed Jan 13 00:00:00 PST 2010	0	false
Thu Jan 14 00:00:00 PST 2010	0	false
Fri Jan 15 00:00:00 PST 2010	0	false
Sat Jan 16 00:00:00 PST 2010	0	false
Sun Jan 17 00:00:00 PST 2010	0	false
Mon Jan 18 00:00:00 PST 2010	0	false
Tue Jan 19 00:00:00 PST 2010	0	false
Wed Jan 20 00:00:00 PST 2010	0	false
Thu Jan 21 00:00:00 PST 2010	0	false
Fri Jan 22 00:00:00 PST 2010	0	false
Sat Jan 23 00:00:00 PST 2010	0	false
Sun Jan 24 00:00:00 PST 2010	0	false
Mon Jan 25 00:00:00 PST 2010	0	false
Tue Jan 26 00:00:00 PST 2010	0	false
Wed Jan 27 00:00:00 PST 2010	0	false
Thu Jan 28 00:00:00 PST 2010	0	false
Fri Jan 29 00:00:00 PST 2010	0	false
Sat Jan 30 00:00:00 PST 2010	0	false
Sun Jan 31 00:00:00 PST 2010	0	false
Mon Feb 01 00:00:00 PST 2010	0	false
Tue Feb 02 00:00:00 PST 2010	33844	true



If you exceed the data limit frequently, you should consider purchasing a license from ArcSight that suits your needs. Please contact your ArcSight sales representative to purchase a license. Once you obtain a new license, follow the instructions in [“Applying a New License on the Software Version of Logger”](#) on page 22 to apply the new license on your Logger.

Supported Platforms and Browsers

You can install the software version of Logger on a platform with the following specifications. For a detailed capacity planning guide, see the *Capacity Planning for Software Version of Logger* document that is available for download from the ArcSight Customer Support site at <https://support.arcsight.com>.

Specification	Details
Certified Operating Systems	<ul style="list-style-type: none"> Red Hat Enterprise Linux (RHEL), version 5.4, 64-bit CentOS, version 5.4, 64-bit <p>NOTE: A VM installation of the above listed operating systems is supported.</p>
Other Supported Operating Systems	<ul style="list-style-type: none"> Oracle Enterprise Linux (OEL) 5.4 Red Hat Enterprise Linux (RHEL), version 4.x, 64-bit CentOS, version 4.x, 64-bit <p>NOTE: A VM installation of the above listed operating systems is supported.</p>
CPU, Memory, Disk Space	<p>For Small to Medium Deployments</p> <ul style="list-style-type: none"> CPU: 1 or 2 x Intel Xeon Quad Core or equivalent Memory: 4 - 12 GB (12 GB is recommended) Disk Space: 100 - 120 GB (120 GB is recommended) <p>For Medium to Large Deployments</p> <ul style="list-style-type: none"> CPU: 2 x Intel Xeon Quad Core or equivalent Memory: 12 - 24 GB (24 GB is recommended) Disk Space: 120 - 400 GB (400 GB is recommended) <p>NOTES:</p> <ul style="list-style-type: none"> The disk space needs to be on the partition where the <code>/opt</code> directory exists. Specifically, most of this space should be available for <code>/opt/data/logger</code> directory. SAN can be used for storing events, however, this LUN must be mapped to the <code>/opt/data/logger</code> directory on the system on which the software version of Logger is installed. Using NFS as primary storage for events on the software version of Logger is not recommended.
Browsers	<ul style="list-style-type: none"> Internet Explorer: Version 8 Firefox: Versions 3.0 and 3.5 <p>An Adobe Flash Player plug-in is required on these browsers for some of the features, such as Histogram and charts, to work.</p>

Specification	Details
Other Applications	For optimal performance, make sure no other applications are running on the system on which you install the software version of Logger.

Installing the Software Version of Logger for the First Time

See the *Logger v4.5 Release Notes* for installation instructions. The release notes are available from the ArcSight Customer Support web site at <https://support.arcsight.com>.

Initializing the Software Version of Logger

See [“Initialization for all Loggers” on page 22](#)

Applying a New License on the Software Version of Logger

To apply a new license on the software version of Logger:

- 1 Unzip the license file that you obtained from ArcSight Customer Support.
- 2 Copy the `arcsight_license` file in the `/etc` directory of the system where the software version of Logger is installed.

Uninstalling the Software Version of Logger

To uninstall the software version of Logger:

- 1 Enter this command in the directory where you installed the software version of Logger:

```
./logger_<build_number>.uninstall
```

Where `<build_number>` is the four digit number that is part of the installation file you used to install the Logger. For example, 4260.

- 2 Reboot your machine.

Initialization for all Loggers

Logger initialization requires planning because there are several initial settings which cannot be changed once they are set.

Planning

Storage Strategy

Logger events can be stored in these ways:

- Locally (on **Logger appliance** and **software version of Logger**)
- Remotely on a Storage Area Network (SAN) **on Logger appliance models** that support SAN. SAN should be available before you bring the Logger online. Only one LUN can be used to store events.

SAN can also be used for storing events on the **software version of Logger**, however only one LUN can be used for storing events, and this LUN must be mapped to the `/opt/data/logger` directory on the system on which the software version of Logger is installed.

Using a Network File System (NFS) as primary storage for events on a Logger is not recommended.

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.

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

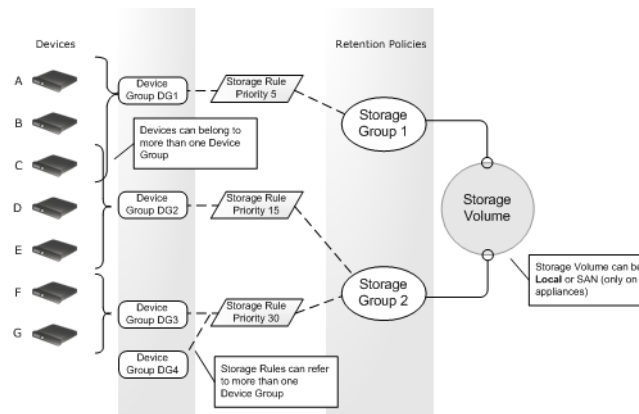


Figure 2-2 Devices participate in Retention Policy

Figure 2-2 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.



An implicit Storage Rule, with lowest priority, maps all Devices to the Default Storage Group.

Peer Loggers

Peer Loggers can be added or removed dynamically.

Initialization Sequence (for all Loggers)

It is very important that you initialize Logger in the sequence shown here. Several of the settings described here cannot be changed once set, therefore, make sure you perform the the initialization steps carefully.



One-time initialization on a Logger **appliance** can only be changed by performing a factory reset (see [Appendix D, Restoring Factory Settings, on page 393](#)). Be sure you know how you want Logger storage set up before performing the first steps of the initialization sequence (up to rebooting).

On the **software version of Logger**, you can uninstall and reinstall the software to restore it to its pre-initialization state.

The following sequence ensures that resources are created and parameters are set in the proper order.

For a Logger appliance follow all of these steps.

For the software version of Logger, skip Steps 1 and 5, and do not perform preallocation in Step 3, Storage Volume.

- 1** License
- 2** SAN (on selected Logger appliance models and the software version of Logger)
- 3** Storage Volume - establish where Logger stores event data
- 4** Storage Groups - apply retention policies to the Storage Volume
- 5** Time Settings
- 6** Index Fields and Full-text Indexing
- 7** Reboot - commit the changes made in previous steps
- 8** Receivers
- 9** Devices
- 10** Device Groups
- 11** Storage Rules

1 License

Skip this step if you are initializing the software version of Logger because the license file for that Logger is applied during the installation of the Logger.

Download valid license files for **all** your **Logger appliances** from your customer directory on the ArcSight Customer Support web site. Then, follow instructions in this section to apply the license file.

If a valid license file is not present on the Logger appliance, only the platform configuration user interface (the System Admin functions) will be available on it. You will not be able to use any of the Logger application functionality.

Please note the following:

- There is no additional charge for the license files.

- A license file contains the serial number of the appliance for which it was generated. Therefore, you need a separate license file for each of your Logger appliances.

If you have multiple Logger appliances, make sure that you install a license file that corresponds to the serial number of an appliance on it. To determine a corresponding license file for a Logger, match the serial number in the license file's name to the serial number on your Logger appliance.

- **Do not** rename the license file you download. A license file with an altered name will not install on Logger.

To apply a license file on a Logger:

- a Download the license file from the ArcSight software download site at <https://software.arcsight.com> to a computer from which you can connect to Logger.



Do not rename the license file you download. A license file with an altered name will not install on Logger.

Caution

- b From the computer to which you downloaded the update file, log in to the Logger's browser-based interface using an account with administrator (upgrade) privileges.
- c Click the **System Admin** tab > **System Update**.
- d Browse to the *license* file you downloaded earlier and click **Upload Update**.

Wait until the user interface displays a message indicating that the upload was successful. You do not need to reboot the Logger after applying a license file.

2 SAN

Skip this step if you will use Logger's built-in storage.

If you are using a SAN as your primary storage for a Logger **appliance**, the SAN must be up before initializing the Logger. Logger can attach to only one LUN (on SAN) at a time for primary storage. (Only certain Logger models support SANs.)

On the Logger **appliance**, you can use either one of the two fiber ports available on the back panel of your SAN-enabled Logger to attach the LUN for primary storage. However, both ports cannot be used simultaneously.

On the **software version of Logger**, any storage technology (including SAN) can be used for storing events. When using SAN, only one LUN can be used for storing events. This LUN must be mapped to the `/opt/data/logger` directory on the system on which the software version of Logger is installed. The software version of Logger does not include SAN management functionality.

You can add more LUNs for event archival, configuration backup, and export. See ["SAN" on page 267](#) for instructions.

3 Storage Volume

Establish the Logger's Storage Volume. See ["Storage Volume" on page 191](#). Choose **Local** to use Logger's built-in storage on a Logger **appliance** or the **software version of Logger**. OR choose **SAN** if your Logger **appliance** model supports SAN. If you will use SAN on the Logger **appliance**, enter a folder path to the SAN. This folder path must already exist on the remote storage. You do not need to enter the folder path for the

software version of Logger because, by default, the `/opt/data/Logger` directory is used and this path cannot be changed. (If you want to use SAN for storage on the software version of Logger, its LUN must be mapped to the `/opt/data/logger` directory on the system on which the Logger software is installed.)

You can choose to pre-allocate your Storage Volume to save time later. Performance is degraded if you don't pre-allocate at least a portion of the storage volume, especially on remote volumes. ArcSight recommends 100% pre-allocation for both local and remote volumes. **Pre-allocation is not needed if you are initializing the software version of Logger.**

Even though 100% pre-allocation can take more time on remote volumes, doing so improves performance on your Logger and protects it from running out of disk space. Allocating lesser than 100% space may result in sub-optimal performance on the Logger.



Note

Storage Volume can be extended but not reduced after initialization. For more information, see [“Storage Volume Size Increase” on page 244](#).

3 Storage Groups

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**. As a result, you have five storage groups available for event storage and one for Logger's internal events.

Once the Storage Volume has been created, you must configure the Default Storage Group, which is created by default. (You cannot change the name of this group.) You are not required to create additional Storage Groups, but ArcSight recommends that you do so even if you don't need them right now because additional **Storage Groups cannot be created once Logger has been initialized. However, a Storage Group's size can be increased and decreased any time; therefore, create additional groups of minimal size even if you don't need them at this point.**

Each Storage Group can have a different retention policy.



Caution

Do not reboot Logger in the next step unless you are certain of your Storage Volume and Storage Group choices.

Maximum number of Storage Groups on Logger (including preexisting groups): 6

Storage Groups created by default: 2 (Default Storage Group and Internal Storage Group)

Number of Storage Groups available for event storage: 5

Number of Storage Groups available for Logger's internal events: 1

Number of Storage Groups you can create: 4

See [“Storage Groups” on page 187](#) for the details of adding Storage Groups.

4 Time Settings

Skip this step if you are initializing the software version of Logger.

Configure the system time manually. Follow instructions in [“Time/NTP” on page 256](#).

Optionally, configure NTP time settings. Configuring an NTP server will ensure precise time stamping of events, which is a key log management function. ArcSight strongly recommends that you use Network Time Protocol (NTP) for system time. See [“Time/NTP” on page 256](#) for more information.

5 Index Fields and Full-text Indexing

As shown in the [Figure 2-3](#), during the initialization process, Logger prompts you to add a recommended set of fields to the index. You are not required to index event fields at this point, but ArcSight strongly recommends that you do so because indexing significantly improves search and reporting performance. When you add fields to the index, search queries yield significantly faster results. You might need to add additional fields to suit your needs.

Additionally, full-text indexing is not enabled by default; to enable it, click **Enable full text indexing**. Once enabled, full-text indexing cannot be disabled. (For full-text indexing, each event is scanned and divided into keywords and stored on the Logger.) See [“Indexing” on page 80](#) for more information.

Once a field has been added, you cannot remove it or unindex it.

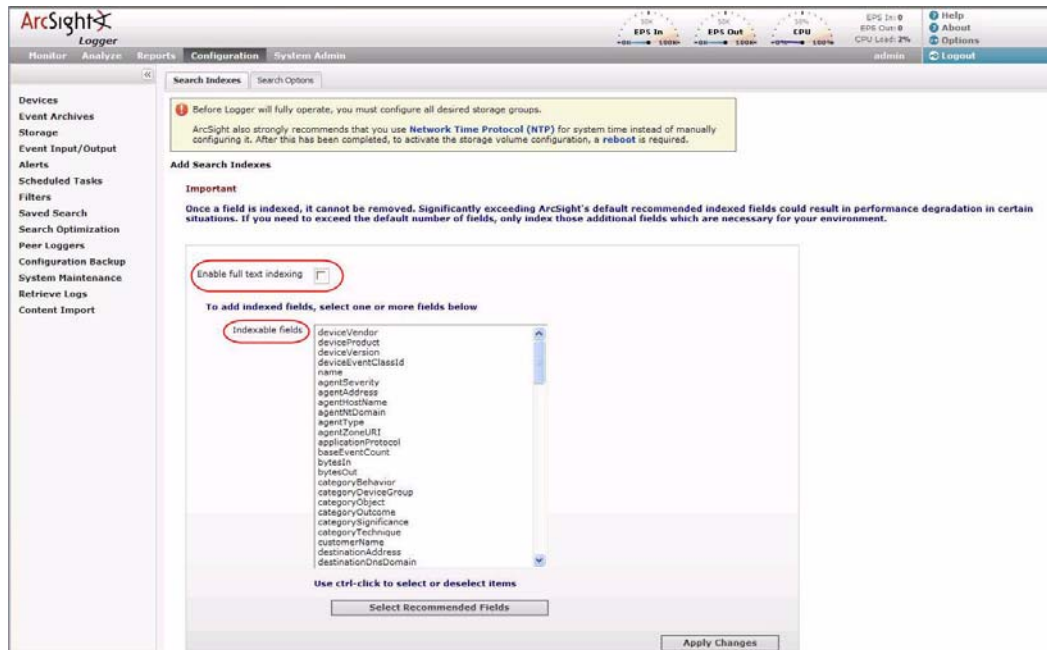


Figure 2-3 Initializing screen that prompts you to select recommended fields to index.

Click **Select Recommended Fields** to highlight the set of fields ArcSight recommends that you add to the index. Then, click **Add** to add those fields to the index.

6 Reboot

After the Storage Volume and Storage Groups have been created, reboot the system to commit changes before other resources can be created and Logger can begin processing events. See [“Reboot” on page 252](#).



When Logger is rebooted, the Storage Volume and Storage Group settings become permanent. Storage Volume size can be extended, but not reduced after initialization. And only certain settings of non-default Storage Groups can be changed. For more information, see [“Storage Volume Size Increase” on page 244](#).

7 Receivers

Now that you have established a Storage and retention policy configuration for Logger, you can create Receivers to listen for events. Unlike the previous configuration choices in this Initialization Sequence, Receivers can be changed and deleted as needed in the future. Receivers can also be disabled and re-enabled later. For more information about setting up Receivers, see [“Receivers” on page 193](#).

8 Devices

When at least one Receiver is enabled, Logger begins storing events. Using a process called auto-discovery, 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 that you expect to be sending events to Logger. You might do this if you don't want to wait for autodiscovery, or if you want to control the initial naming of each Device. (Auto-discovered Devices are named for their host, or if the DNS lookup fails, for their IP address, and their Receiver.) For information about manually creating Devices, see [“Devices” on page 182](#).

9 Device Groups

Device Groups are containers for Devices, in the same way folders (or directories) contain files. Device Groups are a way to give a name to a group of Devices. Each Device Group is associated with a particular Storage Group, which assigns the Device Group a retention policy.

Rather than just creating one Device Group for each retention policy, however, you might want to create more Device Groups as a way to categorize events. You can search for events that match a certain pattern and which belong to a particular Device Group. A given Device can be a member of several Device Groups, as well, which makes them broadly flexible.

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 183](#).

10 Storage Rules

Events are stored in the Default Storage Group unless otherwise specified. Typically, Storage Rules send events from specified Device Groups to Storage Groups other than the Default Storage Group. Therefore, Storage Rules implement your secondary and tertiary retention policy.

If you only implemented extra Storage Groups because ArcSight recommended that you do so (back in step 3), then you do not need to create any Storage Rules and you can skip this step. 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, create Storage Rules that associate the appropriate Device Groups with the Storage Groups that implement the correct retention policy. See [“Storage Rules” on page 189](#) for more information.

Storage Rules are tested in order; 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.

Installing SmartConnectors to Send Events to Logger

ArcSight Logger is a storage solution optimized for extremely high event throughput. Logger stores time-stamped text messages, called events, at high sustained input rates. Unlike ArcSight SmartConnectors, Logger does not “normalize” events. Events consist of an event time, a receipt time, a source (host name or IP address), and an un-parsed message portion. Logger compresses raw data, but can also retrieve it in an unmodified form for forensics-quality litigation reporting.

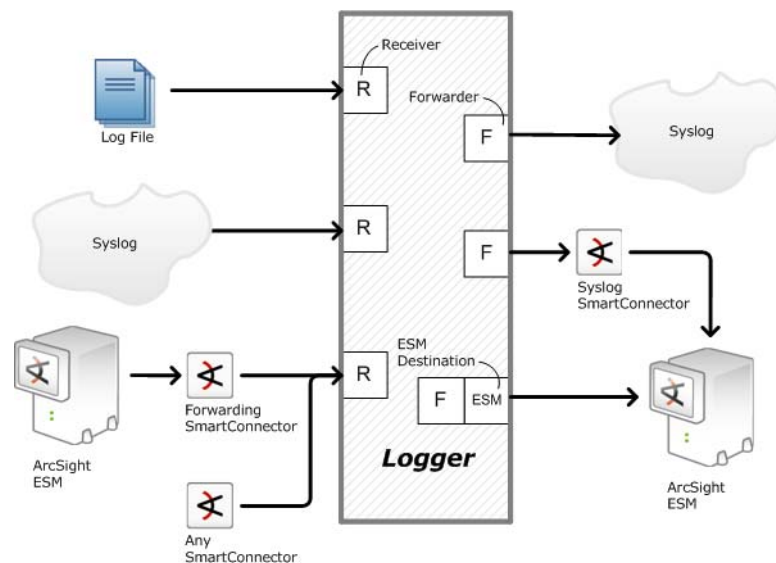


Figure 2-4 ArcSight SmartConnectors interact in a variety of ways with Logger Receivers (R) and Forwarders (F).

Multiple Loggers can work together to support an extremely high event volume. ArcSight Logger can be configured as a peer network with queries distributed across all peer Loggers.

Logger can be configured to receive and log all events from a device, and to forward filtered events on to a destination such as ArcSight ESM. Events can also be filtered by individual SmartConnectors. In such a “funnel,” a device that creates many security events (such as a firewall) might be read by a SmartConnector which filters events of interest (and optionally aggregates events, as well) and sends them to a Logger Receiver. A Logger Forwarder then passes a subset of the received events downstream to ArcSight ESM.

For more information about filtering and aggregation by SmartConnectors, see the *ArcSight SmartConnector User's Guide*.

SmartMessage

SmartMessage is an ArcSight technology that provides an efficient secure channel for Common Event Format (CEF) events between ArcSight SmartConnectors and Logger. SmartConnectors can also send CEF messages in clear to Logger using syslog protocol.

**Caution**

ArcSight recommends installing SmartConnector v4.7.5 or later. If you do not have the current build, download the latest from the ArcSight website.

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

SmartMessage provides an end-to-end encrypted secure channel using 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 secure sockets layer (SSL) protocol to send encrypted events to Logger. This is similar to, but different from, the encrypted binary protocol used between SmartConnectors and ESM Manager.

Use port 443 (instead of ArcSight's traditional port, 8443) because the secure channel uses SSL.

Set up the SmartMessage Receiver on Logger first (see [“Receivers” on page 193](#)) and then configure the SmartConnector as described below.

To configure a SmartConnector to send events to Logger

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

**Note**

Use SmartConnector release 4.7.5 or later for SmartMessages. This version is also required for connectors to connect to Logger in FIPS mode. For CEF and Syslog, older SmartConnectors will work (build 4785 or later).

- 2 Specify the required parameters. Enter the Logger hostname or IP address and the name of the SmartMessage Receiver. (For un-encrypted CEF syslog, enter the Logger hostname or IP address, the desired port, and choose UDP or TCP output.) These settings will need to match the Receiver you create in Logger to listen for events from this connector.

For more information about the Common Event Format (CEF), see [“Common Event Format” on page 373](#).

Forwarding Logger Events to an ESM Manager

Logger can forward these types of events to an ESM Manager:

- Syslog events to an ArcSight Syslog SmartConnector that is connected to an ESM Manager.
- Common Event Format (CEF) events directly to an ESM Manager using Logger ESM Destinations. An ESM Destination appears as a SmartConnector to an ESM Console.

- Events received by file receivers where the type specified is not Other. Such events are forwarded using the ArcSight Streaming SmartConnector.

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

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

- 1 Install the SmartConnector normally. Register the SmartConnector with a running ESM 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 ArcSight Customer Support.

Sending Events from ArcSight ESM to Logger

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

To configure the ArcSight Forwarding SmartConnector to send events to Logger



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.

- 1 Install the SmartConnector component normally, but cancel the installation when the SmartConnector Wizard asks whether the target Manager uses a demo certificate (see Figure 2-5). Confirm that you want to exit, then click **Done** to dismiss the Install Wizard. This will install the SmartConnector, but leave it un-configured.
- 2 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`
- 3 Start the SmartConnector configuration program again using the `$ARCSIGHT_HOME/current/bin/runagentsetup` script (or `arcsight agentsetup -w`).



Figure 2-5 When the first screen of the SmartConnector Configuration Wizard appears, asking about a demo certificate, click **Cancel**.

- 4 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

Parameter	Description
ArcSight Source Manager Host Name	IP or host name of the source ArcSight ESM Manager
ArcSight Source Manager Port	8443 (default)
ArcSight Source Manager User Name	A user account on the source Manager will 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 ESM Manager at the same time, see [“Configuring SmartConnectors to Send Events to Both Logger and an ESM Manager” on page 31](#).

For more information about the Common Event Format (CEF), see [“Common Event Format” on page 373](#).

Chapter 3

Using the User Interface

This chapter describes the user interface portion of the Logger web application. The user interface includes site navigation, and performance monitoring. This chapter includes:

Navigation: see [“Logger User Interface” on page 35](#)

Performance monitoring: see [“Monitor” on page 37](#)

The other tabs, Analyze, Reports, Configuration, and System Admin, are described in later chapters.

Logger User Interface

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.

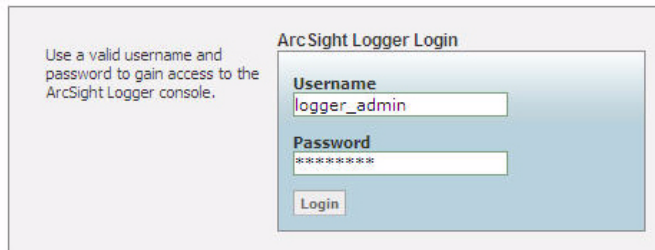


Figure 3-1 Users must login to authenticate themselves to Logger.

Browser Requirements

Logger works with most modern 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/>.

See the Release Notes document to find out the browser versions supported for this release.

Navigating the User Interface

As shown in Figure 3-2, a consistent navigation and information band runs across the top of every page in the user interface.

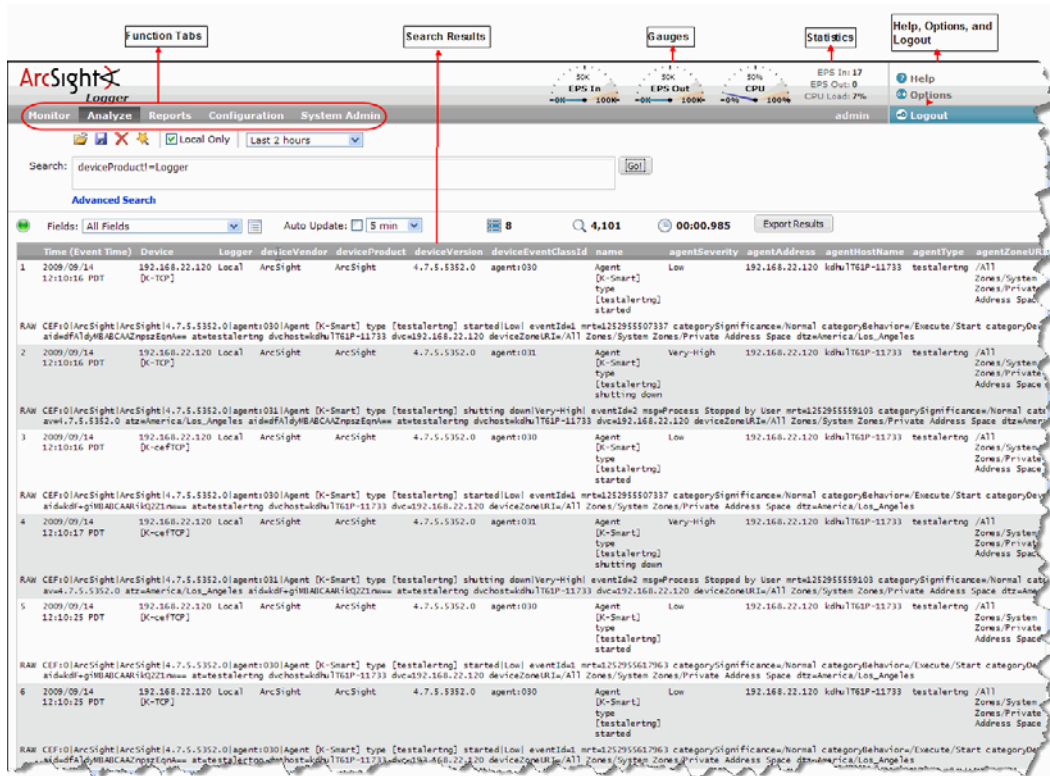


Figure 3-2 Overall layout of the Logger interface.

Gauges at the top of the screen provide an indication of the throughput and CPU usage information available in more detail on the Monitor tab. The range of the gauges can be changed on the Options page. The current logged-in user's name is shown below the statistics.



Figure 3-3 Sub-menus pull down from main function tabs.

The menu list in the upper right includes links for Help, Options, and Logout.

Help

Clicking the Help link on any page displays online help for the current page.

Options

The Options page, shown in [Figure 3-4](#), 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.

The **Default start page for all users** can be set to Monitor Summary (the default), Reports Dashboard, or Analyze to configure which tab will be displayed after a user logs in.

The screenshot shows a window titled "Options". Inside, there are three rows of configuration options, each with a text label and a dropdown menu:

- EPS input rate gauge max: 100K
- EPS output rate gauge max: 100K
- Default start page for all users: Monitor Summary

At the bottom right of the window are two buttons: "Save" and "Cancel".

Figure 3-4 Options, where you specify the range of input and output gauges.

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 279](#).



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.

Monitor

The Monitor tab, shown in [Figure 3-5](#), 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.)

Under the Monitor tab, select monitor pages for Summary, Platform, Network, Logger, Receivers, Forwarders, and Storage.

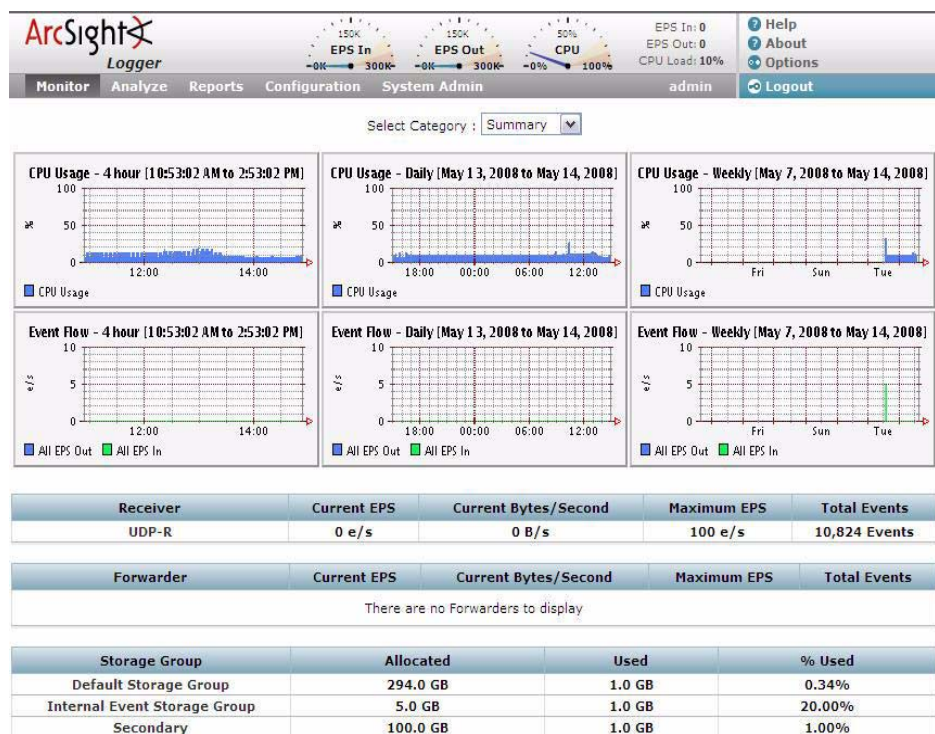


Figure 3-5 The Monitor tab displays summary information by default

Most pages include a Duration control. On these, choose a time span for historical data:

- 4-hours
- Daily
- Weekly

The Summary page displays graphs for each duration as a guide for which duration to choose.

On the Summary page, click on a Receiver, Forwarder, or Storage Group name to jump to the configuration page for that type of resource.

Platform

The Platform monitor page, as shown in [Figure 3-6](#), displays information about CPU usage, memory usage, bytes received and sent on the network, and raw disk reads and writes.

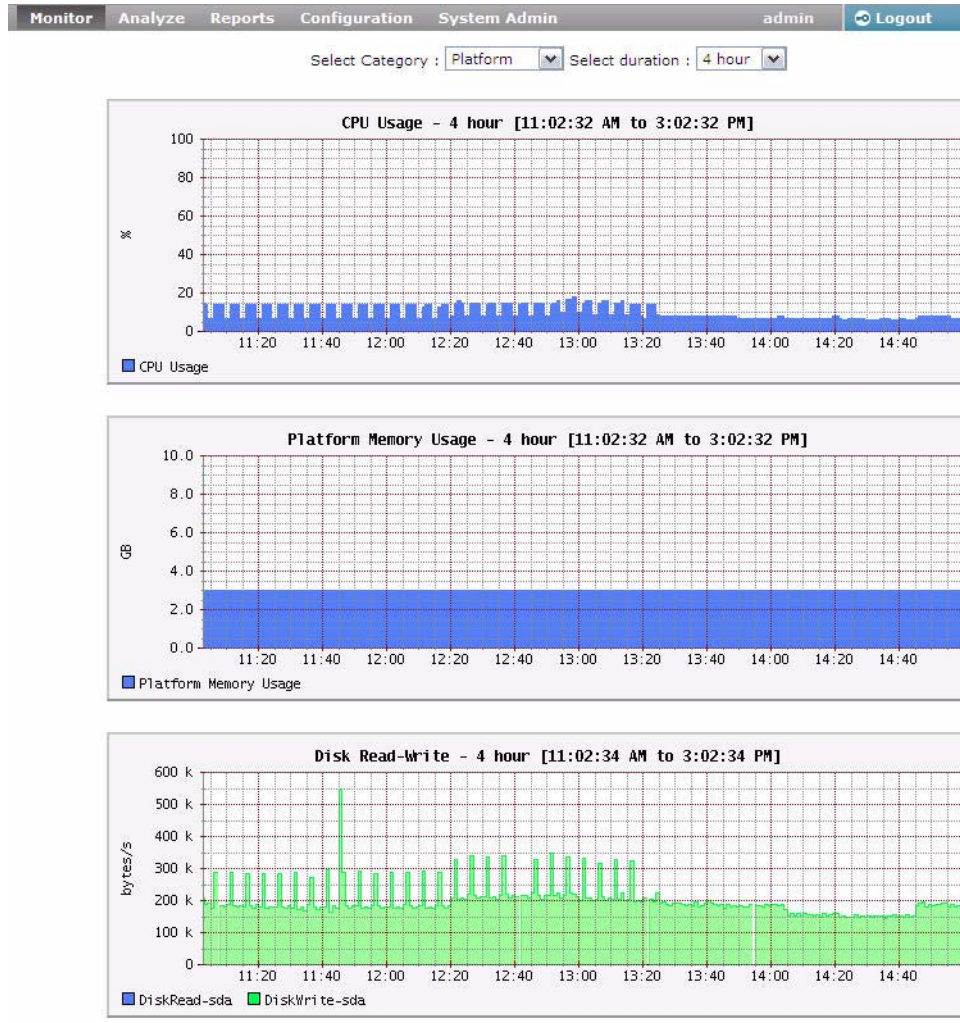


Figure 3-6 Platform page of the Monitor tab

Network

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

Logger

The Logger monitor page, as shown in [Figure 3-7](#), displays details of memory usage as well as information about searches performed.

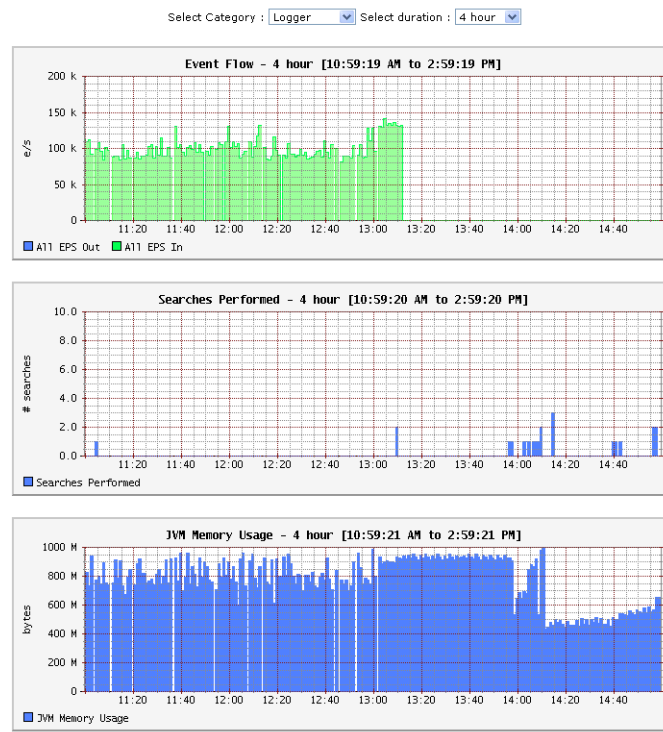


Figure 3-7 Memory usage displayed on the Logger page of the Monitor tab

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 193](#).

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 199](#).

Storage

The Storage monitor page, shown in [Figure 3-8](#), 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 187](#).

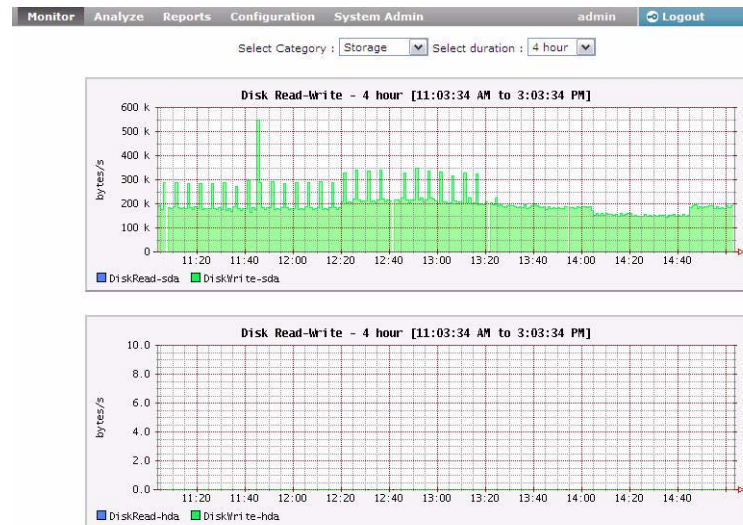


Figure 3-8 Monitor tab, Storage page

Searching and Analyzing Events

This chapter describes how to search for specific events in Logger for analysis. 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 get notified when events matching the criteria you specified are received.

[“The Need to Search Events” on page 43](#)
[“The Process of Searching Events” on page 43](#)
[“Elements of a Search Query” on page 44](#)
[“Syntax Reference for Query Expression” on page 60](#)
[“Using the Search Builder Tool” on page 64](#)
[“Search Analyzer” on page 68](#)
[“Regex Helper Tool” on page 70](#)
[“Searching for Events on Logger” on page 72](#)
[“Understanding the Search Results Display” on page 74](#)
[“Exporting Search Results” on page 77](#)
[“Indexing” on page 80](#)
[“Saving Queries \(Saved Filters and Searches\)” on page 84](#)
[“System Filters/Predefined Filters” on page 86](#)
[“Alerts” on page 89](#)

The Need to Search Events

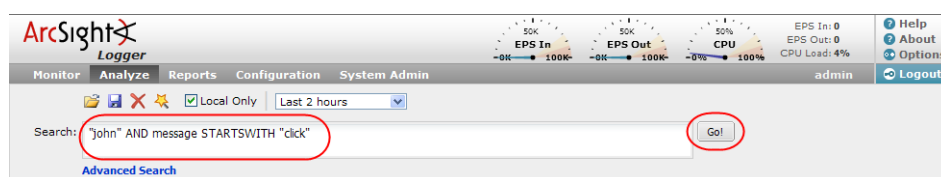
When you need to analyze events matching a specific criteria, include them in a report, or forward them to another system such as ArcSight ESM, you will need to search for them on the Logger.

You need to create queries to search for events. Queries can be as simple as a term to match, such as “login” or an IP address; or they can be more complex, such as events that include multiple IP addresses, ports, and occurred between specific time ranges from devices that belong to a specific device group.

The Process of Searching Events

Searching through stored events is very simple and intuitive on Logger. It uses a flow-based 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.



You enter the keywords or information you are searching for (referred as queries) in the Search text box, select the time range, and click Go, as shown in the previous figure. Logger searches for the data that matches the criteria you specified and displays the results on the same user interface page where you entered your query.

A query can be as simple as a keyword; for example, `hostA.companyxyz.com`. Or a complex query that includes boolean expressions of keywords and indexed fields, and regular expressions; for example:

```
_storageGroup IN ["Default Storage Group"] _deviceGroup IN
["192.168.22.120 [TCPC]"] name="*[4924TestAlert]*" AND ("192.168.*"
OR categoryBehavior CONTAINS Stop) | REGEX=":\d31" | cef name
deviceEventCategory | chart _count by name
```

Additionally, a query can include constraints that limit the search to specific device groups and storage groups.

Logger offers several convenient ways to enter a search query—typing the query in the Search text box, using Logger's Search Builder tool to create a query, or using a previously saved query (referred to as filter or saved search). When you type a query, the auto-suggest facility in the user interface provides suggestions and possible matches for the fields you are entering and the applicable operators for those fields, thus enabling you to quickly build a query expression. The auto-suggest facility is available only for fields in the Logger schema, metadata terms (`_storageGroup`, `_deviceGroup`, `_peerLogger`), and the regular expression term (`|REGEX=`). (See ["Indexing" on page 80](#) for a complete list of fields.)

Although a search query on Logger is as simple as entering a keyword to match, you will utilize the full potential of Logger's search operation if you are familiar with all the elements of a query, as described in the next section, ["Elements of a Search Query" on page 44](#).

Elements of a Search Query

A simple Logger search query consists of these elements:

- Query Expression
- Time range
- Field Set

An advanced Logger search query can also include constraints that limit the search to specific device groups, storage groups, and peer Loggers.

Query Expression

A query expression is a set of conditions that are used to select or reject an event when a search is performed. The expression can specify a very simple term to match such as "login" or an IP address; or it can be more complex, such as events that include several IP

addresses, ports, and occurred between specific time ranges from devices that belong to a specific device group.

A query expression is what you specify in the Search text box on Logger and is specified in 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 expression 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.

Once you run a search query, search results (in tabular form and a histogram) are previewable, that is, immediately displayed on the user interface even if the query has not finished scanning all data. As additional events are matched, the search results table and the histogram are refreshed. Certain search operators such as head, tail, and so on however require a query to finish running before search results can be displayed.

Indexed Search is described in [“Indexed Search” on page 45](#).

Search Operators are described in [“Search Operators” on page 47](#).

Indexed Search

The *Indexed search* uses Logger’s indexing capability to quickly and efficiently search for relevant data, and enables you to specify **keywords**, **indexed**, and **non-indexed fields** in a boolean expression.

Keywords

Keywords are words expressed in plain English. For example, failed, login, and so on. Make sure you understand and follow the requirements and guidelines listed in [“Syntax Reference for Query Expression” on page 60](#).

Multiple keywords can be specified in one query expression by using boolean operators between them. Boolean expressions can be nested; for example, `(John OR Jane) AND Doe*`. Although the boolean operators AND, OR, and NOT can be specified in upper-, lower-, or mixed case when used as an operator, ArcSight recommends that you use uppercase. To search for these words (upper-, lower-, or mixed case) in events, enclose them in double quotes (“”). For example, “and”, “OR”, and so on.

Keyword search is case insensitive.

Indexed and Non-Indexed Fields

The Logger indexing capability allows for *fields* of events to be indexed. The Logger’s search operation and reports utilize these indexed fields to yield significant search and reporting performance gains.

Although you can add indexed and non-indexed fields to a search query, **you will realize the search and reporting performance gains only if all fields in a query are indexed**. (For more information and a list of fields you can index, see [“Indexing” on page 80](#). For discussion on field-based query performance, see [“Performance Optimizations for Indexed Fields in Search Queries” on page 69](#).)

Field search is case sensitive. Make sure you understand and follow other requirements and guidelines listed in [“Syntax Reference for Query Expression” on page 60](#).

You can specify multiple field conditions and also connect keywords to field conditions in a query expression; when doing so, connect them with boolean operators. For example, the following query searches for events with keyword "failed" (without double quotes) or events with "name" field set to "failed login" (lowercase only; without double quotes) and the message field not set to "success" (lowercase only; without double quotes):

```
failed OR (name="failed login" AND message!="success")
```



Note

If a query includes the boolean operator OR and the metadata identifiers (discussed in ["Constraints" on page 58](#)), 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 on the user interface screen.

A complete list of fields you can specify is available in ["Indexing" on page 80](#) section. The operators you can use are listed in the following table. Multiple field conditions can be specified 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"`.

Any literal operator in the following list 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"`.


Field Operator	Example
String Operators	
!=	message!="failed login" message!=failed*login (* means wildcard) message!=failed*login (* is literal in this case)
=	message="failed login" message="failed*login" (* means wildcard)
>	These operators evaluate the condition lexicographically. For example, <code>deviceHostName BETWEEN AM AND EU</code> 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.
<	
>=	
<=	
BETWEEN	
IN	
STARTSWITH	message STARTSWITH "failed"
ENDSWITH	message ENDSWITH "login"
CONTAINS	message CONTAINS "foobar"
Numeric / Timestamp Operators	
=	bytesIn = 32

Field Operator	Example
!=	destinationPort != 100
>	bytesIn > 100
>=	endTime >= "01/13/2009 07:07:21" endTime >= "2009/13/01 00:00:00 PDT" endTime >= "Sep 10 2009 00:00:00 PDT"
<	startTime < "\$now - 1d"
<=	startTime <= "\$now - 1d"
BETWEEN	priority BETWEEN 1 AND 5
SQL Operator	
IS	sessionId IS NULL sessionId IS NOT NULL
Boolean Operators	
AND	name="Data List" AND message="Hello" AND 1.2.3.4
OR	(name="TestEvent" OR message="Hello") AND type=2 AND 1.2.4.3
NOT	NOT name="test 123"
List Operator	
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"]

Search Operators

The *Search Operators* enable you to further refine the data that matched the indexed search filter. Two search operators—**cef** and **rex**—enable you to extract information of interest to you from the events that matched the indexed search filter (the query portion before the first pipeline in the query expression). The **cef** operator is useful for CEF events. It enables you to extract CEF fields from the events. The **rex** operator is useful for syslog (raw) events 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**,

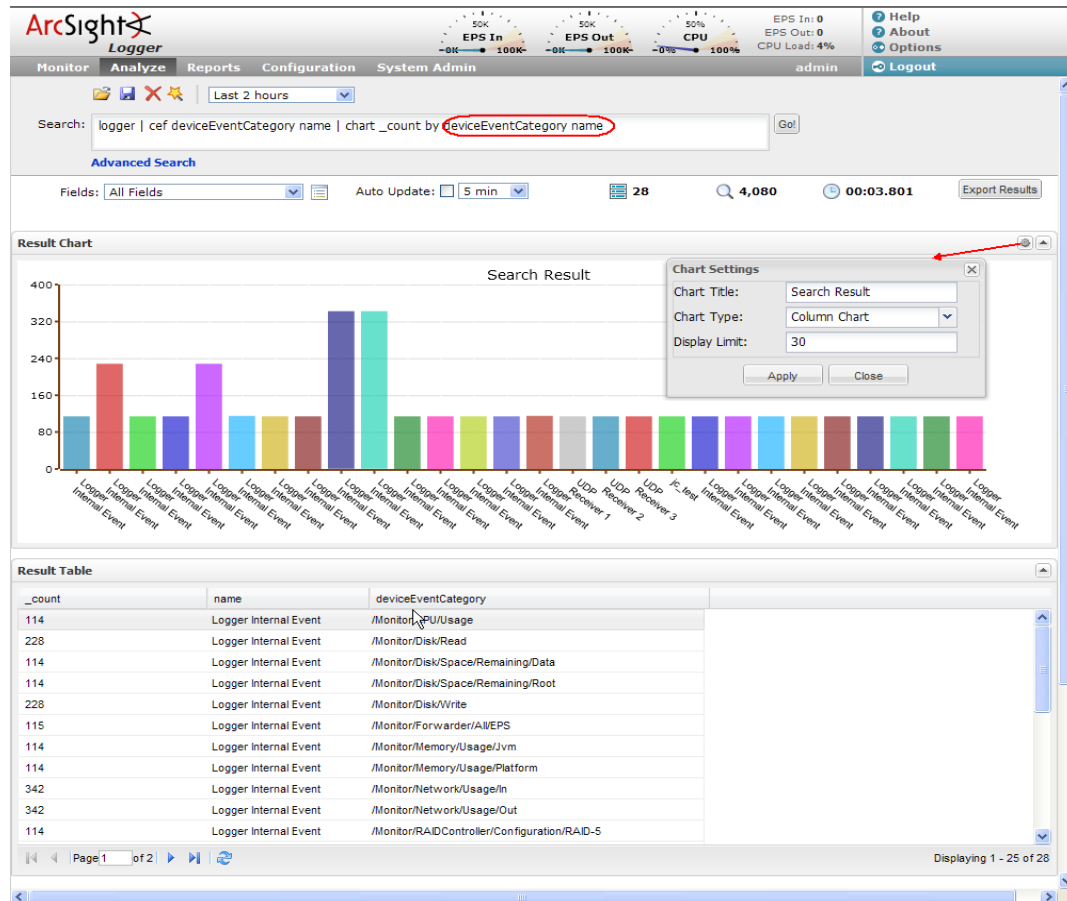
`chart`, `sort`, `fields`, and `eval` are applied to the information you extract using the `cef` or `rex` operator.

Search Operator	Description and Examples
cef	<p>Extracts values for specified fields from matching CEF events. If an event is non-CEF, the field value is set to NULL.</p> <p>Usage: <code>cef field1 field2 field3 ...</code></p> <p>Notes:</p> <ul style="list-style-type: none"> • If multiple fields are specified, separate the field names 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. • If you want to use other search operators such as <code>fields</code>, <code>sort</code>, <code>chart</code>, and so on to refine your search results, you must first use this operator to extract those fields. <p>Examples:</p> <ol style="list-style-type: none"> 1. <code>categorySignificance STARTSWITH "/Normal" AND agentType STARTSWITH "testalert" cef categorySignificance agentType</code> 2. <code>logger cef deviceEventCategory name</code>
chart	<p>Displays search results in a chart form of the count of unique values for the specified fields. When multiple fields are specified, the count of unique sets of all those fields is plotted, as illustrated in example below.</p> <p>By default, a line chart is displayed. You can select from other types of charts, such as bar chart, column chart, pie chart, or area chart.</p> <p>Usage: <code>chart _count by field1 field2 field3 ...</code></p> <p>Notes:</p> <ul style="list-style-type: none"> • A <code>cef</code> or <code>rex</code> operator (to extract fields from matching events) must precede this operator, as shown in the example below. • 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: <ul style="list-style-type: none"> - Title: Enter a meaningful title for the chart. - Type: Column, Bar, Pie, Area, Line - Display Limit: Number of unique values to plot. Default: 10 <p>If the configured Display Limit is less than the number of unique values for a query, the top values equal to the 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.</p> • If multiple fields are specified, separate the field names with a white space or a comma. <p>Example: Example provided in the next row.</p>

Search Operator Description and Examples

Example for Chart Operator, when multiple fields are specified:

logger | cef deviceEventCategory name | chart _count by deviceEventCategory name



Search Operator	Description and Examples
eval	<p>Display 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.)</p> <p>Usage: <code>eval <expression></code></p> <p><code><expression></code> is a mathematical, string, or boolean operation; for example, <code>total_bytes=bytesIn + bytesOut</code>.</p> <p>Note: Typically, a <code>cef</code> or <code>rex</code> operator (to extract fields from matching events) precedes the <code>eval</code> operator, as shown in the examples below. However, you can use the <code>eval</code> operator on a field that has been defined by a previous <code>eval</code> operator in a query.</p> <p>Examples:</p> <ol style="list-style-type: none"> 1. <code>_storageGroup IN ["Default Storage Group"] cef categoryBehavior eval cat=if(categoryBehavior==" /Communicate", "communicate", "notCommunicate")</code> If the Category Behavior is "Communicate", then assign the value "communicate" to a new field "cat"; otherwise, assign the value "notCommunicate" to it. 2. <code>logger cef msg name eval fullname=name + "END"</code> 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". 3. <code>_storageGroup IN ["Default Storage Group"] cef bytesIn bytesOut name eval Plus=bytesIn + 100 sort Plus</code> 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.
fields	<p>Include or exclude specified fields from search results.</p> <p>Usage: <code>fields [(+ -) field]+</code></p> <ul style="list-style-type: none"> + includes only the specified field or fields in the search results. Default. - excludes only the specified field or fields from the search results. <p>Notes:</p> <ul style="list-style-type: none"> Typically, a <code>cef</code> or <code>rex</code> operator (to extract fields from matching events) precedes the <code>fields</code> operator, as shown in the examples below. However, fields might also be defined by other operators such as the <code>eval</code> operator. The + and - can be used in the same expression when multiple fields are specified. For example, <code> fields + name - agentType</code> 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. <p>Examples:</p> <ol style="list-style-type: none"> 1. <code>categorySignificance STARTSWITH "/Normal" AND agentType STARTSWITH "testalert" cef categorySignificance agentType fields - agentType + categorySignificance</code> 2. <code>logger cef name rex "\sInternal(?<eventName>.*\d{1,3}.\d{1,3}.\d{1,2}.*)" fields - name</code>

Search Operator	Description and Examples
head	<p>Displays the first <N> lines of the search results.</p> <p>Usage: <code>head [<N>]</code></p> <p><N> is the number of lines to display. Default: 10, if <N> is not specified.</p> <p>Notes:</p> <ul style="list-style-type: none"> A <code>cef</code> or <code>rex</code> operator (to extract fields from matching events) must precede this operator, as shown in the example below. 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. <p>Example: <code>arcsight cef deviceEventCategory head</code></p>
rare	<p>List 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.</p> <p>When multiple fields are specified, the count of unique sets of all those fields is listed from the lowest to highest count.</p> <p>Usage: <code>rare field1 field2 field3 ...</code></p> <p>Notes:</p> <ul style="list-style-type: none"> A <code>cef</code> or <code>rex</code> operator (to extract fields from matching events) must precede this operator, as shown in the example below. A chart of the search results is automatically generated when this operator is included in a query. If multiple fields are specified, separate the field names with a white space or a comma. <p>Example: <code>arcsight cef deviceEventCategory rare deviceEventCategory</code></p>
regex	<p>Select events that match the specified regular expression.</p> <p>Usages:</p> <p> <code>regex <regular_expression></code></p> <p>OR</p> <p> <code>regex field_name (= !=) <regular_expression></code></p> <p>Note:</p> <ul style="list-style-type: none"> If you are an existing Logger customer, please note that the regular expression syntax has changed. An "equal to" ("=") sign is no longer needed between the <code>regex</code> operator and the regular expression. An "equal to" ("=") or "not equal to" ("!=") sign is only required when equating field names in a regular expression (as in Example #2 below). 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), make sure a <code>cef</code> or <code>rex</code> operator (to extract fields from matching events) precede this operator, as shown in the example below. <p>Examples:</p> <ol style="list-style-type: none"> <code>_storageGroup IN ["Default Storage Group"] regex "failure"</code> <code>logger cef deviceEventCategory regex deviceEventCategory != "fan"</code>

Search Operator	Description and Examples
rex	<p>Extract (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.</p> <p>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 ?<field_name>, where field_name can be a string of alphanumeric characters, beginning with a letter or a "\$" sign. Using an underscore ("_") is not recommended.</p> <p>For example, to extract the IP address from the following event and assign it to a field "clientip", specify <code>"\[client (?<clientip_1>[^\]]*)"</code> as the regular expression. In this regular expression ?<clientip_1> is the field name defined to capture IP address from an event in which the IP address is followed by the literal word "client".</p> <pre>[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]: Can't connect to 10.4.31.4:11211</pre> <p>Usage: <code> rex <regular_expression containing a field name></code></p> <p>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.</p> <p>Usage: <code> cef <field> rex field = <field> mode=sed "s/<string to be substituted>/<substitution value>/g"</code></p> <p>In Example #3 below, 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 the first byte of any destination IP address that starts with "10" is substituted with "xxxx".</p> <p>Notes:</p> <ul style="list-style-type: none"> A detailed tutorial on the rex operator is available at Appendix C, Using the Rex Operator, on page 387. A Regex Helper tool is available for formulating regular expressions of fields in which you are interested. The Regex Helper parses a raw syslog 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 70. 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 "clientip" 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. <p>Examples:</p> <ol style="list-style-type: none"> <code>_storageGroup IN ["Default Storage Group"] rex "\[client (?<clientip>[^\]]*)"</code> <code>_storageGroup IN ["Default Storage Group"] cef deviceEventCategory eventId rex "deviceEventCategory=(?<categories>[^\]]*)"</code>

Search Operator	Description and Examples
rex (contd.)	<p>Example #3. <code>_storageGroup IN ["Default Storage Group"] cef msg rex field=msg mode=sed "s/Authentication/xxxx/g" cef agentAddress rex field=agentAddress mode=sed "s/192/xxxx/g" cef dst rex field=dst mode=sed "s/10./xxxx/g"</code></p>
sort	<p>Sort search results as specified by the sort criteria.</p> <p>Usage: <code> sort [<N>] ((+ -) field)+</code></p> <ul style="list-style-type: none"> + Sort the results by specified fields in ascending order. This is the default. - Sort the results by specified fields in descending order. <p><N> Keep the top N results, where N can be a number between 1 and 10,000. Default: 10,000.</p> <p>Notes:</p> <ul style="list-style-type: none"> A <code>cef</code> or <code>rex</code> operator (to extract fields from matching events) must precede this operator, as shown in the example below. 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, <code> sort + deviceEventCategory - eventId</code> 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 of this release 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. <p>Example: <code>_storageGroup IN ["Default Storage Group"] cef deviceEventCategory eventId sort deviceEventCategory eventId</code></p>
tail	<p>Displays the last <N> lines of the search results.</p> <p>Usage: <code> tail [<N>]</code></p> <p><N> is the number of lines to display. Default: 10, if <N> is not specified.</p> <p>Notes:</p> <ul style="list-style-type: none"> A <code>cef</code> or <code>rex</code> operator (to extract fields from matching events) must precede this operator, as shown in the example below. 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. <p>Example: <code>arcsight cef deviceEventCategory tail</code></p>

Search Operator	Description and Examples
top	<p>List 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.</p> <p>When multiple fields are specified, the count of unique sets of all those fields is listed from the highest to lowest count.</p> <p>Usage: <code>top field1 field2 field3 ...</code></p> <p>Notes:</p> <ul style="list-style-type: none"> • A <code>cef</code> or <code>rex</code> operator (to extract fields from matching events) must precede this operator, as shown in the example below. • If multiple fields are specified, separate the field names with a white space or a comma. • A chart of the search results is automatically generated when this operator is included in a query. <p>Examples:</p> <ol style="list-style-type: none"> 1. <code>arcsight cef deviceEventCategory top deviceEventCategory</code> 2. <code>_storageGroup IN ["Default Storage Group"] cef deviceEventCategory eventId rex "deviceEventCategory=(?<categories>[^\s]*)" top categories</code>
where	<p>Display events that match the criteria specified in the “where” expression.</p> <p>Usage: <code>where <expression></code></p> <p><code><expression></code> can be any valid field-based query expression, as described in “Indexed and Non-Indexed Fields” on page 45.</p> <p>Notes:</p> <ul style="list-style-type: none"> • A <code>cef</code> or <code>rex</code> operator (to extract fields from matching events) must precede this operator, as shown in the examples below. • <code><expression></code> can only be a valid field-based query expression. Arithmetic expressions or functions are not supported. • When a where operator is included in a query, the query performance can be significantly impacted. This is a known issue and will be addressed in a future release of Logger. <p>Examples:</p> <ol style="list-style-type: none"> 1. <code>_storageGroup IN ["Default Storage Group"] cef eventId where eventId is NULL</code> 2. <code>_storageGroup IN ["Default Storage Group"] cef eventId deviceVersion where eventId=10006093313 OR deviceVersion CONTAINS "4.0.6.4924.1"</code> 3. <code>_storageGroup IN ["Default Storage Group"] cef deviceEventCategory eventId rex "deviceEventCategory=(?<categories>[^\s]*)" where eventId >=10005985569 OR categories="/Agent/Started"</code>

Time Range

An event is timestamped with the Logger receipt time when it is received on the Logger. A search query uses this time to search for matching events. 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”, a time range window is created that moves with 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.

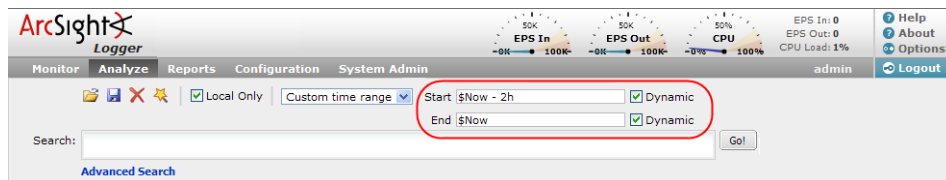
Custom time range: You can specify a time range in a 24-hour format to suit your needs. For example, a custom time range is:

Start: 8/13/2008 13:36:30

End: 8/13/2008 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 mode searches relative to the time when the search 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 4-1 on page 55](#). The <units> portion, if given, consists of an integer and a single, case-sensitive letter, as shown in [Table 4-2 on page 56](#).


Table 4-1 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 4-2 Units

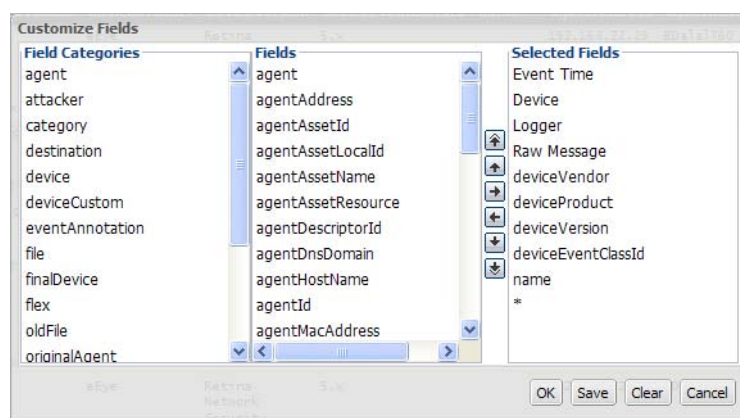
Unit	Description
m (lowercase)	Minutes do not confuse with 'M', meaning months)
h	Hours
d	Days
w	Weeks
M (uppercase)	Months (do not confuse with 'm', meaning minutes)

Field Set

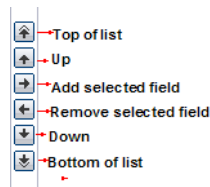
A field set determines the fields that are displayed in the search results for each event that matched a search query. Logger provides a number of predefined field sets, as listed in the following table. To view the fields included in each of the predefined field sets, click the  (Customize Fieldset) icon. When you run the first search operation in a new browser window, you might not be able to select the field sets as they are hidden. The field sets list is displayed after you have run the first search operation.

Field Set	Description
All Fields	To view a list of fields that are included for each field set type, select the field set from the drop-down list and hover your mouse pointer on the Fields: label.
All Fields (w/out raw messages)	
Minimal Fields	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.
Syslog Standard	
Categories	
Base Event Fields	

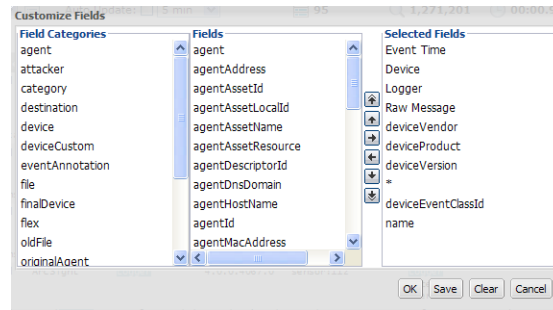
You can also create your own field sets. The Logger user interface offers a simple and intuitive interface to select and move event fields you want to include in a field set, as shown in the following figure.



Use these buttons to create and edit a custom field set.



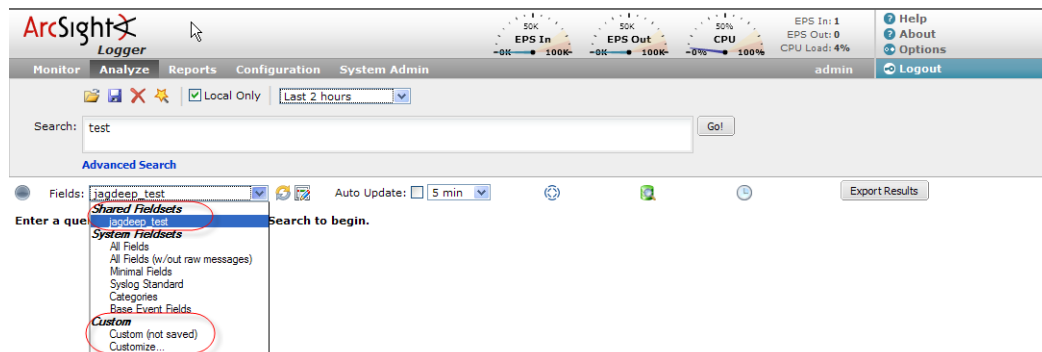
A wildcard field ("*") is available in the Fields list when you create a custom field set. This field includes all fields available in an event that are not individually listed in the custom field set definition. For example, for the following custom field set definition, the search results will list 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 either save the custom field sets you create or use them for the current session.

If you save a custom field set, it appears under the [Shared Fieldsets](#) category and is visible and available to the other users of your Logger, as shown in the following figure. Once a field set is saved, you can edit and delete it.

If you do not save the custom field set, it is temporarily labeled as "Custom (not saved)" and is not visible to other users. Once you log out of the current session, the temporary field set is deleted. You can only create one temporary custom field set at a time.



Field set selection is specific to a Logger user's interface. For example, UserA and UserB are connected to the same Logger and are using the default, All Fields, field set for search results display. UserA changes his selection to a custom field set. This change will only

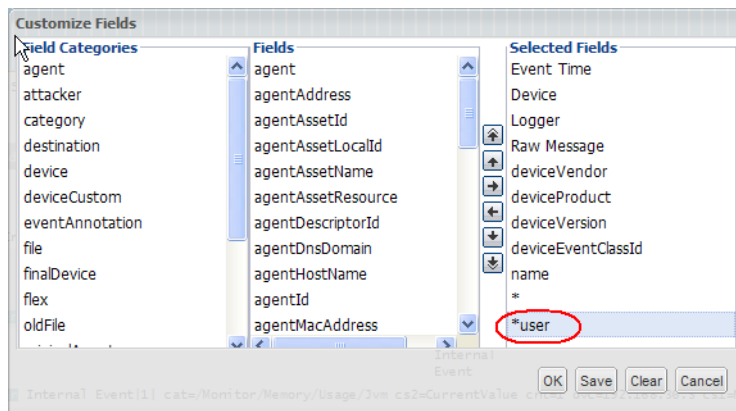
impact UserA's display; UserB will continue to see the search results in the All Fields format.



Field sets are not included in the saved filter definition.

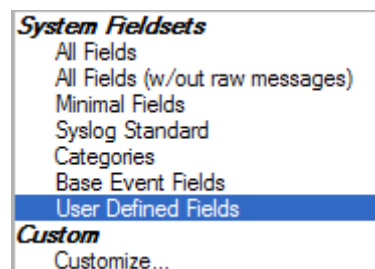
For information about deleting custom field sets, see [“Deleting Custom Field Sets” on page 231](#).

When you use a search operator that defines a new field, such as `cef`, `rex`, or `eval`, a new column for each field is added to the currently selected display. These newly defined fields are displayed by default. A new field, `*user` (shown below), in field sets controls the display of fields defined by search operators. When `*user` is included in the Selected Fields list of a custom field set, the newly defined fields are displayed.



A new field set, User-Defined Fields, is also available that enables you to view only the newly defined fields.

The “User-Defined Fields” field set is available as a drop-down option from the “Fields:” menu on the page where search results are displayed.



Constraints

Constraints enable you to limit a query to events from one or more of the following:

- Devices in a particular device group
- Stored in particular storage groups
- On specific peer Loggers

For example, you might want to search for events for devices in the SMR-1 and SMR-2 device groups on the local Logger only.

Using constraints can speed up a search operation as they limit the scope of data that needs to be searched. 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: Use this to also 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 the metadata identifiers (discussed in [“Constraints” on page 58](#)), 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.

Additionally, if this Logger peers with a Logger that runs v4.0 GA, make sure that the queries sent from that Logger to this Logger also follow this guideline.

- When specifying multiple groups in a constraint, ensure that the group names are enclosed in a square bracket; for example, `_storageGroups IN ["SGA", "SGB"]`.
- You can apply constraints to a search query in these ways:
 - ◆ Typing the constraint in the Search text box

Use Logger's auto-suggest facility to enter a constraint in the Search text box. Once you type `"_s"` (for storage group), `"_d"` (for device group), or `"_p"` (for peerLogger) in the Search text box, Logger automatically provides a drop-down list of relevant terms and operators from which you can select.

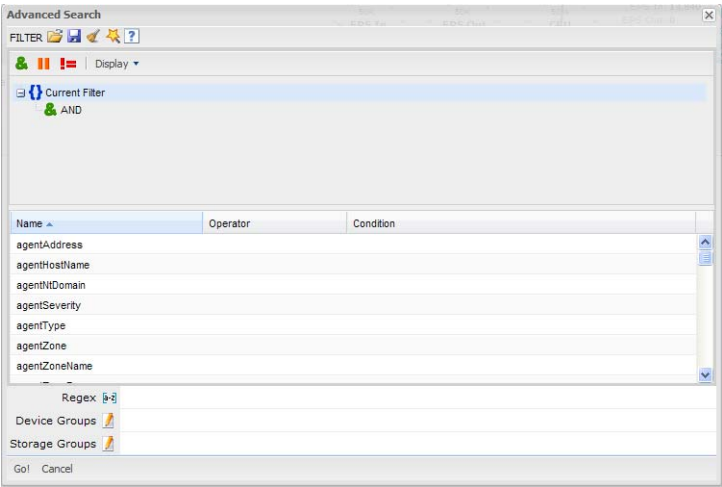


Caution

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 ["10.10.10.10"] name contains abc | REGEX=":\d31"`

- ◆ Using the Search Builder tool as you can select the constraints in it, as shown in the following figure. (To access the Search Builder tool, click **Advanced** to the right of the Search text box where you type query expression.) For more

information about the Search Builder, see [“Using the Search Builder Tool” on page 64](#).



Syntax Reference for Query Expression

You must understand and follow specific requirements for creating query expressions so that you create valid and accurate expressions. The following table lists those requirements.

Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
Syntax	keyword1 boolean_operator keyword2 boolean_operator keyword3...	field_name operator field_value (List of fields in the “Indexing” on page 80 section.) (List of operators in the “Indexed and Non-Indexed Fields” on page 45 section.)	REGEX="<REGEX1>" REGEX="<REGEX2>" ..

Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
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 (discussed in "Constraints" on page 58), 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 "Indexed and Non-Indexed Fields" on page 45 section.</p> <ul style="list-style-type: none"> Unless a value is enclosed between double quotes, a space between value is interpreted as an AND. For example, name=John Doe is interpreted as John AND Doe. 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 (discussed in "Constraints" on page 58), 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 54.</p> <p>Use this operator to AND multiple regular expressions in one query expression.</p>

Behavior	Full Text (Keyword)	Field Search (Indexed)	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 (_storageGroup, _deviceGroup, and _peerLogger), 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 “Indexed and Non-Indexed Fields” on page 45 section to connect and nest field search expressions. Metadata identifiers (_storageGroup, _deviceGroup, and _peerLogger), but can only appear at the top level in a query expression 	<p>Multiple regular expression can be specified in one query using this syntax:</p> <pre> REGEX="<REGEX1>" REGEX="<REGEX2>" ...</pre>
Case sensitivity	<p>Insensitive (Cannot be changed.)</p>	<p>Sensitive* (Can be changed using Tuning options. See “Tuning Advanced Search Options” on page 230.)</p>	<p>Insensitive* (Can be changed using Tuning options. See “Tuning Advanced Search Options” on page 230.)</p>
Wildcard	<p>*</p> <p>Cannot be the leading character; only a suffix or in between a keyword.</p> <p>Examples:</p> <ul style="list-style-type: none"> *log is invalid log* is valid lo*g* is valid 	<p>*</p> <p>Can appear anywhere in the value.</p> <p>Note: Logger v3.0 GA and SP1 did not support the use of wildcard character.</p> <p>Examples:</p> <pre>name=*log (searches for ablog, blog, and so on.) name="*log" name=*log (both search for *log)</pre>	<p>*</p> <p>Can appear anywhere.</p>
Exact Match/Search string includes an operator or a special character	<p>Enclose keyword in double quotes; Otherwise, keyword treated as keyword*.</p> <p>Example:</p> <pre>log (matches log, logging, logger, and so on) "log" (matches only log)</pre> <p>Note: See the list of special characters that cannot be searched even when enclosed in double quotes, later in this table.</p>	<p>Enclose value in double quotes</p> <p>Example:</p> <pre>message="failed login"</pre>	<p>No special requirement.</p>

Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
Escape character	<p>\</p> <p>Use to escape \. You cannot escape any other character.</p>	<p>\</p> <p>Use to escape \, ", and *.</p> <p>Examples:</p> <ul style="list-style-type: none"> name=log\\ger (matches log\ger) name = logger* (matches logger*) 	<p>\</p> <p>Use to escape any special character.</p> <p>Example:</p> <p>To search for a term with the character "[":</p> <p> REGEX="logger\[</p>
Escaping wildcard character	<p>Cannot search for *</p> <p>Example:</p> <p>log* is invalid</p>	<p>Can search for * by escaping the character</p> <p>name=log* is valid</p>	<p>Can search for * by escaping the character</p>
<p>Space</p> <p>Tab</p> <p>Newline</p> <p>,</p> <p>;</p> <p>(</p> <p>)</p> <p>[</p> <p>]</p> <p>{</p> <p>}</p> <p>"</p> <p> </p> <p>*</p> <p>></p> <p><</p> <p>!</p>	<p>Cannot search for these characters.</p> <p>Examples:</p> <p>"John Doe" is invalid</p>	<p>No restrictions.</p> <p>Enclose special character in double quotes. Escape the wildcard character and double quotes.</p> <p>Example:</p> <p>name="John* \"Doe"</p> <p>(matches John* "Doe)</p>	<p>No restrictions.</p> <p>Special regular expression characters such as (,), [,], { , }, " , , and * need to be escaped.</p>

Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
= . : / \ @ - ? # \$ & _ %	<p>You can search for these characters in a keyword. However, enclose the keyword in double quotes.</p> <p>Example: "John="</p>	<p>You can search for these characters. Enclose value in double quotes if value contains any of these characters.</p> <p>Example: name="John="</p>	<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, <code> REGEX="^test\$"</code> will search for events containing the word "test" (without quotes) only. Special regular expression characters such as \ and ? need to be escaped.
Time format, when searching for a specific timestamp	<p>No specific format. The query needs to contain the exact timestamp string. For example, "10:34:35".</p> <p>Note: The string cannot contain spaces. For example, "Oct 19" is invalid.</p>	<p>Use this format to specify a timestamp in a query (including double quotes):</p> <p><code>"mm/dd/yyyy hh:mm:ss"</code></p> <p>OR</p> <p><code>"yyyy/mm/dd hh:mm:ss timezone"</code></p> <p>OR</p> <p><code>"MMM dd yyyy hh:mm:ss timezone"</code></p> <p>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.</p>	No restrictions.

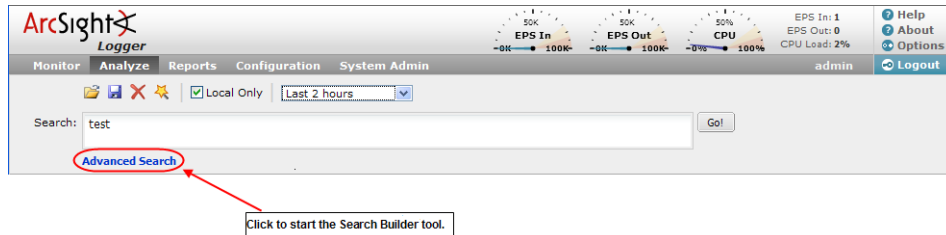
Using the Search Builder Tool

The Logger Search Builder tool is a boolean-logic conditions editor that enables you to quickly and accurately build search queries. The tool provides a visual representation of the conditions you are including in a query. You can specify keywords, field-based conditions,

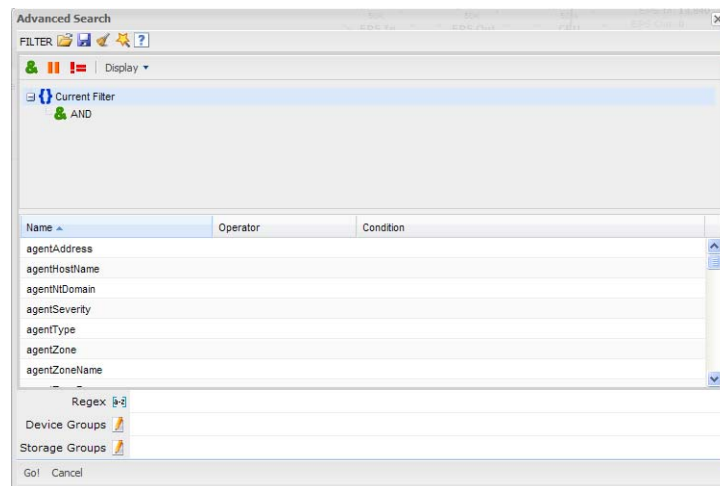
and regular expressions using this tool. In addition, the tool enables you to specify search constraints such as device groups and storage groups (see [“Constraints” on page 58](#)). This section describes how to use the tool.

Accessing Search Builder

To display the Search Builder tool, click **Advanced Search**, below the Search text box, as shown in the following figure.



The Search Builder tool is displayed, as follows:



To build a new search query in the Search Builder tool:

- 1 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

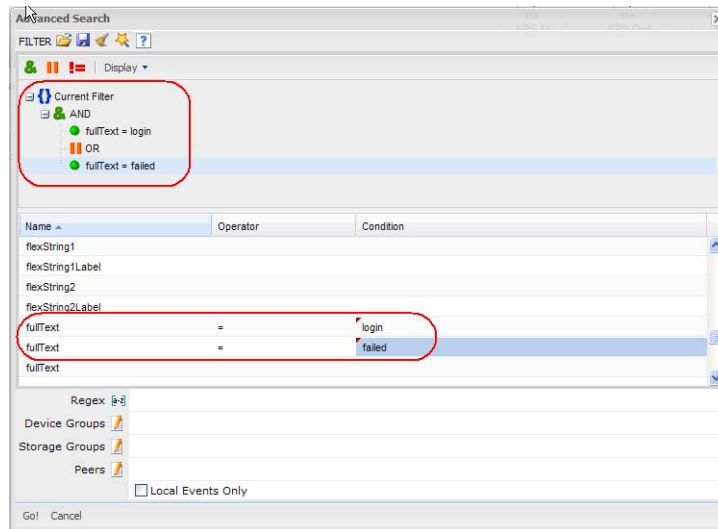
- 2 If you want to load a system or saved filter, or a saved search, click the icon. Select the filter or the saved search from the displayed list and click **Load+Close**.

For more information, see [“Saving Queries \(Saved Filters and Searches\)” on page 84](#) and [“System Filters/Predefined Filters” on page 86](#).

- 3 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**.

Only operators applicable to a field are displayed in the list.

- c 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.

- 4 Repeat [Step 1](#) through [Step 3](#) until you have added all the conditions.
- 5 If your search query will also include a regular expression, type it in the Regex field.
- 6 If you want to constrain your search query to specific device groups, storage groups, and peer Loggers, click the  icon next to the constraint category. Select the relevant groups and peer Loggers. (To select multiple groups, hold the Ctrl-key down.)

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


The peer Logger constraint category is displayed only if peer 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.

- 7 Click **Go**.

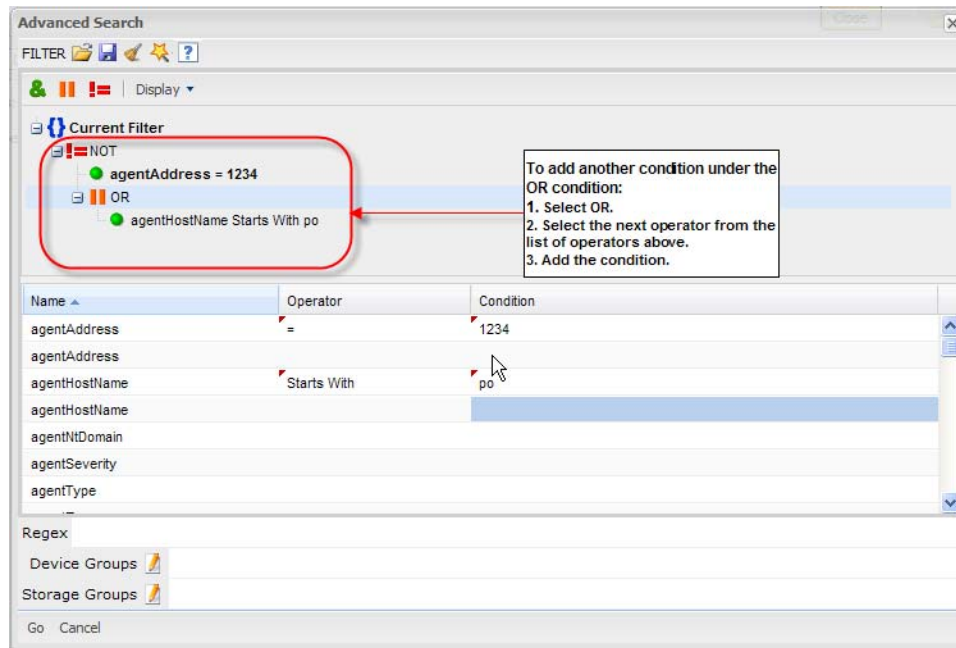
The query is automatically displayed in the Search text box and is ready to be run.

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 84](#).

Nested Conditions

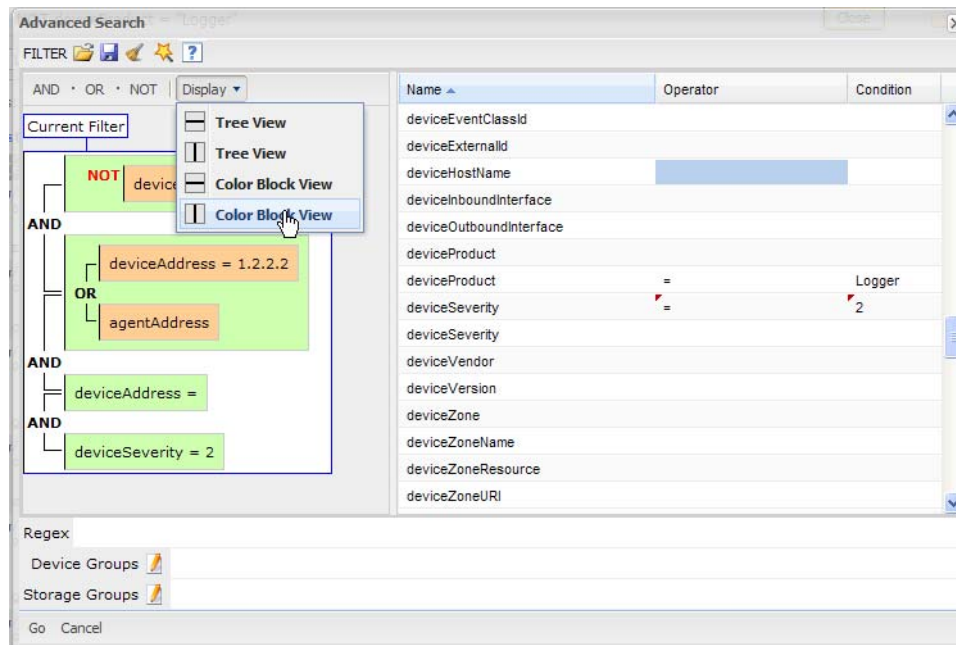
You can create search queries with nested conditions in Search Builder. To do so, click the operator under which you want to nest the next condition and add the condition as described in [“Accessing Search Builder” on page 65](#).



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.

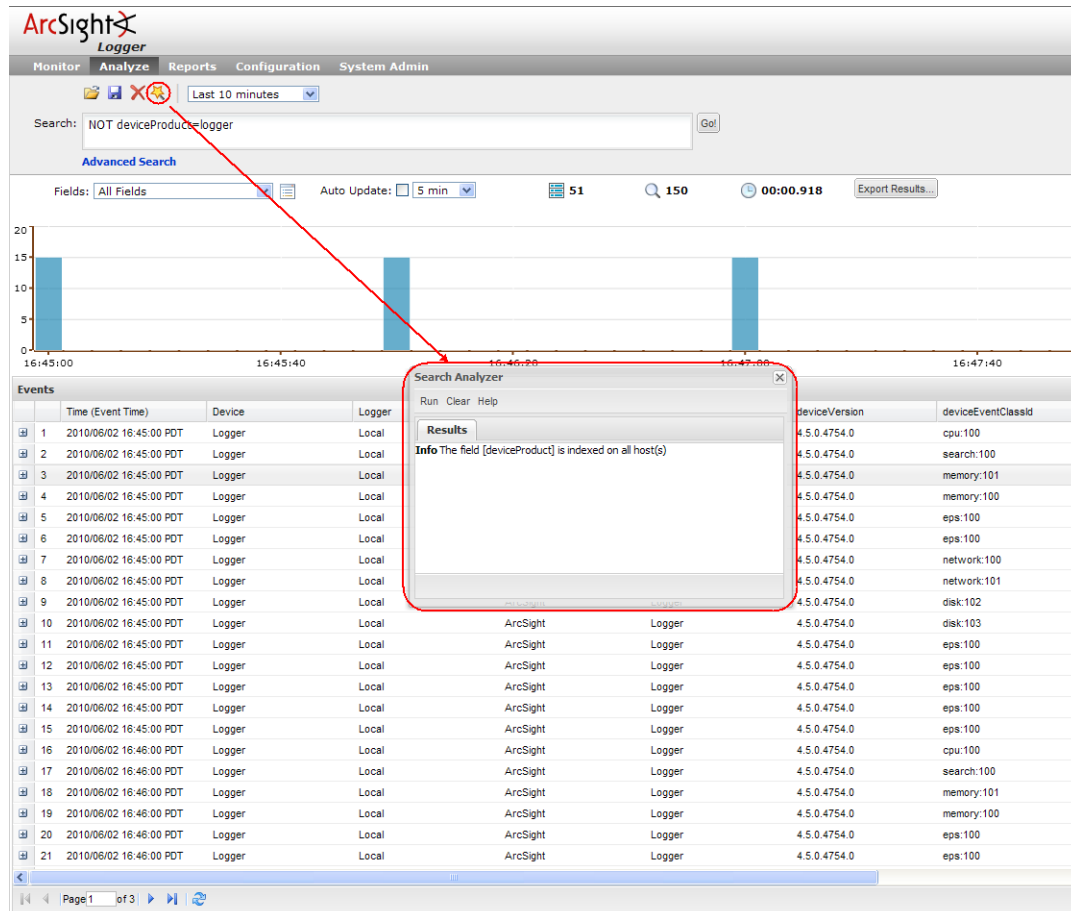
Search Analyzer

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, ArcSight recommends that you identify the fields that you will most commonly use in queries and index all those fields.
- 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.

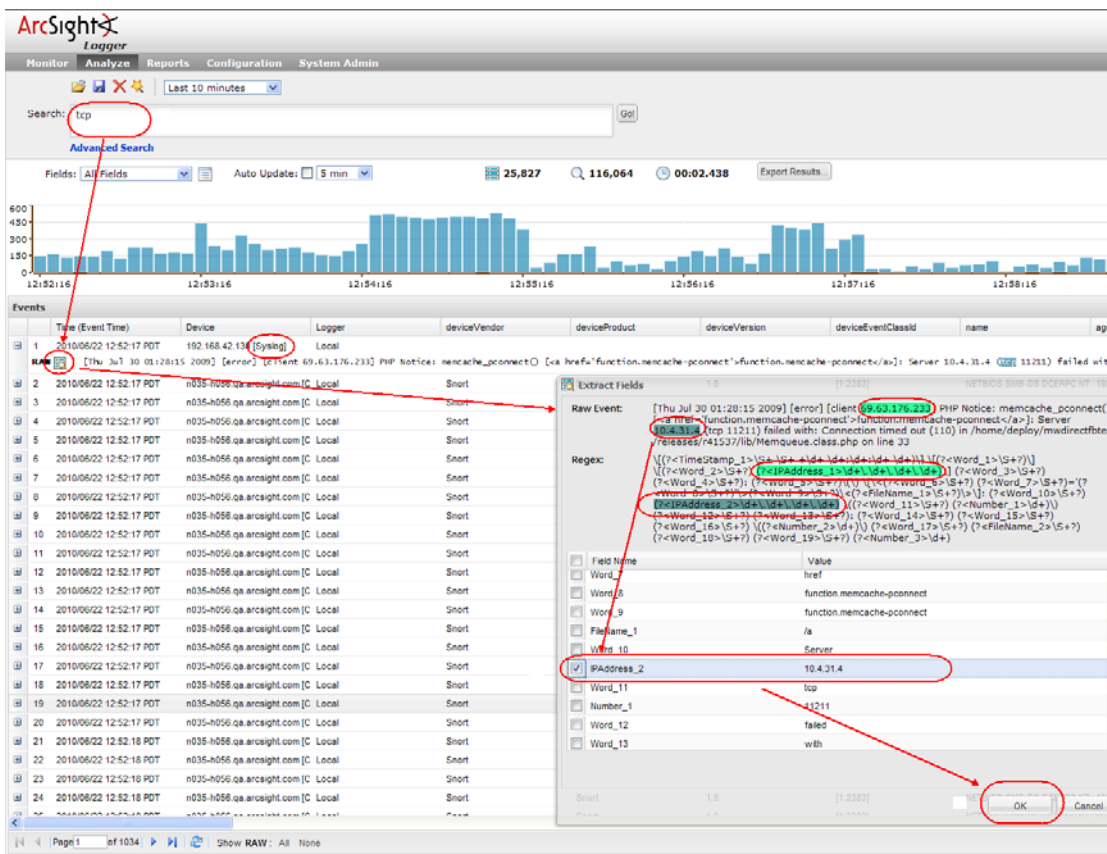
Regex Helper Tool

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 47](#) or [Appendix C, Using the Rex Operator, on page 387](#).) 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 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.

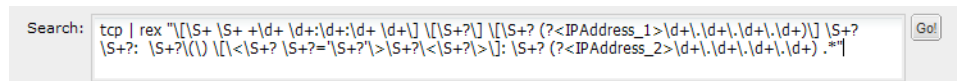
To use the tool, you need to perform the following steps. These steps are also depicted in the figure that follows the steps.

- 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 72](#).)
- 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**.

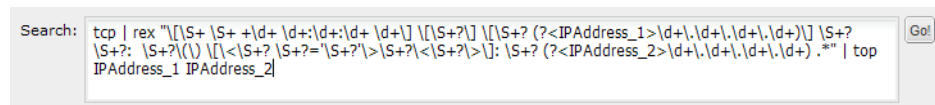


The rex expressions pertaining to the selected fields are automatically entered in the Search query box, as shown in the following figure. In the previous example, the client and server IP addresses need to be extracted from events. Therefore, IPAddress_1 and IPAddress_2 fields were selected in the Regex 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.

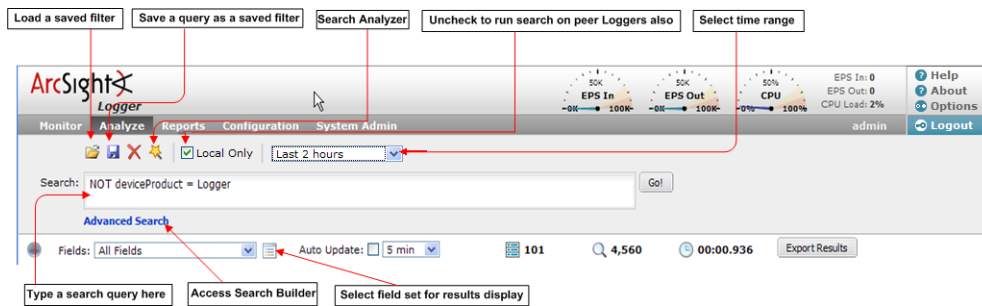


From this point, you can include additional pipeline operators in this query to create charts, identify the top five IP addresses, and so on. In the following example, the above query is modified to identify the top IP addresses.



Searching for Events on Logger

A user needs to belong to a Logger Search Group with the “Search for events” user right set to Yes to perform local searches and “Search for events on remote peers” user right set to Yes to perform peer searches.



To search for events on Logger:

- 1 Click **Analyze > Search**.
- 2 Specify a query expression in the Search text box using one or more of the following methods.


Note: Refer to “[Query Expression](#)” on page 44 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 44.

When you type a query, Logger’s auto-suggest facility enables you to quickly build a query expression by automatically providing suggestions, possible matches, and applicable operators for the following:

- Fields in Logger schema
(See “[Indexing](#)” on page 80 for a complete list of fields.)
- Metadata terms (`_storageGroup`, `_deviceGroup`, `_peerLogger`)
Type “_s” (for storage group), “_d” (for device group), or “_p” (for peerLogger) in the Search text box to obtain a drop-down list of constraint terms and operators.
- Regular expression term (`|REGEX=`)


Note: 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”]`.

- b Click **Advanced** to use the Search Builder tool. (See “[Using the Search Builder Tool](#)” on page 64 for more information.) Also, use this option to specify device groups, storage groups, and peer Loggers to which search should be limited.
- c 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 84 and “[System Filters/Predefined Filters](#)” on page 86.

- 3 Use the following default values or change them suit your needs:
 - a **Local Logger:** By default the query is run on the local Logger only. If you want to run the query on the peer Loggers as well, uncheck the “Local Only” field located to the right of the Go! button.
 - b **Time Range:** By default, the query is run on the data received in the last two hours on the Logger. 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 54](#).
 - c **Field Set:** By default, all fields (All Fields) are displayed in the search results. However, you can select another predefined field set or specify a customized field set. For more information about field sets, see [“Field Set” on page 56](#).
- 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 74](#).

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 84](#).

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 230](#).

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 peer Loggers to which the search should be constrained, as described in [“Searching for Events on Logger” on page 72](#).

Follow these guidelines for searching across peers:

- Peer Loggers can run different versions. However, these are the only supported paths for running a search across peers:
 - ◆ A search from a v4.0.x Logger to v4.5
 - ◆ A search from v4.5 Logger to v4.0.x
 - ◆ A search from v4.5 Logger to v4.5
- A search that is run across peers cannot contain pipeline operators, discussed in [“Search Operators” on page 47](#).
- If peer Loggers do not have identical storage or device group names, a search query operation skips searching for events for those groups on those peers.
- A user needs to belong to these user groups with the listed permissions set to Yes to perform peer searches and view their search results:
 - ◆ Logger Search Group with “Search for events on remote peers” user right set to Yes.
 - ◆ Logger Rights Group with the “View registered peers” and “Edit, save, and remove registered peers” user rights set to Yes.

- When a peer Logger becomes unavailable during a search operation, the one of the following errors might be displayed:

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


These error messages can occur when the peer Logger cannot be reached. Restore the peer relationship and run the search again.




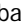
The above listed error messages might still display for the search operation that was in progress even after the peer relationship is restored. However, ignore those messages as these go away when you run a new distributed search.

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.

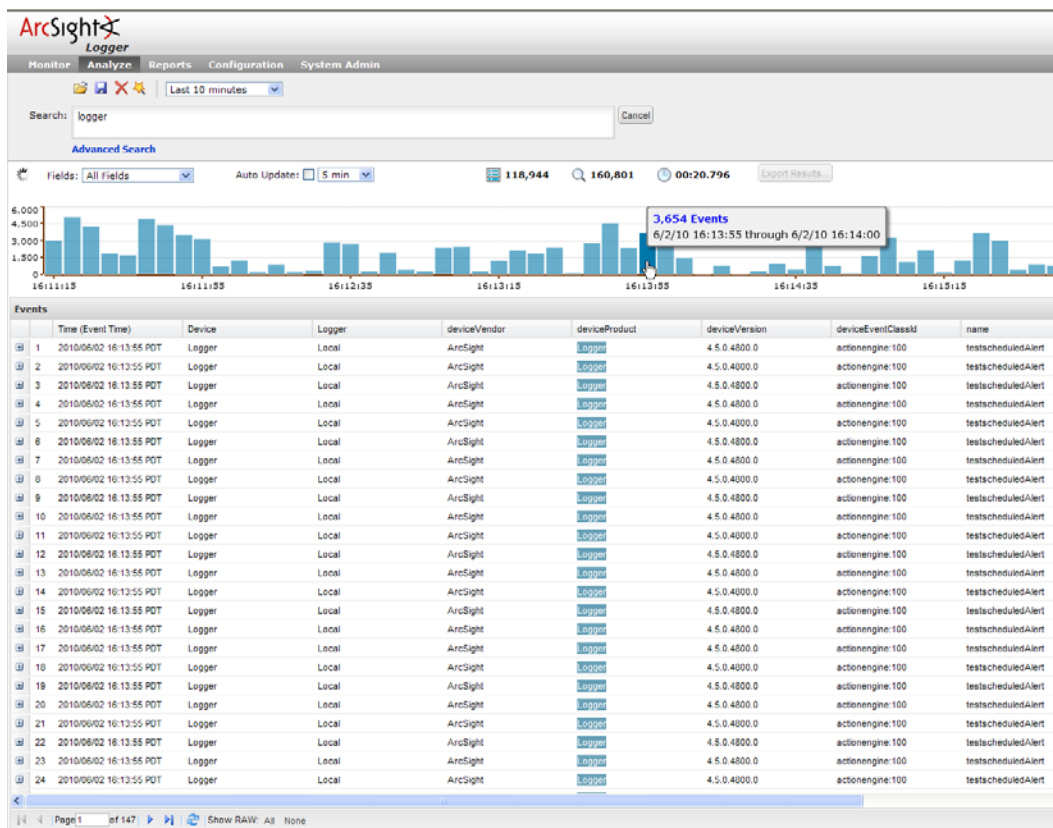
While the search is in progress, the Go! button changes to Cancel—click Cancel to terminate the search early. When a query is running, search results are displayed as matching events are found. Therefore, when you click Cancel, any matching events found so far are displayed as the search results. This 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, or use the histogram to drill down the results. (Note: 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 leftmost column. To see all raw events, click **All** at the bottom 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.

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. Histogram enables you to randomly drill-down to events in a specific time period by clicking the bar representing the time period.

Additionally, the number of events scanned and number of events matching the query and the time it took to run the search is displayed.





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.

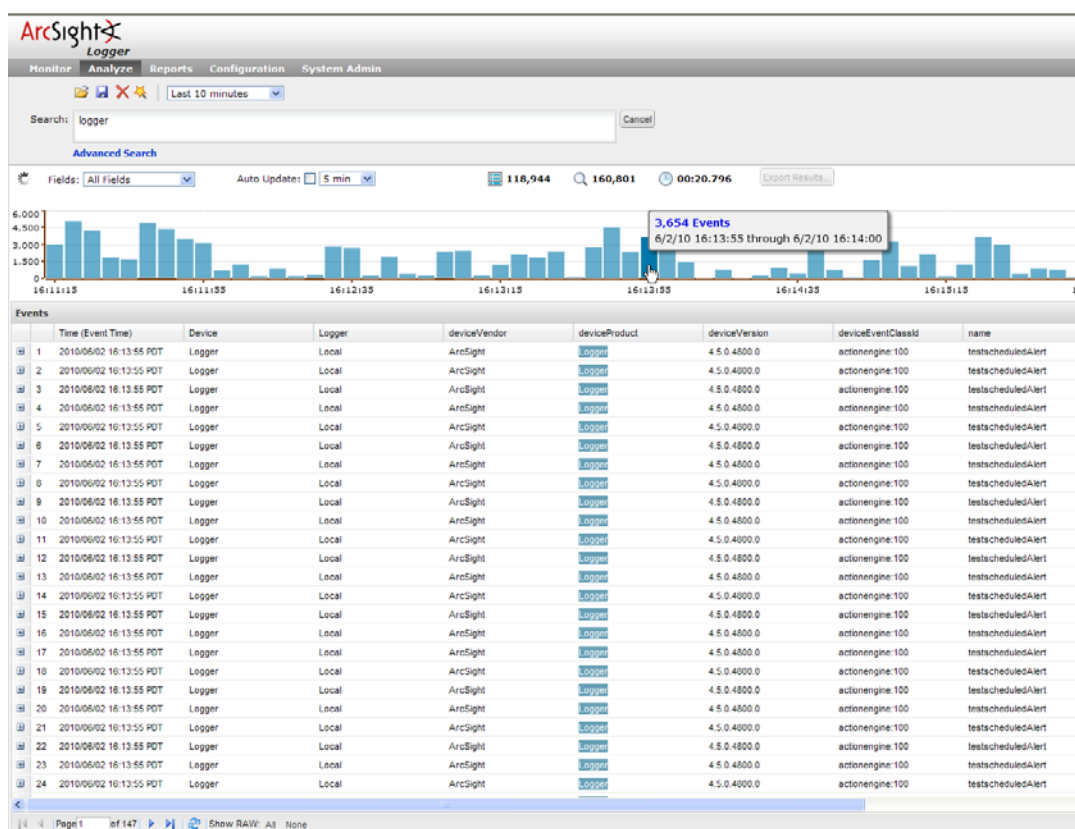
As you roll the mouse over other terms in the events table, they highlight in green. Click a green-highlighted term to add it to the current query. For example, if you search for "login" and roll over the word "fail" in the search results, "fail" will highlight in green. Click the word "fail" to change the query to "login AND fail." You can select only the indexed fields from the search results. Search results are sorted by receipt time.

Guidelines for 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.
- 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.

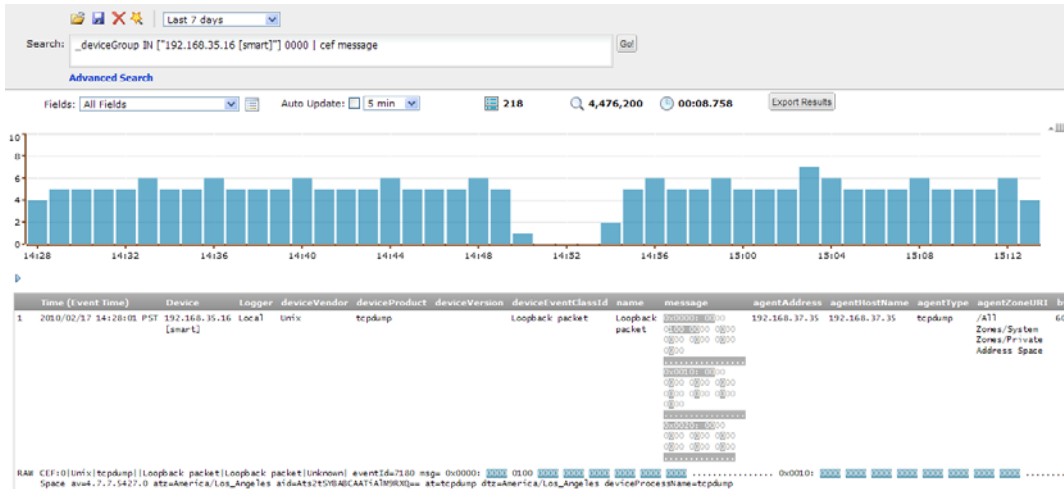
If you need to use the histogram view for event analysis for a search query that matches more than one million events, ArcSight suggests that you adjust the time range specified in your search query such that less than one million are matched to obtain a complete and meaningful histogram or use a pipeline 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 0000 .....
```

Logger v4.x 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.

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 and CEF events, can be included in the exported report.

The PDF export capability is new in Logger v4.x and is described in this document.

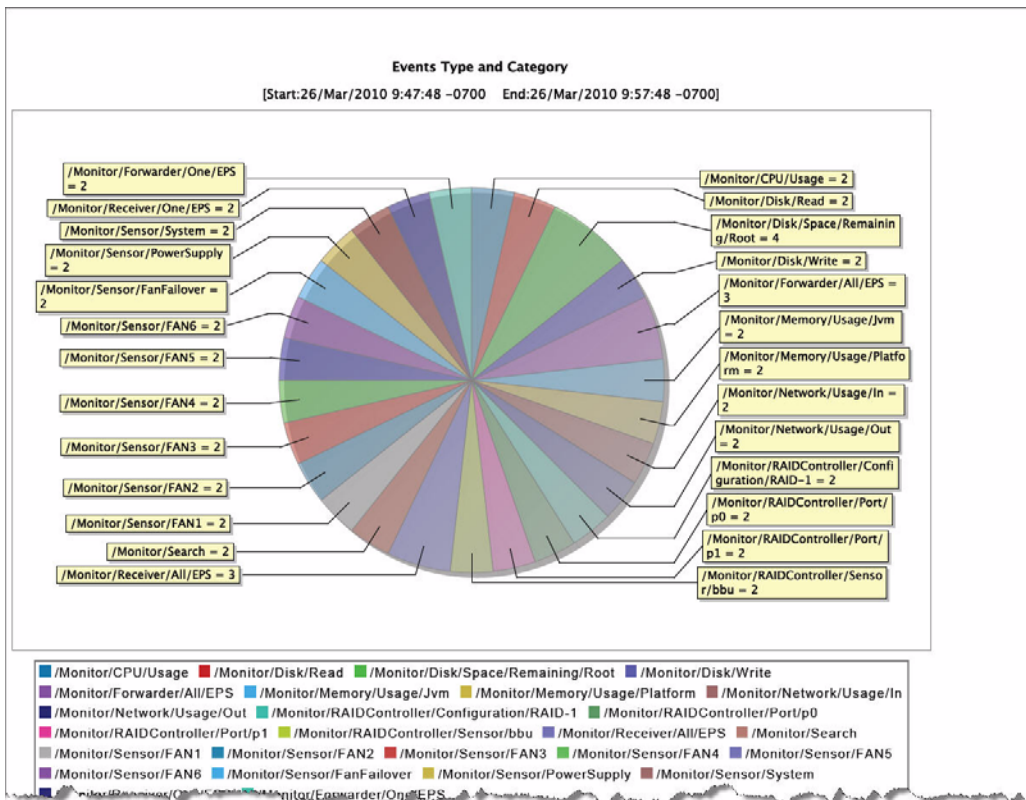
- Comma-separated values (CSV) file—Useful for further analysis with other software applications. The report includes a table of search results. Charts cannot be included in this format.



Note

When you export search results on a Logger in a peer relationship using the dynamic time range option, the query you had run to obtain those results is rerun and the results of the rerun operation are exported. Therefore, the data exported may not exactly match the one displayed in the Search Results screen because the underlying data set would have changed (especially if there is a long delay between the time you run a search query and export its search results, or your Logger is receiving a very high number of events per second).

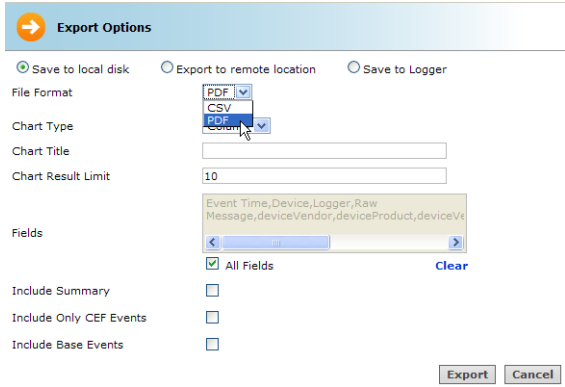
The following is an example of a quick report generated in PDF format. The chart is displayed first, followed by a table of matched events (not shown in this example).



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.

Option	Description
Export to remote location	<p>On a Logger appliance, the file is written to an NFS mount, a CIFS mount, or a SAN system.</p> <p>On the software version of Logger, data is always stored in the <code>/opt/data/logger</code> directory. This directory can reside locally on the system running the Logger software, or on a remote storage system such as SAN, NAS, NFS, or CIFS.</p>
Save to Logger	The file is written to the Logger's local storage.
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. Charts are only included in the PDF file if the search query contains an operator that creates charts, such as <code>chart</code>, <code>top</code>, and so on.</p>
	
Export file name	<p>Specify the name of the file to which events will be exported.</p> <p>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.</p>
Chart Type (for PDF only)	<p>Type of chart to include in the PDF file. You can select from: Column, Bar, Pie, Area, or Line.</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>
Chart Title (for PDF only)	(Optional) A meaningful name that appears on top of the PDF file. If no title is specified, "Untitled" is included.
Chart Result Limit (for PDF only)	<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>
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>

Option	Description
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, ArcSight 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 226](#).

Indexing

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

- Full-text indexing—Each event is tokenized and indexed. See [“Full-text Indexing \(Keyword Indexing\)” on page 81](#).
- Field-based indexing—Event fields are indexed based on a predetermined schema. See [“Field-based Indexing” on page 81](#).

A Logger can have both types of indexing enabled at the same time.

How indexing works

Once you enable indexing on Logger, it starts scanning events automatically and indexing them according to the indexing method you have enabled. You can have both methods—full-text and field-based—enabled at the same time. Once indexing is enabled on Logger, it cannot be disabled.

Events are indexed from the point at which you enable indexing. An event is timestamped with the Logger receipt time when it is received on the Logger. Logger uses the receipt time of an event and the time when a field was added to the index (for field-based indexing) and when full-text indexing was enabled (for full-text indexing) 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 (for field-based indexing) and when full-text indexing was enabled (for full-text indexing), the event is indexed; otherwise, it is not.

The following events are not indexed:

- Existing non-indexed events on a Logger that is upgraded to v4.0.
- Events received on a Logger before indexing was initiated on it.
- Events that are archived.

Full-text Indexing (Keyword Indexing)

For full-text indexing, each event is scanned and divided into keywords and stored on the Logger. **Full-text indexing is not enabled by default**; you are prompted to enable it at initialization time (described at [“Initialization Sequence \(for all Loggers\)” on page 24](#)). Once you do so, Logger automatically indexes incoming events from that point on.

If you do not enable full-text indexing at initialization time, you can do so at any time on the Search Optimization page (**Configuration > Search Optimization**). Once enabled, full-text indexing cannot be disabled. For details about enabling full-text indexing, see [“Enabling Indexing” on page 83](#).

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 is not enabled by default; therefore, no fields exist in the index by default on new Logger; however, it is automatically enabled once you add at least one field to an index.

Although you can add fields to an index at any time, you are presented with a list of recommended fields during the initialization sequence of Logger (described at [“Initialization Sequence \(for all Loggers\)” on page 24](#)). Once you index those fields, indexing is enabled at initialization time. You can add more fields to an index at any time. Once a field has been added, you cannot remove it.



- ArcSight strongly recommends that you index fields that you will be using in search and report queries. Additionally, ArcSight recommends that you index the recommended fields at the initialization time to optimize search performance.
- The `requestUrl` field is available for search and report queries; however, this field cannot be indexed.

Once you enable indexing on a Logger, Logger starts indexing the event metadata fields—event time, Logger receipt time, and device address—for every event in addition to the fields you added to the index. The event metadata fields are also referred to 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 ArcSight recommends to you to add during 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.

Indexable Fields

agentAddress	deviceCustomDate2	flexDate1
agentHostName	deviceCustomDate2Label	flexDate1Label
agentNtDomain	deviceCustomNumber1	filePath
agentSeverity	deviceCustomNumber1Label	flexNumber1

Indexable Fields		
agentType	deviceCustomNumber2	flexNumber1Label
agentZone	deviceCustomNumber2Label	flexNumber2
agentZoneName	deviceCustomNumber3	flexNumber2Label
agentZoneResource	deviceCustomNumber3Label	flexString1
agentZoneURI	deviceCustomString1	flexString1Label
applicationProtocol	deviceCustomString1Label	flexString2
baseEventCount	deviceCustomString2	flexString2Label
bytesIn	deviceCustomString2Label	message
bytesOut	deviceCustomString3	name
categoryBehavior	deviceCustomString3Label	priority
categoryDeviceGroup	deviceCustomString4	requestClientApplication
categoryObject	deviceCustomString4Label	requestContext
categoryOutcome	deviceCustomString5	requestMethod
categorySignificance	deviceCustomString5Label	requestUrlFilename
categoryTechnique	deviceCustomString6	requestUrlQuery
customerName	deviceCustomString6Label	sessionId
destinationAddress	deviceEventCategory	sourceAddress
destinationDnsDomain	deviceEventClassId	sourceHostName
destinationHostName	deviceExternalId	sourceMacAddress
destinationMacAddress	deviceHostName	sourceNtDomain
destinationNtDomain	deviceInboundInterface	sourcePort
destinationPort	deviceOutboundInterface	sourceProcessName
destinationProcessName	deviceProduct	sourceServiceName
destinationServiceName	deviceReceiptTime	sourceTranslatedAddress
destinationTranslatedAddress	deviceSeverity	sourceUserId
destinationUserPrivileges	deviceVendor	sourceUserName
destinationUserId	deviceVersion	sourceZone
destinationUserName	deviceZone	sourceZoneName
destinationZone	deviceZoneName	sourcezoneResource
destinationZoneName	deviceZoneResource	sourceZoneURI
destinationZoneResource	deviceZoneURI	startTime
destinationZoneURI	endTime	transportProtocol
deviceAction	eventId	type
deviceAddress	externalId	vulnerabilityExternalID

Indexable Fields

deviceCustomDate1	fileName	VulnerabilityURI
deviceCustomDate1Label		

Guidelines for Field-based Indexing

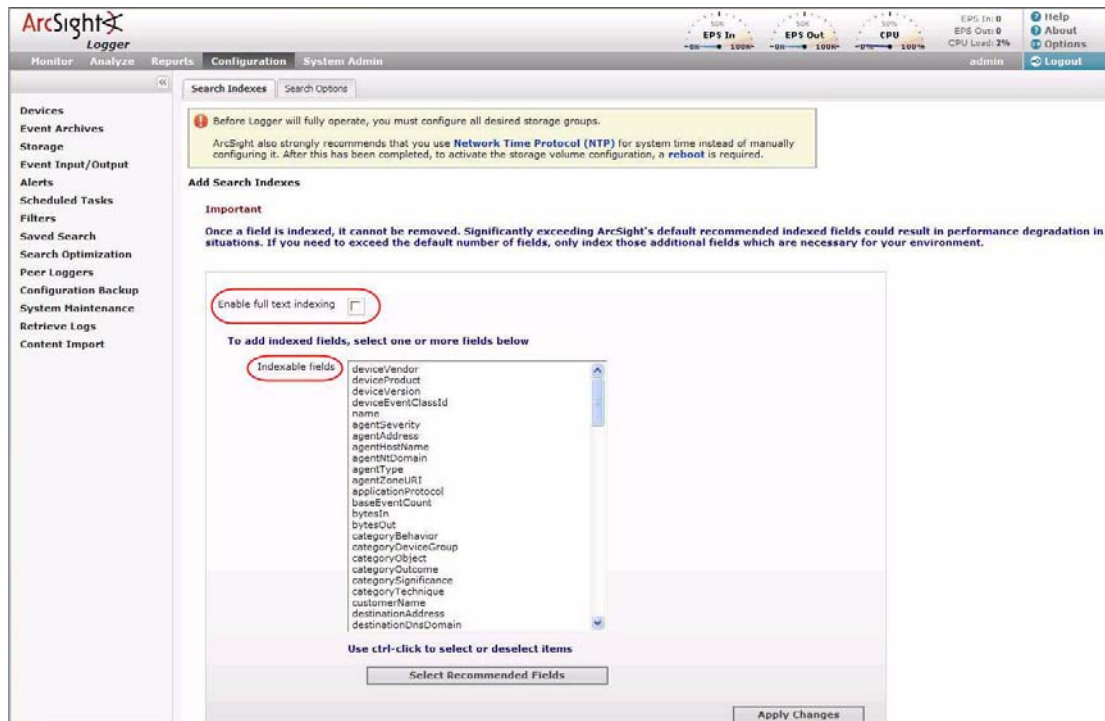
Make sure you are familiar with these guidelines before you index any fields:

- Events are indexed by any fields you add to the index and the default event metadata fields—event time, Logger receipt time, and device address.
- 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.
- Once you initiate indexing on your Logger, it starts indexing events it receives from that point on. Any existing events are not indexed.
- 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 peer 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

If you did not enable indexing on Logger at initialization time (described at [“Initialization Sequence \(for all Loggers\)” on page 24](#)), you can do so using these instructions.

To enable indexing:



1 Click **Configuration > Search Optimization > Search Indexes**.

2 To enable full-text indexing:

a Click **Enable full text indexing**.

3 To enable field-based indexing:

a Select the fields from the Indexable Fields list.

To select multiple fields at the same time, hold the Ctrl key down and click on the fields.

b Click **Add**.

Saving Queries (Saved Filters and Searches)

If you need to run the same search query regularly, you can save it in these ways:

■ As a filter

A Filter saves the query expression, but does not save the time range or the field set information.


■ As a saved search

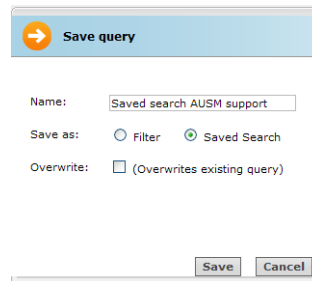
A saved search saves the query expression and the time range that you specified.

For information about Saved Search Alerts, see [“Creating and Managing Saved Search Alerts” on page 212](#).

Saving a Query

To save a query:

- 1 Define a query as described in [“Searching for Events on Logger” on page 72](#) or [“Using the Search Builder Tool” on page 64](#).
- 2 Click the Save icon () and enter a name for the query in the Name field, as shown in the following figure.




The 'Save query' dialog box contains the following fields and options:

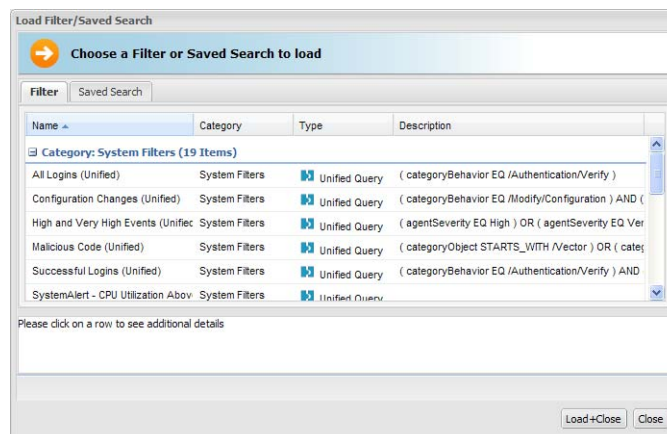
- Name:** Saved search AUSM support
- Save as:** ☐ Filter ☒ Saved Search
- Overwrite:** ☐ (Overwrites existing query)
- Buttons:** Save, Cancel

- 3 Select whether you want to save this query as a filter or as a saved search.
- 4 Click **Save**.

Using a Saved Filter or a Saved Search

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 shows a table of saved filters and searches. The 'Filter' tab is selected.

Name	Category	Type	Description
Category: System Filters (19 Items)			
All Logins (Unified)	System Filters	Unified Query	(categoryBehavior EQ /Authentication/Verify)
Configuration Changes (Unified)	System Filters	Unified Query	(categoryBehavior EQ /Modify/Configuration) AND (
High and Very High Events (Unified)	System Filters	Unified Query	(agentSeverity EQ High) OR (agentSeverity EQ Ver
Malicious Code (Unified)	System Filters	Unified Query	(categoryObject STARTS_WITH /Vector) OR (cate
Successful Logins (Unified)	System Filters	Unified Query	(categoryBehavior EQ /Authentication/Verify) AND
SystemAlert - CPU Utilization Above	System Filters	Unified Query	

Please click on a row to see additional details

Buttons: Load+Close, Close

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 textbox below.

- 3 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. Filters 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 System Alert filters (listed in 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.

	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)
	Configuration Changes	System configuration changes (CEF format)
Event Count use case	High and Very High Events	High and Very High Severity CEF events
	Event Counts by Source	
Malicious Code use case	Malicious Code	Malicious Code (CEF format)
Firewall use case	FW-Deny (Firewall Deny)	
	FW-Drop (Firewall Drop)	
	FW-Permit (Firewall Permit)	
Network use case	Net-DHCP Lease Events	
	Net-Port Links Up and Down	
	Net-Protocol Links Up and Down	
Syslog	Syslog-Failed Logins	

	Unified Query Filters	Regular Expression Query Filters
Connector System Status use case	SysStatus-CPU Utilization by Connector Host SysStatus-Disk Utilization by Connector Host SysStatus-Memory Utilization by Connector Host	
UNIX Server use case	Unix-CRON related events Unix-IO Errors and Warnings Unix-PAM and Sudo Messages Unix-Password Changes Unix-SAMBA Events Unix-SSH Authentications Unix-User and Group Additions Unix-User and Group Deletions	
Windows Events use case	Win-Windows Events (CEF)	
System Alerts	<p>System Alerts: 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 88.</p> <p>NOTE: Although these filters are displayed on the software version of Logger, these do not apply to it.</p> <p>CPU Utilization Above 90 Percent</p> <p>CPU Utilization Above 95 Percent</p> <p>Disk Space Below 10 Percent</p> <p>Disk Space Below 5 Percent</p> <p>Device Configuration Changes</p> <p>Filter Configuration Changes</p> <p>High CPU Temperature</p> <p>Low Fan Speed</p> <p>Power Supply Failure</p> <p>RAID Status Battery Failure</p> <p>RAID Status Disk Failure</p> <p>Storage Configuration Changes</p> <p>Zero Events Incoming</p> <p>Zero Events Outgoing</p> <p>Zero Fan Speed</p>	

Using a System Filter

To use a predefined system filter, follow instructions in [“Using a Saved Filter or a Saved Search” on page 85](#).

Monitoring System Health

You can monitor your Logger’s health in two ways:

- Using a predefined system filter, as listed in [“System Filters/Predefined Filters” on page 86](#). The predefined system health filters are based on the system health events listed in [“System Health Events” on page 88](#).
- Searching for system health events in Logger’s Internal Storage Group, as listed in [“System Health Events” on page 88](#). If a predefined 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 with the desired e-mail address destination (see [“SMTP Settings” on page 258](#)) or create an SNMP Destination (see [“SNMP Destinations” on page 218](#)) or Syslog Destination (see [“Syslog Destinations” on page 219](#)).
- 2 Create an Alert that uses one or more System Alert Filters or defining 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 207](#)).
- 3 Enable the new Alert.

System Health Events

The following table lists all of the system health events that Logger generates. These events are stored in Logger’s Internal Storage Group. The predefined System Filters that provide system health status are based on some of these events. If a predefined filter does not suit you needs, create an alert using one of the following events.

Group	Device Event Category	Device Event Class ID
CPU	/Monitor/CPU/Usage	cpu: 100
Disk	/Monitor/Disk/Space/Remaining/Data	disk: 100
	/Monitor/Disk/Space/Remaining/Root	disk: 101
	/Monitor/Disk/Read	disk: 102
	/Monitor/Disk/Write	disk: 103
EPS	/Monitor/Receiver/All/EPS	eps: 100
	/Monitor/Forwarder/All/EPS	eps: 100
Memory	/Monitor/Memory/Usage/Platform	memory: 100
	/Monitor/Memory/Usage/Jvm	memory: 101
Network	/Monitor/Network/Usage/In	network: 100
	/Monitor/Network/Usage/Out	network: 101

Group	Device Event Category	Device Event Class ID
Raidcontroller	/Monitor/RAIDController/Configuration/RAID-5	raidcontroller: 100
	/Monitor/RAIDController/Port/p0	raidcontroller: 101
	/Monitor/RAIDController/Port/p1	raidcontroller: 102
	/Monitor/RAIDController/Port/p2	raidcontroller: 103
	/Monitor/RAIDController/Port/p3	raidcontroller: 104
	/Monitor/RAIDController/Sensor/bbu	raidcontroller: 105
Search	/Monitor/Search	search: 100
Sensor	/Monitor/Sensor/CPU1	sensor: 100
	/Monitor/Sensor/CPU2	sensor: 101
	/Monitor/Sensor/System	sensor: 102
	/Monitor/Sensor/DIMM	sensor: 103
	/Monitor/Sensor/CPU1Core	sensor: 104
	/Monitor/Sensor/CPU2Core	sensor: 105
	/Monitor/Sensor/3.3V	sensor: 106
	/Monitor/Sensor/5V	sensor: 107
	/Monitor/Sensor/12V	sensor: 108
	/Monitor/Sensor/-12V	sensor: 109
	/Monitor/Sensor/Battery	sensor: 110
	/Monitor/Sensor/FAN1	sensor: 111
	/Monitor/Sensor/FAN2	sensor: 112
	/Monitor/Sensor/FAN3	sensor: 113
	/Monitor/Sensor/FAN4	sensor: 114
	/Monitor/Sensor/FAN5	sensor: 115
	/Monitor/Sensor/FAN6	sensor: 116
	/Monitor/Sensor/FAN7	sensor: 117
	/Monitor/Sensor/FAN8	sensor: 118
	/Monitor/Sensor/PowerSupply	sensor: 119
Storage Group	/Monitor/StorageGroup/Space/Used	storagegroup: 100
	Note: The size of the storage group, indicated by the "fsize" field is in GB.	

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.

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 207](#)), 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 54](#) for more detail.

Receiving Alerts for Events

To receive alerts:

- 1 Configure the Logger’s SMTP with the desired e-mail address destination (see [“SMTP Settings” on page 258](#)) or create an SNMP Destination (see [“SNMP Destinations” on page 218](#)) or Syslog Destination (see [“Syslog Destinations” on page 219](#)).



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 84](#).)

- 3 Create an Alert that uses the new Filter and specify match count and threshold (see [“Alerts” on page 207.](#)) Enable the new Alert.

Monitor Analyze 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

Base Event Fields

Events that are labeled 'Action Engine' are Alert events. Other events are base events--that is, the events which 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 72](#), [“Understanding the Search Results Display” on page 74](#), [“Viewing Alerts” on page 90](#), and [“Advanced Search Options” on page 73](#).

Chapter 5

Reporting

This chapter describes Logger reporting features.

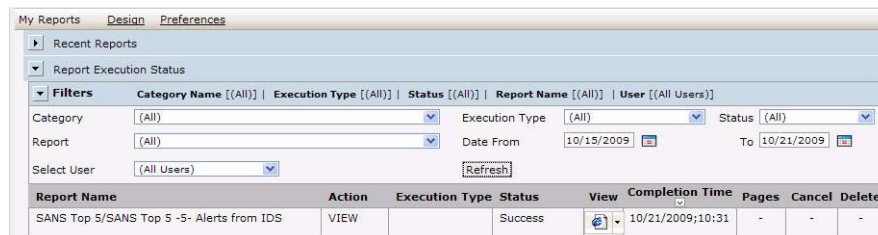
Reports are captured views or summaries of events which you can view from the Logger Reports tab or export for sharing in a variety of file formats. Reporting is an essential tool for communicating the state of your network security to internal and external stakeholders.

["Navigating to Reports" on page 93](#)
["Report Groups" on page 94](#)
["Reports Home Page" on page 99](#)
["Using the Dashboard" on page 100](#)
["Running, Viewing, and Publishing Reports" on page 113](#)
["Designing Reports" on page 124](#)
["Scheduling Reports" on page 173](#)
["Deploying a Report Package" on page 177](#)
["Report Server Administration" on page 178](#)
["Backup and Restore of Report Content" on page 180](#)

Navigating to Reports

To access the Reporting home page, click **Reports** on the Logger navigation bar.

If there is no Dashboard display configured and selected, the Reports home page shows the execution status of recently run or accessed reports as the default view.



Report Name	Action	Execution Type	Status	View	Completion Time	Pages	Cancel	Delete
SANS Top 5/SANS Top 5 -5- Alerts from IDS	VIEW		Success		10/21/2009;10:31	-	-	-

Figure 5-1 Reports Home Page Showing Recently Run Reports (No Dashboard)

If a dashboard is configured to display, the Reports Home page shows the selected **Dashboard** view.



Figure 5-2 Reports Home Page with a Dashboard

For information about designing and selecting dashboard views, see [“Using the Dashboard” on page 100](#).

Report Groups

Logger supports the following report groups:

- [“Foundation Reports” on page 95](#)—This report group contains ready-made reports that address common security use cases. This report group is displayed by default.
- [“Solution Reports” on page 97](#)—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.
- [“Device Monitoring Reports” on page 97](#)—This report group contains ready-made reports that address common device monitoring use cases for systems and devices on your network. For example, top infected systems, failed login attempts, VPN connections denied by address, and so on. This report group is displayed by default.
- [“User Reports” on page 98](#)—This report group contains the custom reports built using the provided tools and templates. This report group is displayed by default.



More Foundation and Solution Reports may be available for download as *report packages* on the Customer Support Web site. (For information about deploying report packages, see [“Deploying a Report Package” on page 177](#).)

These report groups are listed in the left panel menu of the Reports page. Under each report group, the report categories for the report group are listed. For example, under the Foundation Reports report group, the SANS Top 5 report category is listed. Under each report category, a set of reports are listed. For example, the *SANS Top 5 - 1 - Number of Failed Logins* report is listed under the SANS Top 5 report category.

To view reports, click a report category on the Reports page left panel menu.

Foundation Reports

As a starting point for thorough and effective monitoring and compliance, ArcSight Logger provides packages of pre-built reports for common security use cases. These reports are listed in the Foundation Reports report group.



More Foundation Reports may be available for download as *report packages* on the Customer Support Web site. (For information about deploying report packages, see [“Deploying a Report Package” on page 177.](#))

The Foundation Reports include [“SANS Top 5 Reports” on page 95](#), [“Network Monitoring Reports” on page 96](#), [“Intrusion Monitoring Reports” on page 96](#), [“Configuration Monitoring Reports” on page 97](#).

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 173](#).

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 113](#) and [“Task Options on Available Reports” on page 114.](#)

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 Foundation Reports | **SANS Top 5** on the Reports left panel menu.

Category List > SANS Top 5 > Standard							
Standard		Custom					
S.No.	Report Name	Quick Run	Run	Published	Edit	Description	Delete
1.	SANS Top 5 -1- Number of Failed Logins						
2.	SANS Top 5 -1- Top Users with Failed Logins						
3.	SANS Top 5 -2- Failed Resource Access by Users						
4.	SANS Top 5 -2- Failed Resource Access by Users Drilldown						
5.	SANS Top 5 -2- Failed Resource Access Events						
6.	SANS Top 5 -2- Failed Resource Access Events Drilldown						
7.	SANS Top 5 -3- Password Changes						
8.	SANS Top 5 -3- User Account Creations						
9.	SANS Top 5 -3- User Account Deletions						

For information how to run, view, and publish these reports, see [“Running, Viewing, and Publishing Reports” on page 113](#).

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

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

Intrusion Monitoring Reports

Logger provides reports that address intrusion monitoring. To access these reports, click Foundation Reports | **Intrusion Monitoring** on the Reports left panel menu.

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 113](#).

Configuration Monitoring Reports

Logger provides reports that address configuration monitoring. To access these reports, click Foundation Reports | **Configuration Monitoring** on the Reports left panel menu.

For information how to run, view, and publish these reports, see [“Running, Viewing, and Publishing Reports” on page 113](#).

Solution Reports

Solutions Reports are available as add-on packages to Logger for specific compliance requirements or scenarios.



More Solution Reports may be available for download as *report packages* on the Customer Support Web site. (For information about deploying report packages, see [“Deploying a Report Package” on page 177](#).)

For information on deploying Solutions Packages, see [“Deploying a Report Package” on page 177](#). 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**.

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 173](#).

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 113](#) and [“Task Options on Available Reports” on page 114](#).)

Device Monitoring Reports

ArcSight Logger provides packages of pre-built reports for common device monitoring use cases such as top infected systems, failed login attempts, VPN connections denied by address, and so on. These reports are listed in the Device Monitoring Reports group.

The Device Monitoring Reports include [“Anti-Virus Reports” on page 98](#), [“Cross Device Reports” on page 98](#), [“Database Reports” on page 98](#), [“Firewall Reports” on page 98](#), [“Identity Management Reports” on page 98](#), [“IDS-IPS Reports” on page 98](#), [“Network Reports” on page 98](#), [“Operating System Reports” on page 98](#), [“VPN Reports” on page 98](#).

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 173](#).

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 113](#) and [“Task Options on Available Reports” on page 114](#).)

Anti-Virus Reports

These reports 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 Reports | **Anti-Virus** under the Device Monitoring Reports section on the left panel menu.

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 Reports | **Cross Device** under the Device Monitoring Reports section on the left panel menu.

Database Reports

These reports provide information on database errors and warnings.

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.

Network Reports

These reports provide information on activity involving network infrastructure, including interface status, device errors, SNMP authentication failures, and so on.

Operating System Reports

These reports provide information on activity involving operating systems, such as login errors per user, and user and user group creation and modification events.

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.

User Reports

The reports you create and save are displayed in the User Reports pages. Reports with custom-built queries and one or more data sources, typically obtained from ArcSight or other custom developer sources in a *report package* are also listed on this page. If no user reports have been created yet, the report lists on these pages will be blank.

To navigate to user reports, click Reports | **User Reports** on the left panel menu.

Reports > Default Reports								
New Adhoc Report								
S.No.	Report Name	Quick Run	Run in Background	Run	Published	Edit	Description	Delete
1.	Failed Logins							
2.	Intrusion Attempts							

Figure 5-3 User Reports

Some reports, such as the ones obtained from ArcSight or other custom developer sources might not be editable. For such reports, the Edit column icon () is gray, as shown in the following example:

S.No.	Report Name	Quick Run	Run in Background	Run	Published	Edit	Description	Delete
1.	SecurityDBReport							
2.	SecurityDashBoardRpt							
3.	Top Attacker Detail							
4.	Access Events by Resource							
5.	Attack Events by Destination							

For information how to run, view, and publish reports, see [“Running, Viewing, and Publishing Reports” on page 113](#).

For information on using the Report Designer to create reports, see [“Designing Reports” on page 124](#).

For information on deploying Custom Packages, see [“Deploying a Report Package” on page 177](#).

Reports Home Page

If you click the **Reports** tab from elsewhere in the Logger UI, the Reports Home page is displayed. Also, if you click **Dashboard** on the Reports left panel menu from within Reports, the Reports Home page is displayed.

If a dashboard is configured and selected for display, then the Dashboard View page *is* the Reports Home page and the selected dashboard is shown (for example, see [Figure 5-4 on page 101](#)).

To get started by creating a dashboard to show as your default Reports Home page, see [“Using the Dashboard” on page 100](#), [“Designing Dashboards” on page 101](#), and [“Setting Dashboard Preferences” on page 112](#).

If no dashboard is configured and selected for display, the default Reports home page shows the **Report Execution Status** page that lists the status of currently running, recently run, or accessed reports, as shown in the following figure. By default, all reports except the completed scheduled reports are displayed, however, you can restrict the list by defining filter criteria. (To view completed scheduled reports, click **Configuration > Scheduled Tasks > Finished Tasks > Filter by Job Type/Report.**)

If a report is in run in the background, the Execution Type column indicates it. Otherwise, the column is left blank.

The screenshot shows the 'Report Execution Status' window. It includes a 'Filters' section with dropdowns for Category, Report, and Select User, and date pickers for Date From and To. Below the filters is a table with the following data:

Report Name	Action	Execution Type	Status	View	Completion Time	Pages	Cancel	Delete
SANS Top 5/SANS Top 5 - Alerts from IDS	VIEW		Success		10/21/2009;10:31	-	-	-

To get started by running and viewing reports, see [“Running, Viewing, and Publishing Reports” on page 113](#) and [“Scheduling Reports” on page 173](#).

Using the Dashboard

Dashboards display reporting data to provide a quick view of the latest information about network events. You can assemble various reports, common network monitoring use cases, and external links onto a dashboard to provide network status at-a-glance.

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 view. 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.

Viewing the Dashboard

If no dashboard is configured and selected for display, the default Reports home page shows the **Report Execution Status** page that lists the status of recently run or accessed reports, as shown in the following figure. In this case, clicking on **Dashboard** in the Reports left panel menu will show only the Report Execution Status list. By default, all reports except the completed scheduled reports are displayed. (To view completed scheduled reports, click **Configuration > Scheduled Tasks > Finished Tasks > Filter by Job Type/Report.**)

If a dashboard is configured and selected for display, it is shown on the Dashboard **View** 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 **View** (Reports Home page).

The Dashboard View page displays the contents of various items placed on the dashboard during design time. If the dashboard includes reports, reports will show current data from recently run reports.

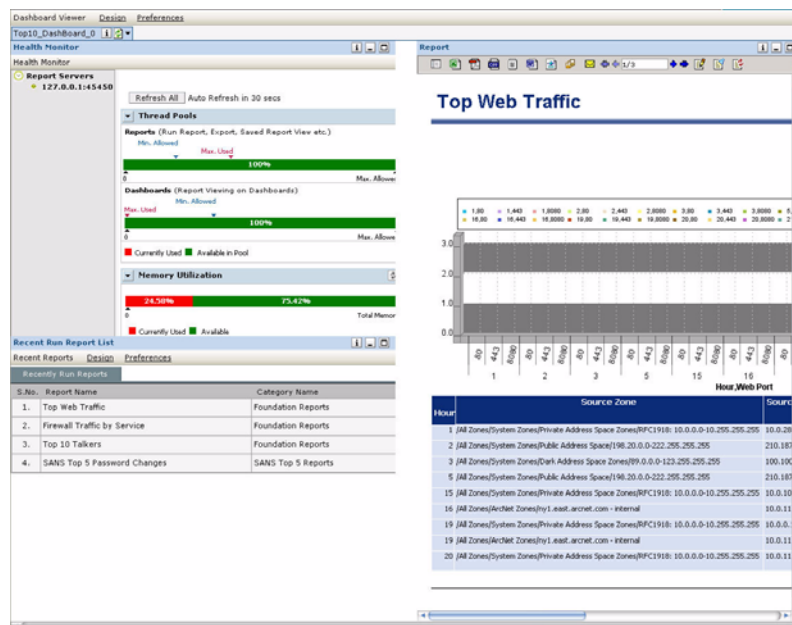


Figure 5-4 Dashboard View with Health Monitor and Reports



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. For more information, see [“Designing Dashboards” on page 101](#) and [“Scheduling Reports” on page 173](#).

To run a report manually, click User Reports and select (Run), (Quick Run), or (Run in Background) button, set the parameters, and click **Run** or **Run Report**, respectively. For more information on running and publishing reports, see [“Running, Viewing, and Publishing Reports” on page 113](#).

Designing Dashboards

Use the **Dashboard Designer** 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 **View** 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 Reports Dashboard Designer, click **Design** on the Dashboard navigation bar.

Click "Design" to open
Dashboard Designer

Recent Reports **Design** Preferences



Dashboards are optional. If you do not create at least one dashboard and select it for display, then the Reports home page simply shows a list of recently run reports by default.

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

Quick Start - 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.


- 1 Add a new, empty dashboard.

To do this navigate to **Dashboard > Design** on the Reports menu bar, and click  (Click here to create dashboard) on the Dashboards list title bar in upper left. This brings up a dashboard with an empty layout.

- 2 Under Dashboard Properties, specify a **Name** for the new dashboard and other dashboard properties, as needed.



- 3 Place items onto the dashboard in the **Widgets** provided in the Layout area.

To do this, 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.

- 4 For each item (widget) placed, specify Widget Properties, as needed.

- 5 To automatically set the dashboard as the default Dashboard View, click (check) **Add to my preferred list**.
- 6 Click **Save** to save the dashboard.


**Note**

Once saved, new dashboards become available in the **Dashboard > Preferences** list of "Available Dashboard(s)".

See ["Selecting a Dashboard View" on page 112](#) for information on how to display the new dashboard you just created or set the default display to a different dashboard.

Add an Empty Dashboard

Dashboards are created on the Dashboard Design page.

- 1 On the Reports menu bar, navigate to **Dashboard > Design**, and click the Add button  on the Dashboards list title bar in upper left.



This brings up a dashboard with an empty layout.

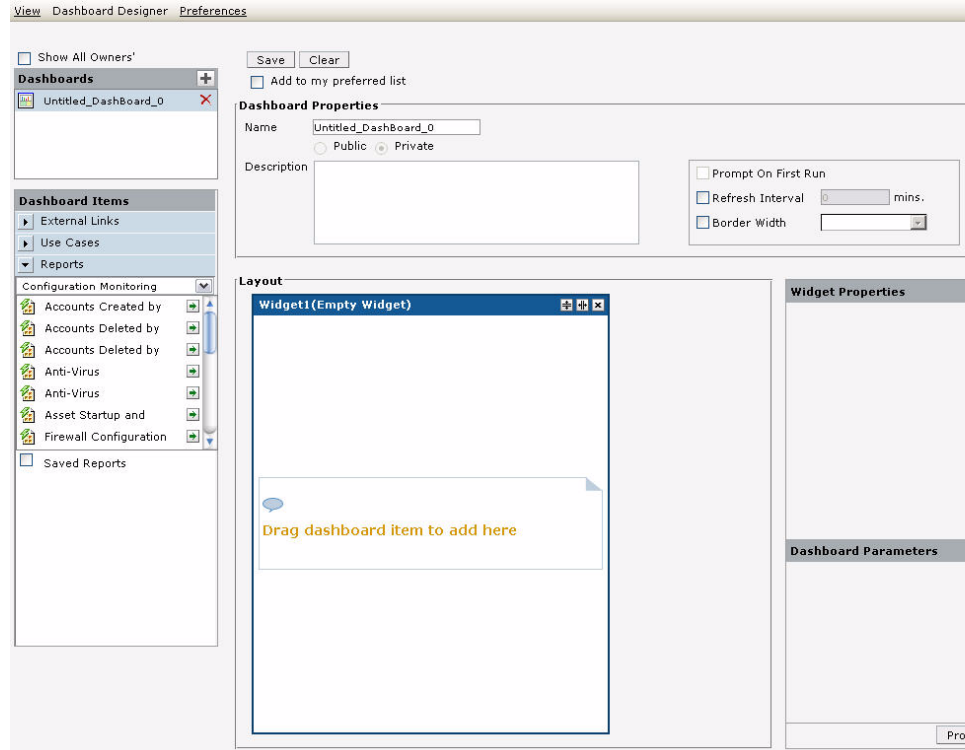


Figure 5-5 New Dashboard Layout

- 2 In the **Name** field, specify a unique name to identify the dashboard.
- 3 Specify Dashboard Properties. (For details, see [“Dashboard Properties” on page 104.](#))
- 4 Click **Save** button.

The new dashboard name is added to the list of Dashboards on the Design page, and also to the list of Available Dashboard(s) on the Preferences page.

- 5 To place the dashboard in the list of Selected Dashboard(s) and set it as the default view, click **Add to my preferred list**. Keep in mind that only one dashboard at a time can be displayed as the default view. (The “default view” Selected Dashboard can also be set on the Dashboard Design Preferences page, as described in [“Selecting a Dashboard View” on page 112.](#))



Clicking **Clear** erases the selected dashboard and gives a clean slate on which to start over. This clears the layout area, dashboard parameters if any, and widget properties.

Dashboard Properties

Dashboard Properties

Name:

☐ Public ☒ Private

Description:

☐ Prompt On First Run

☐ Refresh Interval: mins.

☐ Border Width:

Figure 5-6 Reports Dashboard Properties

The Dashboard Properties are described in the following table.

Table 5-1 Dashboard Properties Description

Property	Description
Name	Name of the dashboard.
Description	Descriptive information about this dashboard.
Refresh Interval	Sets the time in minutes to refresh results for all the reports on dashboard. 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 173.)
Border Width	Check (select) the checkbox to have border around the widgets of the dashboard and select the width from drop-down box.



Creating Widgets

When a new dashboard is created, it has one widget on the layout. Each dashboard item must be placed in its own widget for display on the dashboard.

To get a new widget, simply split the existing widget either vertically or horizontally, depending on the layout you want. (See [“To get a new widget” on page 105.](#))

You can also delete widgets you do not need. (See [“To remove a widget” on page 105.](#))

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  (Remove Widget) on the widget you want to remove.

Placing Dashboard Items on the Layout

Reports, use cases, and external link objects are available under “Dashboard Items” (to the left of the Layout area).



Figure 5-7 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.



Keep in mind that 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.

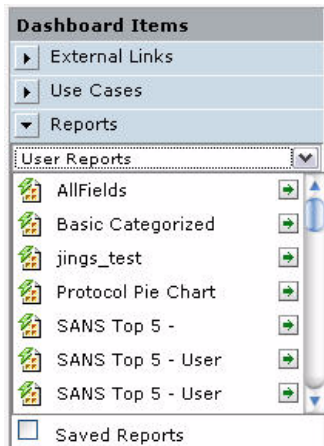



Figure 5-8 Reports under Dashboard Items

- 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. (See [“Widget Properties for Reports” on page 106.](#))

Widget Properties for Reports


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 5-9 Widget Properties for Reports on a Dashboard

The following table describes Widget Properties settings for Reports dashboard items.

Table 5-2 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 173.)</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 173.)</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.

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 108 .
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 5-10 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. (This 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 Widget 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.

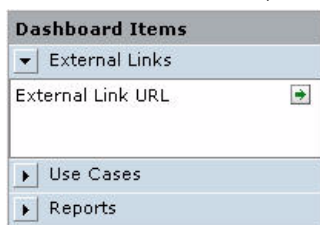
- 5 Click **OK** to save the settings and close the dialog.


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 5-11 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. (See [“Widget Properties for Use Cases” on page 109.](#))

Widget Properties for Use Cases

 A screenshot of the 'Widget Properties' dialog box. It has a title bar 'Widget Properties' and several fields: 'Name' (text input with 'Health Monito'), 'Refresh Interval (in mins.)' (text input with '15'), 'Auto Refresh' (dropdown menu with 'YES' selected), 'Show Scrollbar' (dropdown menu with 'NO' selected), and 'Description' (text area).
Figure 5-12 Widget Properties for Use Cases on a Dashboard

The following table describes Widget Properties settings for Use Case dashboard items.

Table 5-3 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.

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.

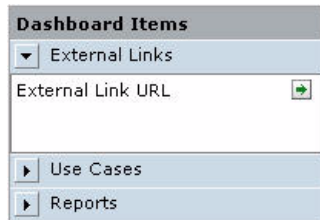



Figure 5-13 External Link under Dashboard Items

- 2 Click and drag a 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. (See ["Widget Properties for External Links" on page 110.](#))

Widget Properties for External Links

Widget Properties	
Name	External Link
Refresh Interval (in mins.)	15
Auto Refresh	YES <input type="checkbox"/>
Show Scrollbar	NO <input type="checkbox"/>
URL	www.arcsight.com
Description	

Figure 5-14 Widget Properties for an External Link on a Dashboard

The following table describes Widget Properties settings for External Links dashboard items.

Table 5-4 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 5-15 Swapping Widgets on a Dashboard Design

In the above example, the Recent Run Reports List item is swapped to the position of the 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 Dashboard Preferences

In Dashboard Preferences, you can specify:

- The dashboard to be made available for viewing
- Decide how dashboards are to be displayed

To navigate to Dashboard **Preferences**, click **Dashboard** on the left panel, then click **Preferences** in the navigation sub-menu at the top.

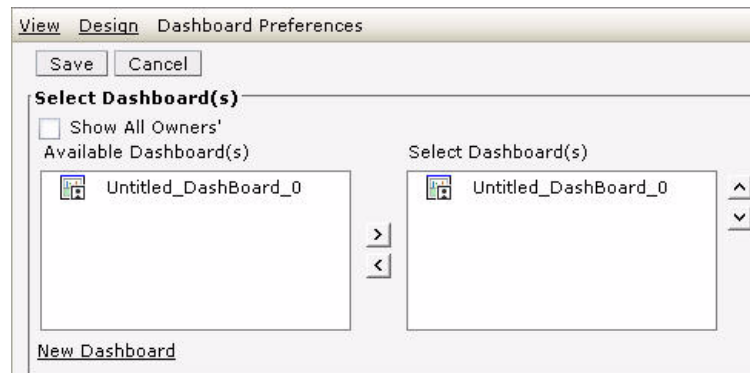


Figure 5-16 Dashboard Preferences

Working with Available Dashboards

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.

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.

Also, 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.

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.



You must have at least one dashboard in order to set a preference for the Dashboard View. For a quick summary of steps to create a dashboard, see ["Quick Start - Creating a New Dashboard" on page 102.](#)

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.



Only one dashboard at a time can be displayed as the default 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 [Step 5](#) in "Quick Start - Creating a New Dashboard" on page 102.

- 3 Click **Save** to save your preferences and display the selected dashboard.

To remove or change the currently displayed dashboard

- 1 Return to the Dashboard **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.


To start designing a new dashboard

To create a new dashboard, click the **New Dashboard** link. This opens a new, empty dashboard in the Dashboard Designer. (This is another way to start designing a new

dashboard, as an alternative to clicking  on the Dashboards list in the designer). For full detail on creating a new dashboard, see ["Designing Dashboards" on page 101](#).

Modifying or Removing Existing Dashboards

To edit existing dashboards, navigate to the Dashboard Designer (**Dashboard > Design**).

- To modify an existing dashboard, select one of the dashboards under **Dashboards** list on the left side. It's current configuration is displayed in the Layout panel, Widgets, and so forth, and you can modify then save settings as needed.
Follow the procedures for working with layout, widgets, and dashboard items described in ["Designing Dashboards" on page 101](#).
- To delete a dashboard, click  (Click here to delete the dashboard) next to the dashboard you want to remove.

Running, Viewing, and Publishing Reports

Reports are deployed (made available) under their respective categories. (See ["Report Groups" on page 94](#))

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.

You can run a report on-demand from any of the reports categories and from **Scheduled Reports** lists.

You can also run a report from the “Recent Reports” list displayed as the default Reports home page on Loggers for which no dashboard is implemented.



There are no options available to *run* reports from a Dashboard view. On a Dashboard view, you can *view* saved or published reports but not run them.

Best Practices

ArcSight 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.

ArcSight 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 173](#).

Finding Reports

You can find reports on the following pages within the Logger **Reports** tab:

- The Foundation Reports, Device Monitoring, User Reports, and Solution Reports groups contain report categories that provide lists of reports. If you are looking for a published version of one of those reports, click into one of those lists. (See [“Report Groups” on page 94](#).)
- You can set a Dashboard View to include “Use Cases” such as “Saved Report List” or “Recent Run Report List”. (See [“Placing a Use Case on a Dashboard” on page 108](#).) If you have one of these lists displayed on a dashboard and you know the report is published, you can find it on the dashboard.
- If the report you are looking for is a scheduled report and it’s been run and published, you can find it in the Scheduled Reports list. (See [“Scheduling Reports” on page 173](#).)



The Search feature on the Logger “Analyze” page (described in [Chapter 4, Searching and Analyzing Events, on page 43](#)) does not search on resources such as reports. It searches only on events in the database.

Task Options on Available Reports

Standard		Custom					
S.No.	Report Name	Quick Run	Run	Published	Edit	Description	Delete
1 .	Accounts Created by User Account						
2 .	Accounts Deleted by Host						
3 .	Accounts Deleted by User Account						
4 .	Anti-Virus Updates-All-Failed						
5 .	Anti-Virus Updates-All-Summary						
6 .	Asset Startup and Shutdown Event Log						
7 .	Firewall Configuration Changes						
8 .	Firewall Configuration Events						
9 .	Firewall Misconfigurations						

Figure 5-17 Task Options on All Reports

The following task options are provided for reports in all categories.








Note

Your access to various reports and report options (view, publish, edit, etc.) depends on the access rights associated with your user role 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 Logger **System Admin** page.

For more information on Logger Report Group management, see [“Setting Access Rights on Reports” on page 140](#), and [“Groups” on page 284](#) in [“System Admin” on page 251](#).

Table 5-5 Task Options on Reports

Button	Description
Quick Run 	<p>Runs the report using default data filtering configuration, which was set at report deploy time.</p> <p>Provides options to change start and end time parameters, storage groups, and devices included in the scope of the report run.</p> <p>See also “To run and view a report” on page 116 and “Quick Run / Run In Background Report Parameters” on page 117.</p>
Run in Background 	<p>Use this option to run reports that take long time to generate or the ones that are not required online immediately.</p> <p>See also “To run and view a report” on page 116 and “Quick Run / Run In Background Report Parameters” on page 117.</p>
Run 	<p>Provides options to modify the data filter criteria used by the report query for this run.</p> <p>You can specify a maximum number of rows to include in the report, and perform various comparison and logical operations on event fields.</p> <p>See also “To run and view a report” on page 116 and “Run Report Parameters” on page 119.</p>
Published 	<p>Displays the list of previously-generated reports that are not yet expired. You can view the user (user name) who generated the report, generate time, and expiry time of the report.</p> <p>The report can be viewed as well as deleted from the saved report list.</p> <p>See also “Viewing the Output of a Published Report” on page 124, “Quick Run / Run In Background Report Parameters” on page 117, and “To publish a report” on page 121.</p>
Edit 	<p>Opens the Report Designer for the associated report, where you can make changes to the underlying query the report uses.</p> <p>See also “Editing a Report” on page 138.</p>
Description 	<p>Description of the report specified at report deployment time.</p>
Delete	Delete a report.

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 from the Reports page left menu, and then choose a report within the category.




For more information about available reports, see [“Foundation Reports” on page 95](#), [“Solution Reports” on page 97](#), and [“User Reports” on page 98](#).

About the Pagination of Reports

If a report contains more columns than can be displayed horizontally across a screen using the default width specified in the report query (Reports > Design > Queries), 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 (Reports > Design > Queries).

To run and view a report

- 1 Click a report category in the left menu and select  (Run Report),  (Run in Background), or  (Quick Run) button next to the report you want to run.
- 2 Set the parameters, and click **Run Now** or **Run in Background**, depending on the report run option you selected in the previous step.


Note: Even if you selected Run Report in the previous step, 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).

Top 10 Talkers			09/28/2007 4:00 PM
Source Zone Name	Source Address	Count	
/All Zones/System Zones/Private Address Space Zones/RFC1918: 10.0.0.0-10.255.255.255	10.210.16.39	3417	
/All Zones/System Zones/Private Address Space Zones/RFC1918: 10.0.0.0-10.255.255.255	10.210.14.124	1618	
/All Zones/System Zones/Private Address Space Zones/RFC1918: 10.0.0.0-10.255.255.255	10.0.111.202	1197	
/All Zones/System Zones/Private Address Space Zones/RFC1918: 10.0.0.0-10.255.255.255	10.0.111.160	620	
/All Zones/System Zones/Public Address Space/198.20.0.0-222.255.255.255	216.136.56.28	93	
/All Zones/ArcSight System/Public Address Space Zones/RIPE NCC 2	83.97.81.58	87	
/All Zones/Site Zones/Alliant Corporate/Alliant/Alliant ICN	142.134.157.38	46	
/All Zones/System Zones/Public Address Space/60.0.0.0-69.255.255.255	68.98.158.138	39	
		19	
/All Zones/System Zones/Public Address Space/60.0.0.0-69.255.255.255	66.26.208.187	15	



Figure 5-18 Results of a Report Run

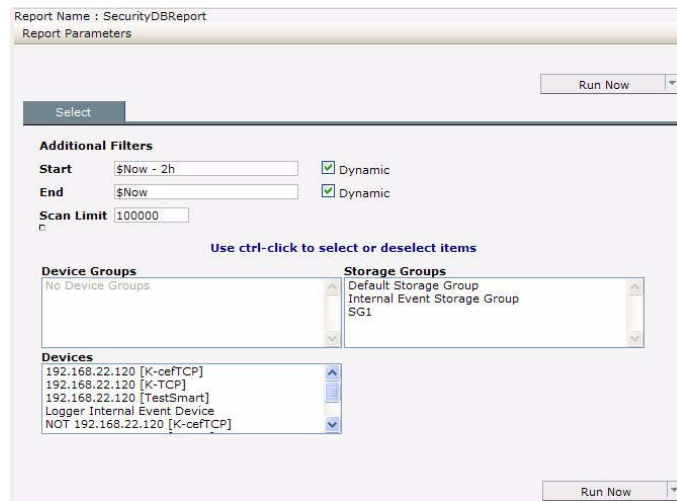
At this point, the results of this report generation is available as a file for viewing only by you. If you close the file without saving or publishing it, the results are no longer available.

If you want to make the results of this run available for others, publish it. To do this, leave the file open, click  (Publish report), and follow the steps in [“Publishing Reports” on page 121](#).

For information about other delivery options available to you at this point, see [“Report Delivery Options” on page 122](#).

Quick Run / Run In Background Report Parameters

When you click or  (Quick Run) or  (Run in Background) 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 storage groups over which the report runs.


Figure 5-19 “Quick Run” / “Run in Background” Report Parameters

The following table describes Quick Run / Run in Background report parameters.

Table 5-6 “Quick Run” / “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>

Option	Description
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	Select the device group(s) on which to run the report query, if any. (See “Selecting Device Groups, Storage Groups, or Devices” on page 118.)
Storage Groups	Select the storage group(s) on which to run the report query. (See “Selecting Device Groups, Storage Groups, or Devices” on page 118.)
Devices	Select the device(s) on which to run the report query. (See “Selecting Device Groups, Storage Groups, or Devices” on page 118.)

Selecting Device Groups, Storage Groups, or Devices

The following figure shows how to select or de-select items on Device Groups, Storage Groups, or Devices as a part of setting Report “Quick Run” and “Run in Background” parameters.

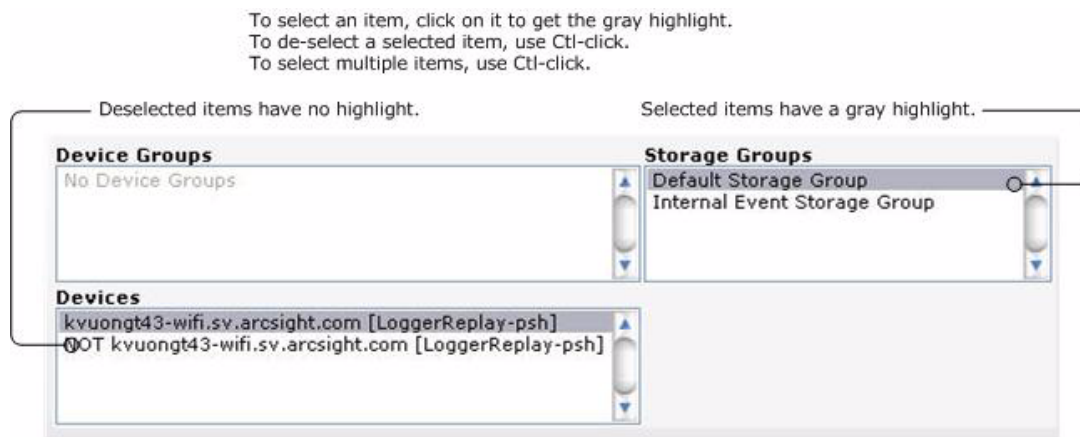



Figure 5-20 Selection Model for “Quick Run” or “Run in Background” Report Scope of Storage and Devices

- Items with a gray 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 Ctl-Click.
- To de-select a currently selected item, use Ctl-Click.
- If none of the items are selected, all items are included in the report query.
- The selected items in the Device Groups and the Devices lists are ORed in the report query, and these items are ANDed with the other selected items such as Storage Groups.

Run Report Parameters

When you click  (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.

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, as indicated in the following figure.

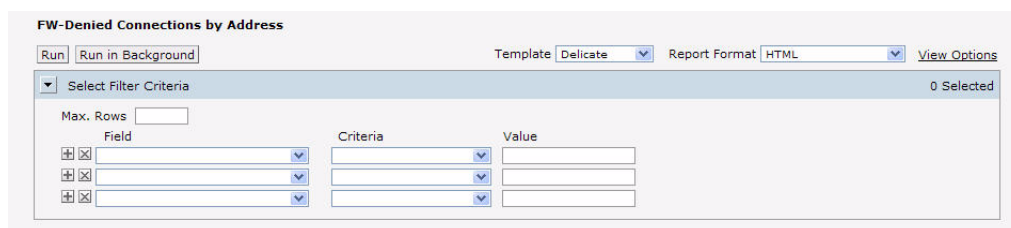


Figure 5-21 “Run” Report Parameters

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 multi-page option if applicable to the chosen file format.</p> <p>ArcSight strongly recommends using the multi-page option for all reports. This option is the default.</p> <p>For descriptions of report format see “Report File Formats” on page 120</p>

Option	Description
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 username(s) or involving specified IP address(es).</p> <p>For details on how to create these filters (with Field, Criteria, and Value fields), see “Select Filter Criteria” on page 130 in “Designing New Reports” on page 128.</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>

When you click **Run** on this first “Parameters” dialog, you then get the same dialog you get for a Quick Run (or Run in Background) report where you can specify filters on timeframe and storage groups on which to run the report. (See [“Quick Run / Run In Background Report Parameters” on page 117](#) for details on this “Select Additional Filters” dialog. Click **Run Report** on this second dialog runs the report.

Report File Formats

Report file formats include:

- HTML (Web page format)
- PDF (Adobe PDF)
- Microsoft Excel
- Comma Separated (Delimiter separated file. The delimiter is usually a comma.)
- Microsoft Word
- Interactive
- XML

For most formats, you can select Multipage option. ArcSight 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 140](#) 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 116](#)), the output results for that run of the report are saved for subsequent.



Tip

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 173](#) and [“Scheduling Reports” on page 173](#) and [“Add Report Job Settings” on page 176](#).

To publish a report

- 1 In a generated report output file you get from running a report, click  (Publish report) 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.



File Name:

Expires on:  hh:mm

(Blank date stands for never expires)

☒ Public ☐ Private

Figure 5-22 Publish Report Settings

The following table describes the publish report options.

Table 5-8 Publish Report Settings

Option	Description
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).
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 124](#).

Report Delivery Options

When you run a report (as described in [“Running and Viewing Reports” on page 116](#)), many options are available to you in terms of delivery options for generated output.

The most common next step is to publish the resulting report (described in [“Publishing Reports” on page 121](#)), 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.

Refreshing a Report

To re-run the report and get an updated result set, click  (Refresh).


E-mailing a Report

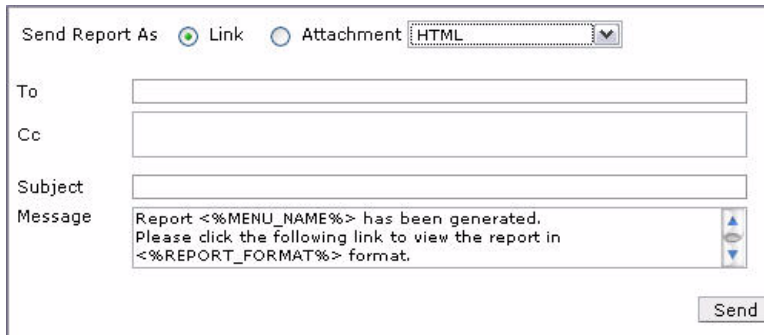
You can send a report via e-mail as either a Web link or an attachment.



You can also configure these same e-mail options on *scheduled reports*, as described in [“Scheduling Reports” on page 173](#) and [“Add Report Job Settings” on page 176](#).

To e-mail a report

- 1 Click the  (Email report) button.
- 2 Specify the following information about the e-mail.



Send Report As ☒ Link ☐ Attachment HTML

To

Cc

Subject

Message

Send

Figure 5-23 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	Chose 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.
To and CC	Specify e-mail addresses to which to send the report.
Subject	Provide e-mail Subject header.







Option	Description
Message	For the body of the e-mail, either use the default message provided, modify it, or enter your own message.

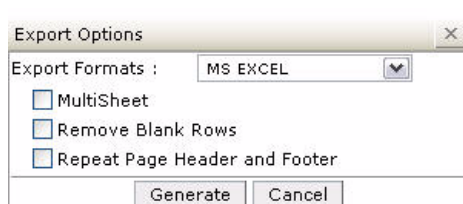
- 3 Click **Send** to send the report.

Exporting and Saving a Report

You can export a report to a file format of your choice and save it.

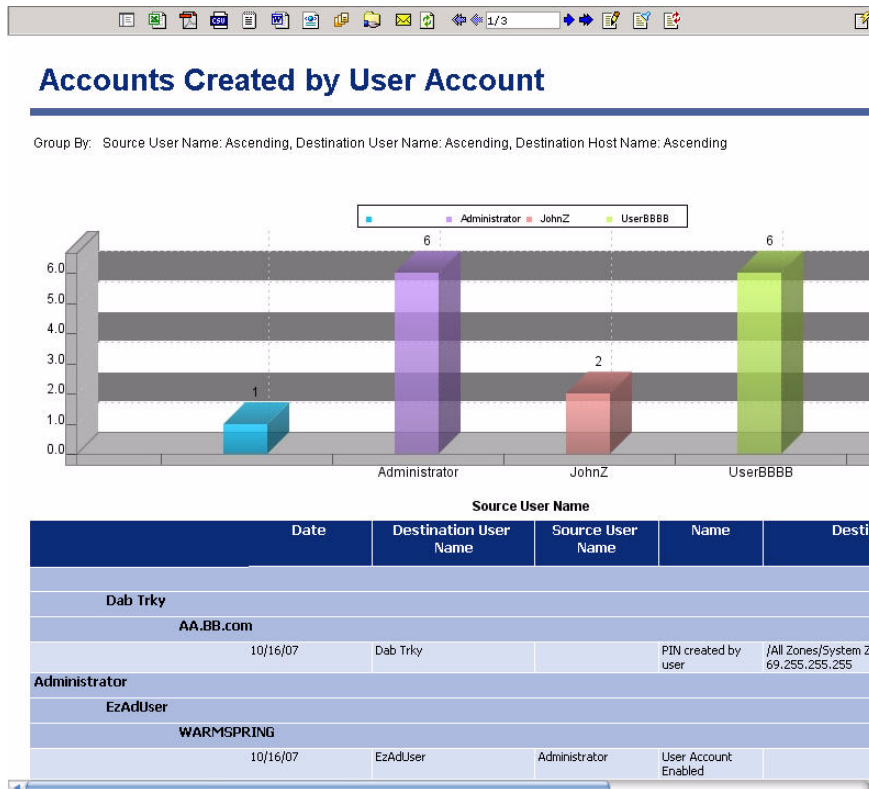
To export and save a report

- 1 Click the  (Export) button or click one of the file formats 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**.

- 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

- Navigate to the report for which you want to view output results. (See [“Finding Reports”](#) on page 114 if you need help locating a report.)
- Click the “Published” button  (Navigate to list of published outputs for this report) next to the report you are interested in.

Saved Report List : Top 10 Talkers

S.No.	File Name	Generated By	Generated Time	Expiry Time	View	Delete
1.	Top 10 Talkers	admin	09/28/2007;16:00 RECENT	10/05/2007;24:00:00		

Figure 5-24 List of Published Report Outputs for a Selected Report

From this dialog you can select various options on any of the listed reports, including options to:


- View report outputs in various formats (HTML, PDF, Microsoft Word, and so on)
- 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 Edit button  for a report in a reports list. (See [“Report Groups” on page 94](#) and [“Task Options on Available Reports” on page 114](#) for more information on available reports and how to get to their task option buttons, respectively.

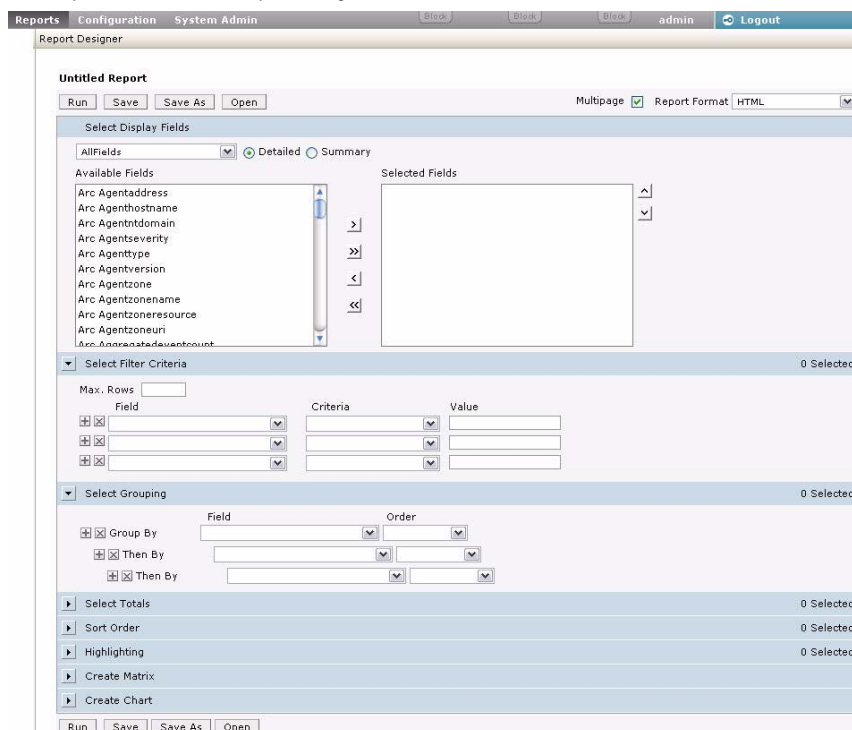


Figure 5-25 Report Designer (click **New Report** or edit an existing report)

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 125](#).

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 128](#).

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.



Modifications to reports and other ArcSight-defined content may be overwritten without warning when the content is upgraded. It is not a 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.

To create a new report based on an existing report

- 1 Navigate to the report you want to use as a starting point. (See [“Report Groups” on page 94](#) for an overview of available reports.)

- 2 Click the Edit button (Customize report) for a report in a reports list.

This opens the report in the Report Designer.

Note: Some reports, such as the ones obtained from ArcSight or other custom developer sources might not be editable. For such reports, the Edit column icon () is gray, as shown in the following example:

S.No.	Report Name	Quick Run	Run in Background	Run	Published	Edit	Description	Delete
1.	SecurityDBReport							
2.	SecurityDashBoardRpt							
3.	Top Attacker Detail							
4.	Access Events by Resource							
5.	Attack Events by Destination							

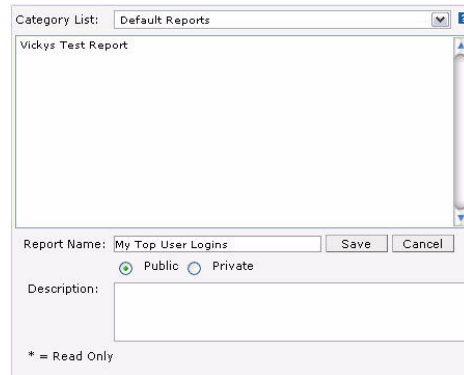
- 3 In the Report Designer for the selected report, click **Save As**.

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 5-26 Save Report Layout As dialog for an Existing Report

- 4 In the Category List at the top of the dialog, select **Default Reports** as the location where you want to save the copy of this report.

Choosing Default Reports provides a view of the reports in that category.



The dialog box titled 'Save Report Layout As' shows the 'Category List' set to 'Default Reports'. The report name is 'My Top User Logins'. The 'Public' radio button is selected, and the 'Private' radio button is unselected. The 'Description' field is empty. A legend at the bottom indicates '* = Read Only'.

Figure 5-27 Save Report Layout As dialog for an Existing Report

- 5 Provide a Report Name for your new report (in the example, we named the report My Top User Logins).


Also select **Public** (if you want everyone to have access to the report) or **Private** (to make the report available only to you), and add a Description, if needed.

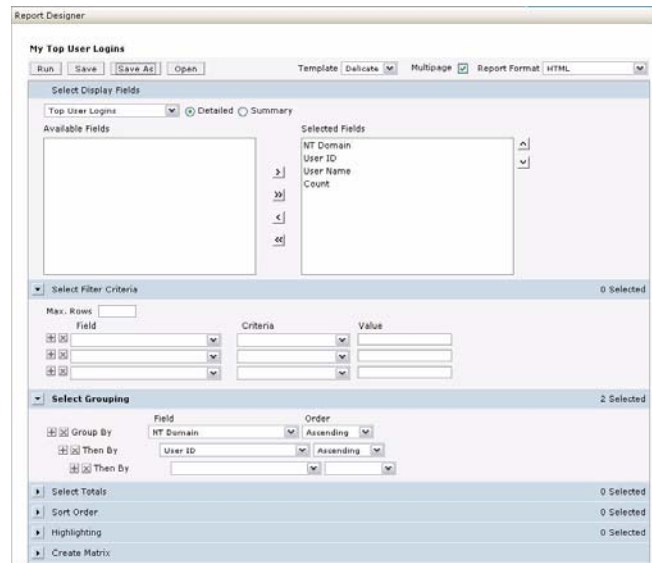
- 6 Click **Save** to save the report.

Click **OK** on the confirm dialog telling you that the report was saved successfully.

- 7 On the left menu under User Reports, click **Default Reports**.

Your new report is shown in the right panel.

- 8 Click the Edit button  (Customize report) to start modify the new report to suit the a specific scenario. (See the next section, ["Designing New Reports" on page 128.](#))



The 'Report Designer' window shows the 'My Top User Logins' report. The 'Select Display Fields' section has 'Top User Logins' selected. The 'Available Fields' list is empty, and the 'Selected Fields' list contains 'NT Domain', 'User ID', 'User Name', and 'Count'. The 'Select Filter Criteria' section shows '0 Selected'. The 'Select Grouping' section shows '2 Selected' with 'Group By' set to 'NT Domain' and 'Order' set to 'Ascending'. The 'Then By' section is empty. The 'Select Totals', 'Sort Order', 'Highlighting', and 'Create Matrix' sections are all set to '0 Selected'.

Figure 5-28 Editing a Report

Designing New Reports

To access the Report Designer to create a new report from scratch, do one of the following:

- Click Design | **New Report** on the Reports page left panel menu.
- On the list of **User Reports** | **Default Reports**, click the New Adhoc Report button



This brings up the Report Designer with a blank template.

The screenshot shows the 'Report Designer' window with a title bar containing 'Reports', 'Configuration', 'System Admin', and user controls. The main area is titled 'Untitled Report' and includes buttons for 'Run', 'Save', 'Save As', and 'Open'. A 'Multipage' checkbox is checked, and the 'Report Format' is set to 'HTML'. The 'Select Display Fields' section has a dropdown for 'AllFields' and radio buttons for 'Detailed' (selected) and 'Summary'. Below this is a list of 'Available Fields' including Arc Agentaddress, Arc Agenthostname, Arc Agentntdomain, Arc Agentseverity, Arc Agenttype, Arc Agentversion, Arc Agentzone, Arc Agentzoneame, Arc Agentzoneresource, Arc Agentzoneuri, and Arc Agentzoneusername. To the right is a 'Selected Fields' list. The 'Select Filter Criteria' section shows '0 Selected' and a table with columns for Field, Criteria, and Value. The 'Select Grouping' section also shows '0 Selected' and options for Group By, Then By, and Order. At the bottom are sections for 'Select Totals', 'Sort Order', 'Highlighting', 'Create Matrix', and 'Create Chart', all showing '0 Selected'. A final row of 'Run', 'Save', 'Save As', and 'Open' buttons is at the very bottom.

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 **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.</p> <p>See “Applying Report Template Styles” on page 172 for more information on working with templates.</p>

Option	Description
Report Format	Select the default format for the report. For information on available formats, see “Report File Formats” on page 120 .
View Options	Select whether report should be Multipage (to split a longer and wider report in multiple pages).

Select Display Fields (Base Query and Fields)

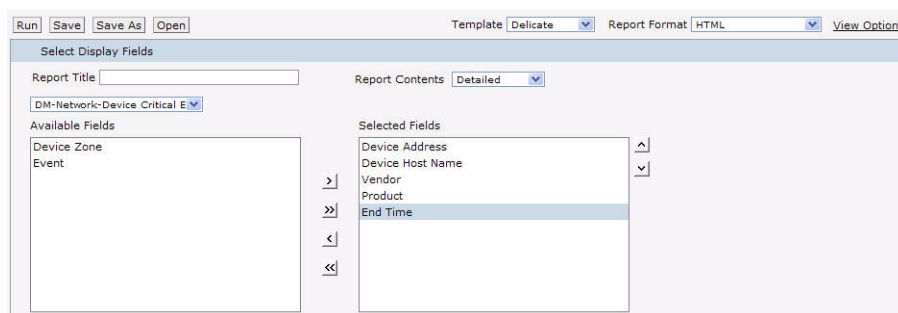


Figure 5-29 Report Display Fields

Each report is built on a base query. Available queries are provided in the drop-down menu under “Select Display Fields” on the Report Designer. 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. (For information on building new queries, see [“Setting up Queries” on page 140](#).)



Note


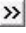

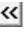


In addition to the fields in the WHERE clause of the query, the fields in the SELECT clause also need to be indexed to yield faster report generation. For more information about indexing fields, see [“Indexing” on page 80](#).

Under **Select Display Fields**, 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 in 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 140](#). 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



Field	Criteria	Value
+ -		
+ -		
+ -		

Figure 5-30 Report Filter Criteria

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 username(s) or involving specified IP address(es).

Also, you can limit the number of rows in a report by defining a Max. Rows value.

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 119](#).





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 149](#) under [“Setting up Queries” on page 140](#).

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> <p>Notes:</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 132), 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. You can, however, 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.

Option	Description
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>Notes:</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 (e.g., <i><Field> Between 90 and 100</i>)</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 149 under "Setting up Queries" on page 140.
Criteria	Select a logical operator. (For example, Is, Is Not, Starts With, Ends With, Contains, and so forth.)
Value	Select a value to complete the conditional filter expression.

Select Grouping

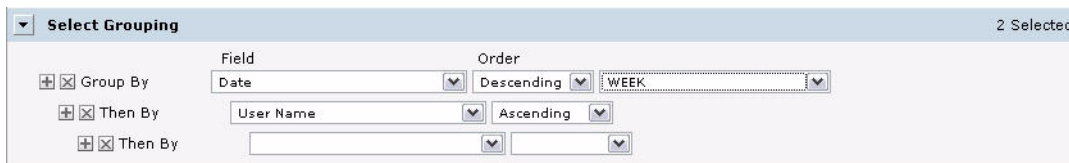


Figure 5-31 Grouping Items by Field in a Report

Define group requirements to arrange the report information into logical groups based on particular fields you are interested in. 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.





Note

If you select set Max. Rows under **Select Filter Criteria** (as described in "[Select Filter Criteria](#)" on [page 130](#)) and also specify grouping, you may get a different result than if you just specified Max. Rows without grouping.

To define a group

- 1 Select a field by which you want to group (as described in [Table 5-20 on page 176](#)).
- 2 Select the order of arrangement of group (as described in [Table 5-20 on page 176](#)).

Table 5-12 Select Grouping Options

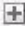

Option	Description
Field	<p>The Fields will be populated with event data fields specified in the base query.</p> <p>Select a field by which to create a group.</p> <p>To add another field for a grouping, click  (Add Group).</p> <p>To remove a group-by field, click  (Remove Group).</p>
Order	<p>Select the order of arrangement of group:</p> <ul style="list-style-type: none"> • Ascending • Descending

- 3 Select the method of arrangement of records within the group.

The value that you can specify for arrangement depends on the type of the group-field:

Value	Char	Num	Date	Explanation
Day			Yes	Day of the month.
Week			Yes	Week number of the month.
Month			Yes	Month number
Quarter			Yes	Quarter number
Year			Yes	Number indicating the year
Numeric range		Yes		A number indicating entries in the range. For example, 10 means, 0-9, 10-19, etc.

- 4 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 134](#).

Select Totals

Figure 5-32 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.



Note

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 132](#)), but you do expect “sorted” results, then specify a Sort Order (instead of grouping).

Figure 5-33 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 5-34 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** check box 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 5-35 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 menu provided.
- 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").

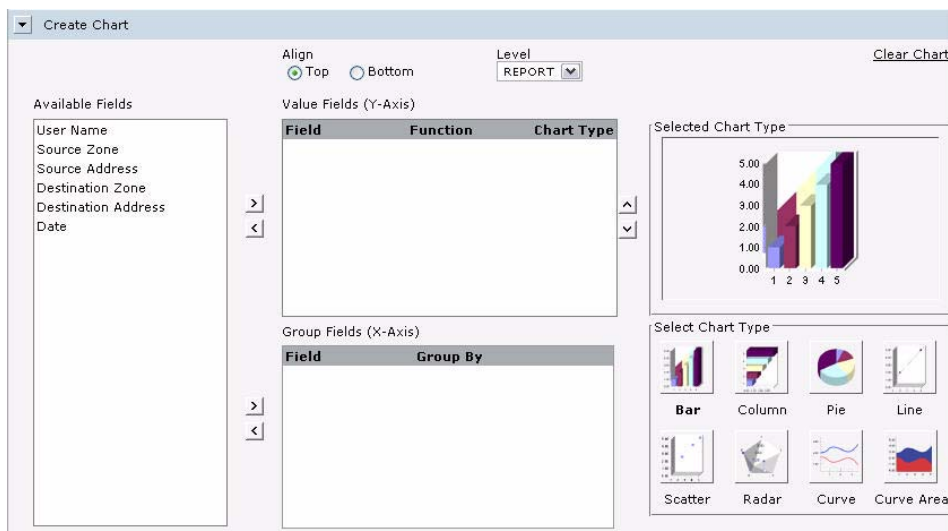


Figure 5-36 Adding a Chart to a Report

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:

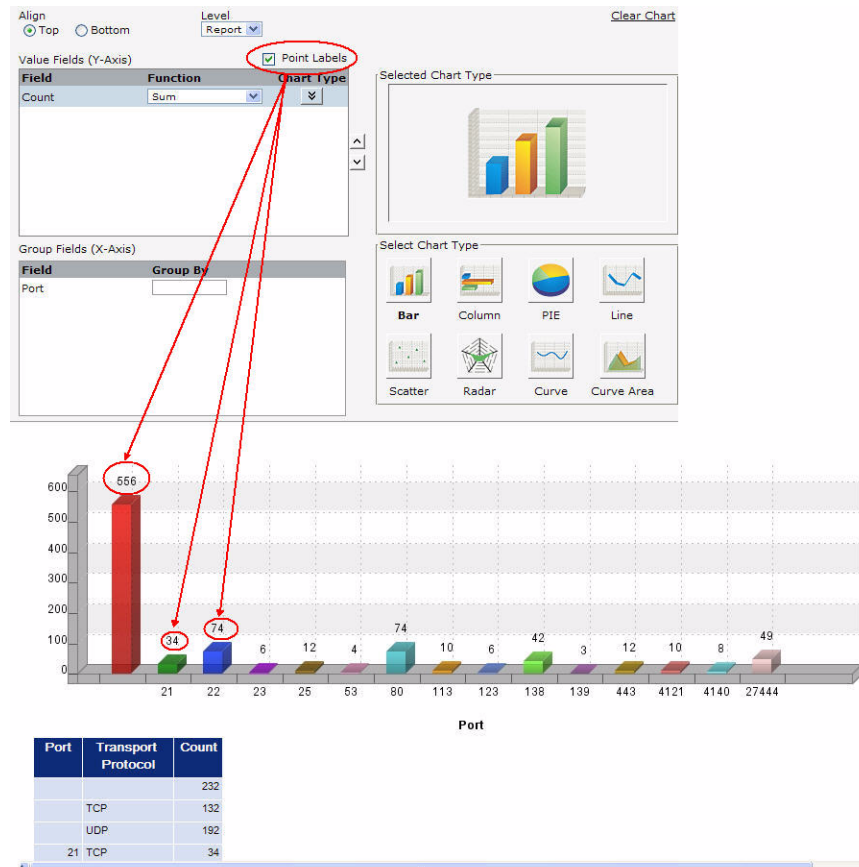
- Select **Top** to place the chart above other components.
- Select **Bottom** to place the chart below other components.
- In **Level**, select PAGE to plot chart having page level data. Select REPORT to plot chart from data that has come from entire report.

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 chart type different than the selected one, click the button on the right to open a box having 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  (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 Reports. (The supplied reports are not editable.)

To edit an existing report

- 1 From any Report list, click the Edit button  (Customize report) 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 125](#)).
- 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 139](#)).
- 4 Click **Save**.



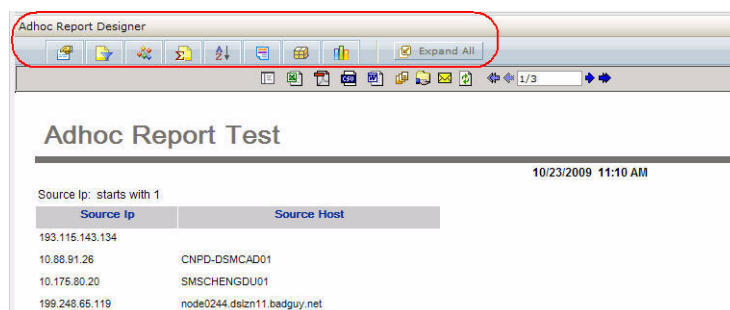
If a user has access rights to “view, run, and schedule all reports”, you can create **private** reports. If you do not have permissions to edit a **public** report that you want to modify but you do have permissions to create private reports, then you can save the public report as a private one and edit the private report.

For more about publishing a report as “public” or “private”, see [Table 5-8 on page 121](#). For more about “access rights” on reports, see [“Setting Access Rights on Reports” on page 140](#).

See also [“Quick Start: Base a New Report on an Existing One” on page 125](#).





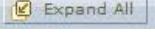
Adhoc Report Designer

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 129 for more information.
	Specify filter criteria. See “Select Filter Criteria” on page 130 for more information.
	Specify grouping. See “Select Grouping” on page 132 for more information.
	Specify the summary (total) fields. See “Select Totals” on page 134 for more information.
	Specify sort order. See “Sort Order” on page 134 for more information.
	Set up highlighting. See “Highlighting” on page 135 for more information.
	Include a matrix. See “Create Matrix” on page 135 for more information.
	Create a chart. See “Create Chart” on page 136 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 114.](#))

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 [“Groups” on page 284](#) in [Chapter 7, System Admin, on page 251.](#)

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 163.](#)

To view and work with Logger Report queries, click Design | **Queries** on the Reports left menu bar. The contents for the selected query is displayed. To view the contents of a different query, select a query name in the **Queries** list on the left. In [Figure 5-37 on page 141](#), the query “SANS Top 5 - Password Changes” is selected.

Query Object List

Save Cancel Import

Queries

(Starts With) >>

SANS

- SANS Top 5 -2- Failed Res Access Events
- SANS Top 5 -2- Failed Resource Access
- SANS Top 5 -3- Password Changes**
- SANS Top 5 -3- User Account Creations
- SANS Top 5 -3- User Account Deletions
- SANS Top 5 -3- User Account Modifications
- SANS Top 5 -4- Vulnerability Scanner Logs
- SANS Top 5 -5- Alerts from IDS
- SANS Top 5 -5- IDS Signature Destinations
- SANS Top 5 -5- IDS Signature Sources
- SANS Top 5 -5- Top 10 Types of Traffic
- SANS Top 5 -5- Top Destination IPs
- SANS Top 5 -5- Top Target IPs

Name SANS Top 5 -3- Password Cha

SQL

Edit Load in New Window

```
SELECT events.arc_destinationUsername "User Name",
events.arc_sourceZoneURI "Source Zone",
events.arc_sourceAddress "Source Address",
events.arc_destinationZoneURI "Destination Zone",
events.arc_destinationAddress "Destination Address",
events.arc_endTime "Date"
FROM events
WHERE events.arc_categoryBehavior = "/Authentication/Modify"
AND events.arc_categoryOutcome = "/Success"
AND events.arc_name like "password%"
```

Mandatory Filtering

On Field (Any)

Fields

User Name

Field User Name

Caption User Name

Data Type CHAR

Format

Width 10

Align Left

Output Format

Input Format

Group Label (Select to add group label)

Lookup Values

Lookup Key Field

SQL Predefined

User Defined SQL Edit

Fetch on Every Use

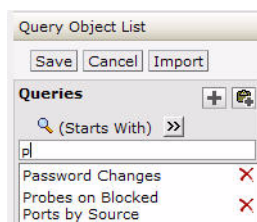
Display Column

Value Column

Figure 5-37 Report Queries Object List

ArcSight Logger Reporting provides a set of pre-built queries, which are used as the basis for the Foundation Reports and Solutions Reports to address common security use cases (as described in [“Report Groups” on page 94](#)).

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.



Caution

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 either using plain English keywords, field names, or regular expressions. See [“Searching for Events on Logger” on page 72](#) 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




Note

You can search for an existing query. 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.

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 **Queries** list, select the name of the query that you want to copy.
- 2 Click  (Add Like), then click **OK** on the resulting message dialog to confirm the copy.


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”. The new query is selected and displayed on the Query Object List page.
- 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).
- 4 Click **Save**.



Caution

You must click **Save** to save the new query object to the Query Object List. Before you save the new query for the first time, it is only a temporary object. If you navigate away from this page before clicking Save, the copied query object will not be retained.

Designing a New SQL Query

- 1 Click  (Add) button.
- 2 In **Name** field, specify a unique name for this query object.
- 3 Under **SQL**, click **Edit** to design SQL.

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.)

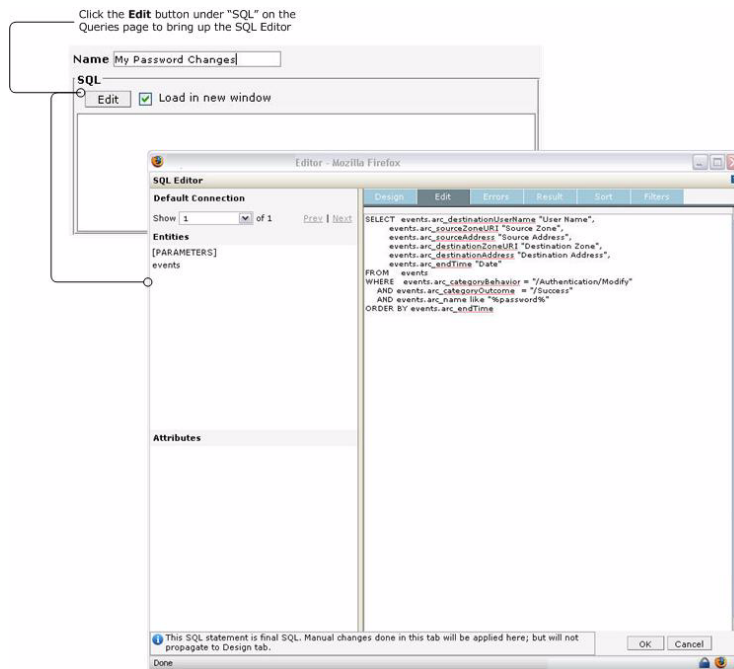


Figure 5-38 Query SQL Editor

- 4 Use the SQL Editor to define the query statement. (See [“Defining SQL in the Editor”](#) on page 156.) Report queries are case insensitive.



Caution

If a report query uses single quotes (' ') in the SELECT clause, the report designer starts refreshing continuously and does not allow you to proceed further. Therefore, if a SELECT clause uses quotes, make sure you alias those fields. For example:

```
Select events.arc_deviceSeverity,
sum(IF (events.arc_name = 'allow', 1, 0)) as Sum1,
sum(IF (events.arc_name = 'object', 1, 0)) as Sum2
From events
group by events.arc_deviceSeverity
```

- 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** button to save your work as part of the query object.



Caution

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 the that field.



Figure 5-39 Query Field Attributes


You can set the following properties on fields in a query.

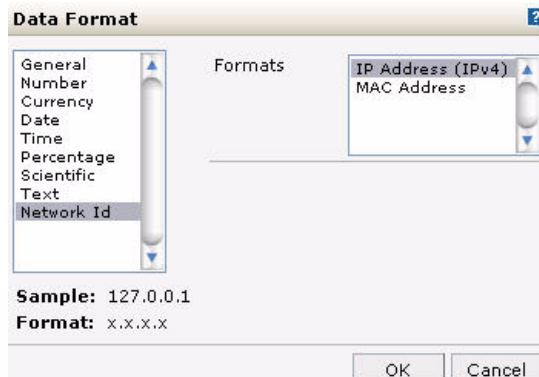
Table 5-13 Query Field Attributes

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 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.

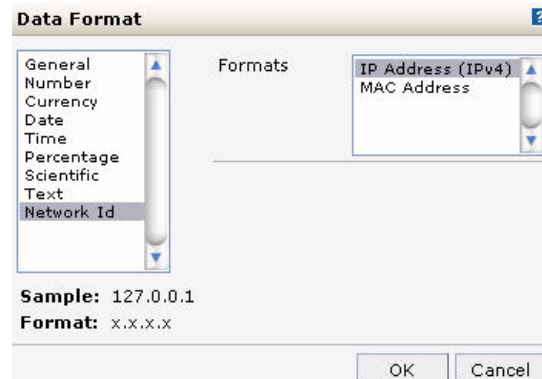
- 4 Click **OK** to accept the changes and close the dialog.

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.

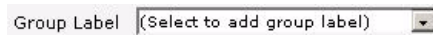


Grouping Fields

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 click (Select to add group label) option.



- 2 Specify 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.

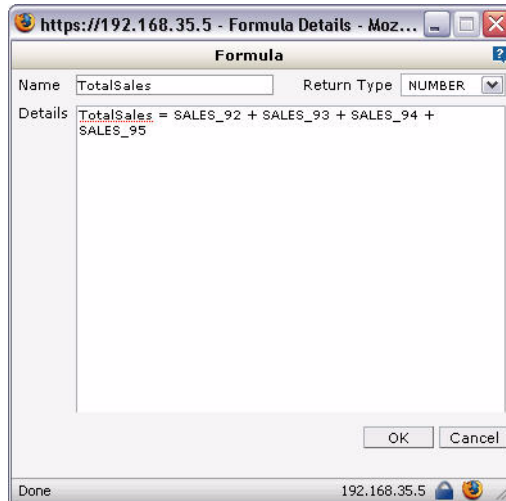
Formula Fields

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 147.](#))





- ◆ 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 147.](#))
- 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  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 serve 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 On the Queries list, 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.

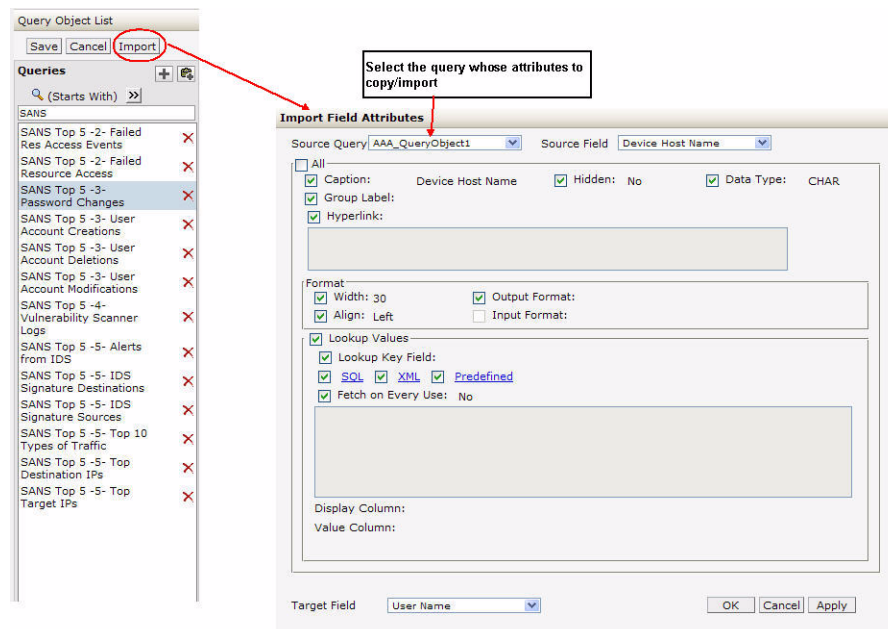


Figure 5-40 Importing Attributes from One Query into Another

- 2 From the **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 **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) 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 130](#) under [“Designing Reports” on page 124](#).)

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 5-41 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 which 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  button next to the **Hyperlink** option to launch the Hyperlink Options dialog.

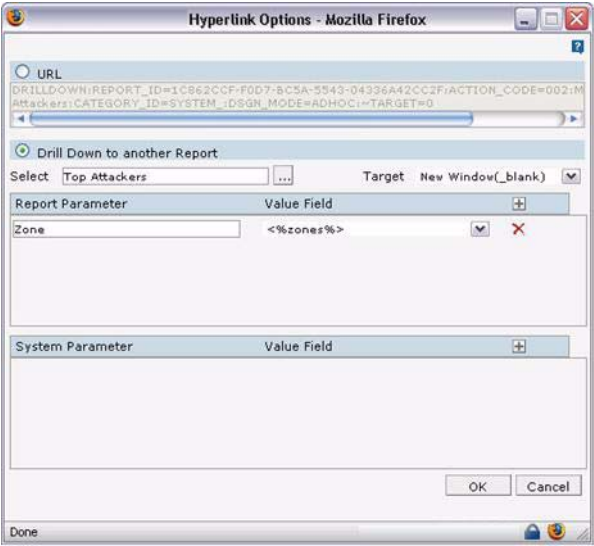




Figure 5-42 Making a Field in a Report a Hyperlink

- 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 152. <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: <code>REQ_SD</code></p> <p>Value Field: <code><%REQ_SD%></code></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.)

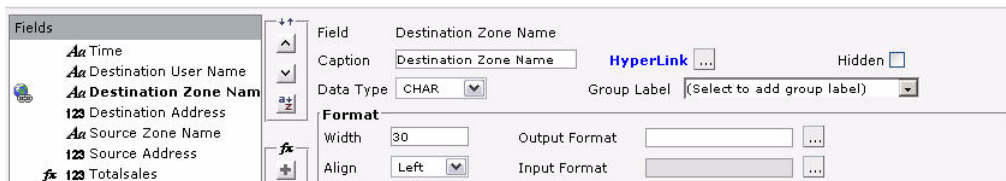


Figure 5-43 Hyperlink Options

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 120)

Parameter	Description and Values
Report Connection Name	Report type and database. We recommend leaving this set to Default .
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.

Specifying 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.

☒ **Lookup Values**


Lookup Key Field:

☐ SQL ☒ **Predefined**

Display: Value:

Below the input fields are two empty list boxes with scrollbars. At the bottom, there is a section labeled 'Values' with a large empty text area.

Figure 5-44 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.

Specifying SQL 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** link.

Figure 5-45 Setting SQL Lookup Values in a Query

- 4 Optionally, Check (select) the **User Defined SQL** 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 is user defined).
- 7 From the Value Column drop-down box, select the column to be used in the filter (only when SQL 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 142](#).


To modify an existing query

- 1 On the **Reports** right panel menu, click **Queries** to bring up the Query Object List.
- 2 In the **Queries** list, select the name of query that you want to modify.
- 3 Edit the query as needed (via the settings described in [“Setting up Queries” on page 140](#)) 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

- 1 On the **Reports** right panel menu, click **Queries** to bring up the Query Object List.
- 2 In the **Queries** list, select the name of query that you want to delete.
- 3 Click  (Delete) next to the query you want to remove.
- 4 Click the **Save** button to make the change permanent. (If you do not click **Save**, the query object will not be deleted.)

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. ArcSight Logger supports SQL queries, and provides an interactive, SQL Editor in which to define SQL statements.

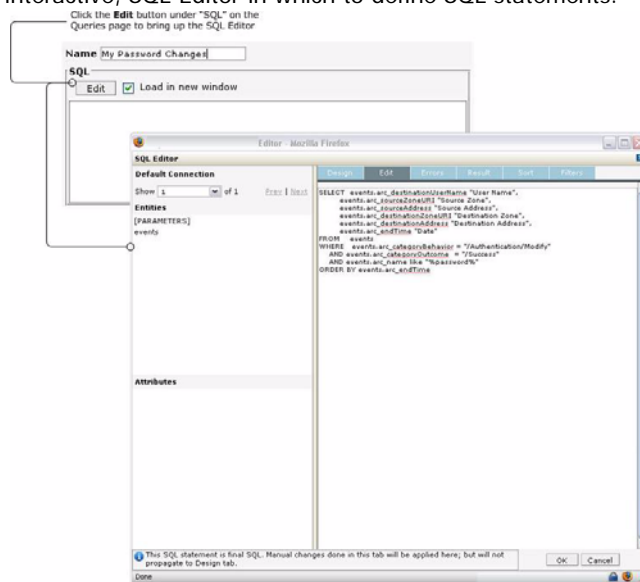


Figure 5-46 Accessing SQL Editor on the Reports | Queries page

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.

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.

Option	Description
Errors	Shows errors, if any, in the SQL statement.
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. ArcSight 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”.

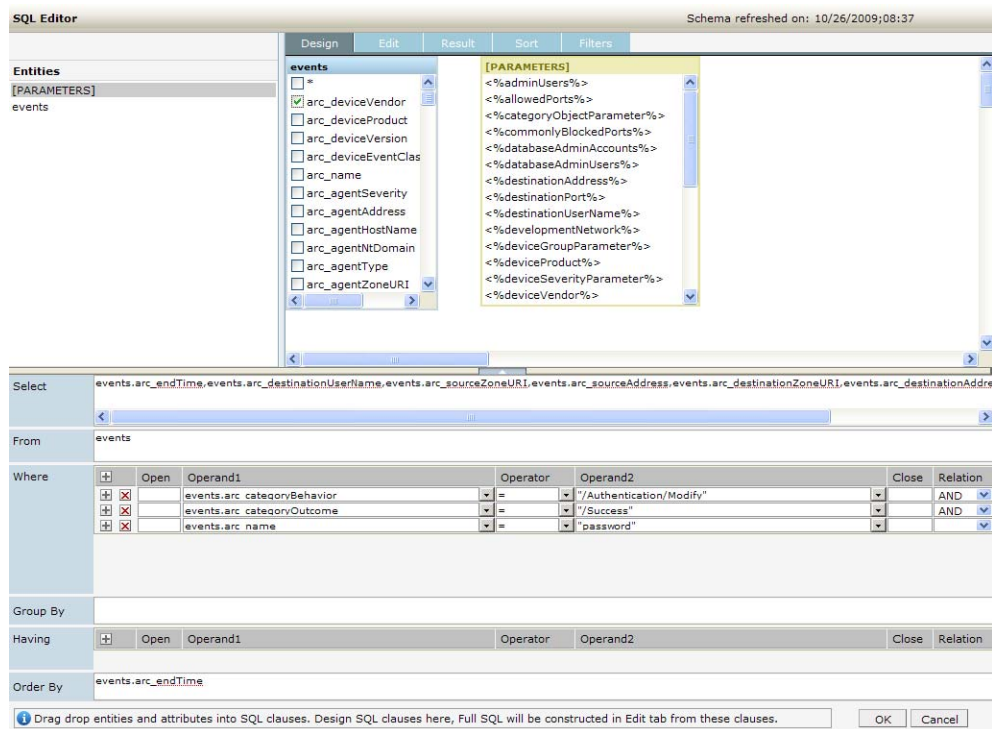


Figure 5-47 SQL Editor: Design Tab

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 are 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.

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 158](#).



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, ArcSight 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.

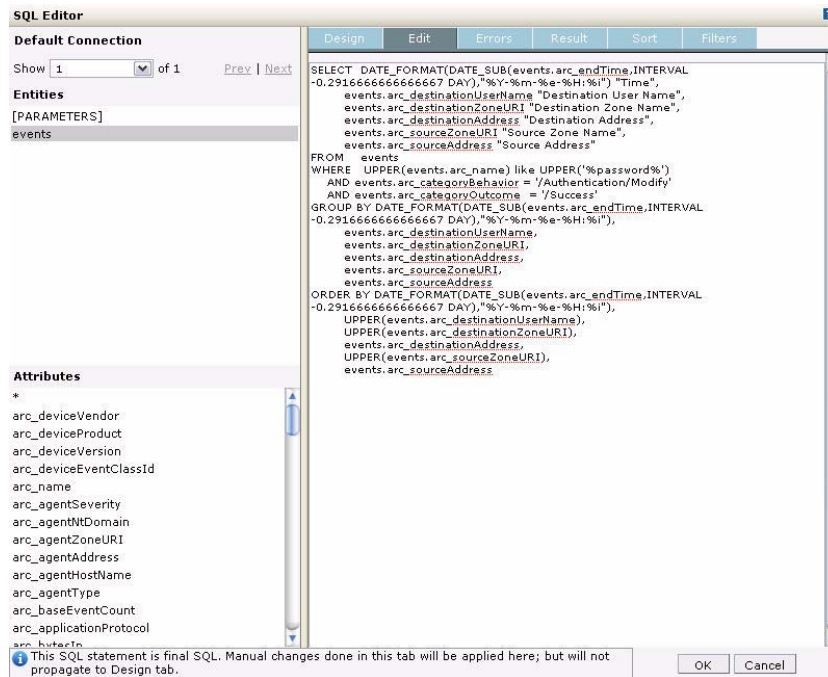


Figure 5-48 SQL Editor: Edit Tab

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.

Errors Tab

The **Errors** tab shows errors compilation errors, if any, in the SQL statement as currently written.

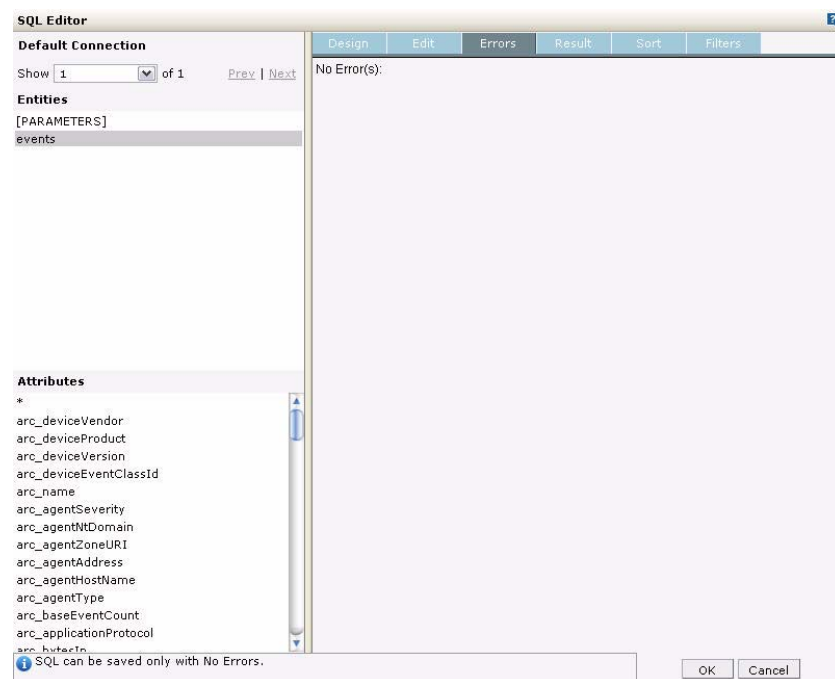


Figure 5-49 SQL Editor: Errors Tab

When you select the Errors tab, the defined SQL statement will be compiled. A message will be displayed on successful compilation, and will also give the details for compilation error(s) if any. This would help you in finding the exact location of error(s) and rectify them before using the SQL results for the report.

If the SQL has used one or more parameters, you will be prompted to provide the values for each of them.

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.

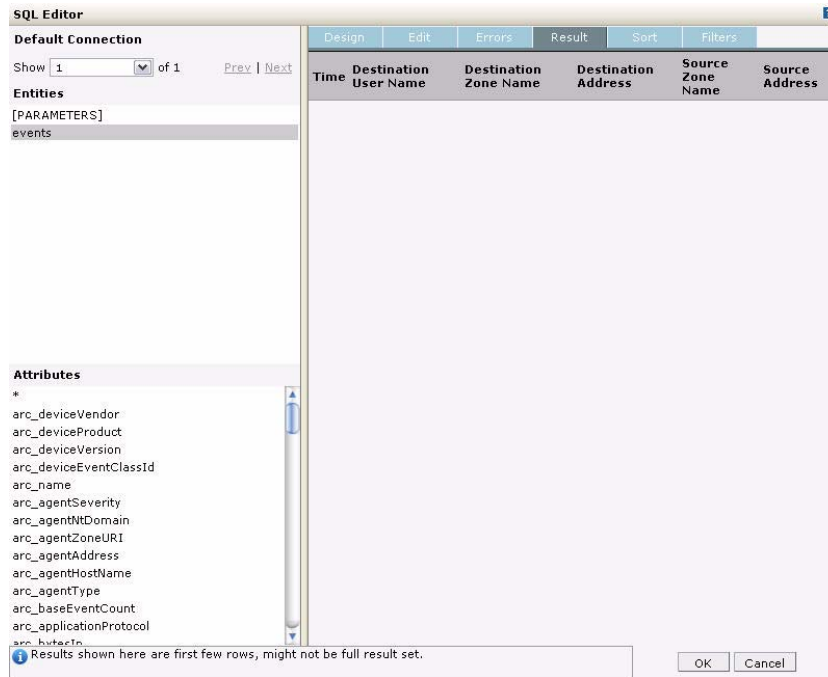


Figure 5-50 SQL Editor: Result Tab

Sort Tab

Click the **Sort** tab to specify levels of sorting at report run time.

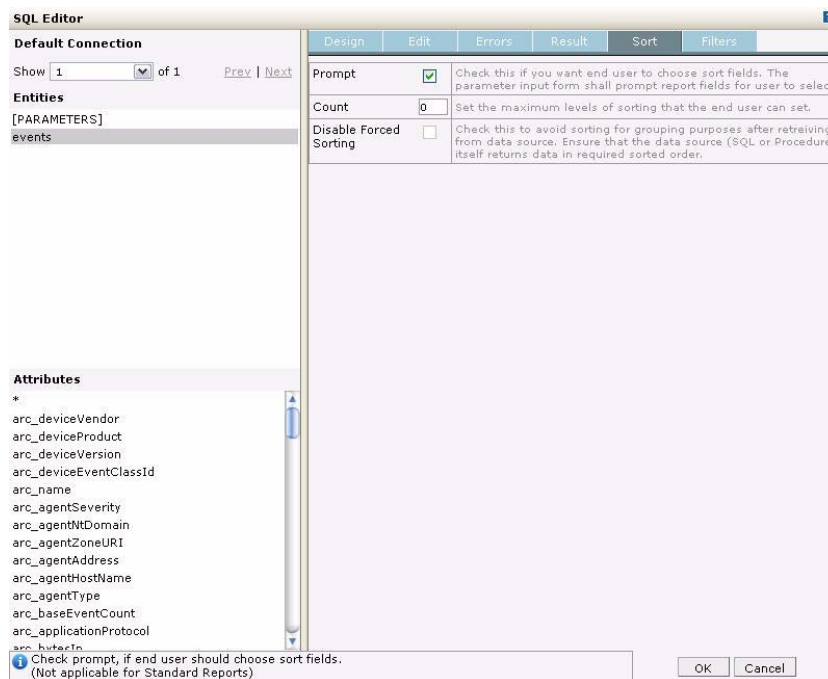


Figure 5-51 SQL Editor: Sort Tab

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.

Run-time Effect: When you specify sorting, the run-time report displays a dialog with one or more selection boxes (the number specified in "Count"). From each selection box, the user can select one field on which the report is to be sorted.

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.

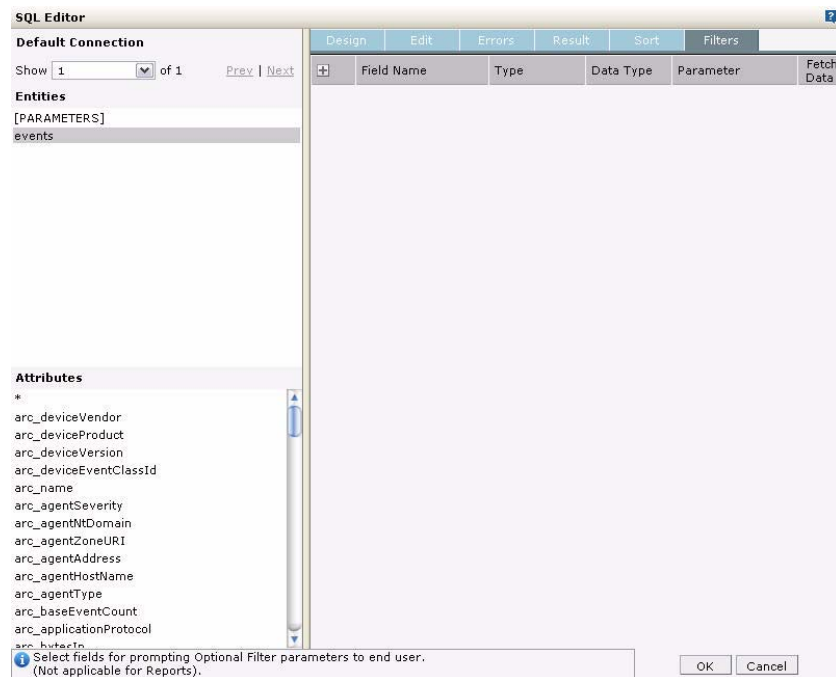




Figure 5-52 SQL Editor: Filters Tab


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.
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.

Run-time Effect: When you add a filter, all the values that the user selects at report run time are added to the SQL query as part of “where” clause. At run time, if the user selects “All” check box, all the optional values are added to the SQL query as IN.

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 140.](#))

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** on the Reports left menu bar.

The screenshot shows the 'Parameters' configuration window. On the left, a list of parameters is shown, with 'categoryObjectParameter' selected. The right pane displays the configuration for this parameter. The 'Name' field is 'categoryObjectParameter', 'Prompt' is 'Resource Type', 'Data Type' is 'CHAR', and 'Size' is '300'. The 'Default Value' is '/Host/Application/D'. The 'Pre Defined List' is populated with various host and application paths. The 'Input Type' is set to 'Combo'. The 'Multi Select' checkbox is checked, and the 'Select list' is set to 'Maximum Selectable Values 0'. The 'Select Default Values' are set to 'Selected'.

Figure 5-53 Report Parameters Object List

Creating New Parameters

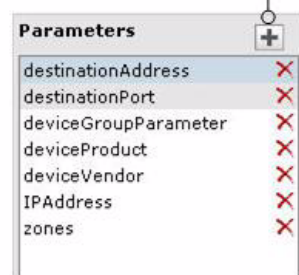


You can search for an existing parameter. To do so, either

- Enter the first few letters with which the parameter name begins (if the "Starts With" search criteria is selected) in the text box above the list of existing parameters, OR
- Enter a word or part of a word that the parameter name contains (if the "Contains" search criteria is selected) in the text box above the list of existing parameters.

- 1 On the Parameter Object List, click at the top right of the **Parameters** list box.

Click to add a parameter



- 2 Specify values for the new parameter. (Details are given in the topics below.)
- 3 After providing all required values, click **Save**.

The parameter is added to the Parameters list.


Setting Parameter Name, Data Type, and Default Values

Name	<input type="text" value="deviceProduct"/>
Prompt	<input type="text" value="Device Product"/>
Data Type	<input type="text" value="CHAR"/> ▼
Size	<input type="text" value="30"/>
Format	<input type="text" value=""/> ...
Default Value	<input type="text" value="'FoundScan'"/>


Figure 5-54 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.
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 167.)
Size	Specify number of characters or digits this parameter should accept. <p>Note: This is only applicable to CHAR and NUMBER data types, not for Boolean or Date type parameters.</p>
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. <p>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.</p>

Default Value for Date Type Parameter

For a date type parameter, the **Default Value** field provides an 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.

The screenshot shows a configuration window for a parameter. It has three main sections: 'Data Type' with a dropdown set to 'DATE' and a 'Size' field set to '12'; 'Format' with a text field containing 'MM/dd/yyyy' and a small menu icon; and 'Default Value' with a dropdown set to 'CURRENT_DATE' and a '±Days' button with a calendar icon.

At report run time, a parameter with a "Date" format will display with the default date set here.

Defining Input Type

The screenshot shows a horizontal control bar labeled 'Input Type'. It contains three radio buttons: 'Text Box' (unselected), 'Combo' (selected), and 'Option' (unselected).

Figure 5-55 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.

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 Multiselect checkbox so that user can select multiple values from the box.

The screenshot shows a configuration window for the 'Multiselect' option. It has a tabbed interface with 'Multi Select' selected. Below the tabs, there are two main sections: 'Select list' with a 'Maximum Selectable Values' field set to '0' and 'Enclosed By' and 'Separator' fields; and 'Select Default Values' with radio buttons for 'Selected' (selected), 'All', and 'None', and a list box containing several host/resource paths with 'Host/Resource/File' selected.

See ["Setting Multiple Default Values" on page 169](#) to configure other settings for this option.

Option	Description
Option	<p>Select “Option” if you want the user to select values represented as options.</p> <p>Select the Multiselect checkbox to have value options in the form of checkboxes.</p> <p>Keep Multiselect checkbox clear to have options in the form of radio buttons.</p>

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).

The screenshot shows a configuration window for a parameter. The 'Data Type' dropdown is set to 'BOOLEAN'. Below it, 'Size' is 30, 'Format' is empty, and 'Default Value' is 1. At the bottom, the 'Values' section has two input fields: 'Checked' with the value 1 and 'Unchecked' with the value 0.

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.

The screenshot shows a set of checkboxes for parameter options. 'Mandatory' and 'Restrict to List' are checked, while 'Visible', 'Pass Values Using Tables', and 'Forced' are unchecked.

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	<p>Select this checkbox if you want the parameter to be visible (displayed) on the input form at report run time.</p> <p>Keep this unchecked (clear) if the value for this parameter be populated from another report or if you want the parameter to use the default value in all cases.</p>

Option	Description
Restrict to List	<p>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.</p> <p>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.</p>
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.
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.

To Set Predefined Values

Pre Defined List

Display Name	Value
Host	/Host
Host/Operating System	/Host/Operating System
Host/Application	/Host/Application
Host/Application/Database	/Host/Application/Database
Host/Application/Database/Data	/Host/Application/Database/Data
Host/Application/Service	/Host/Application/Service
Host/Application/Service/Email	/Host/Application/Service/Email
Host/Application/Service/Instant	/Host/Application/Service/Instant
Host/Application/Service/MMS	/Host/Application/Service/MMS
Host/Application/Service/Peer to Peer	/Host/Application/Service/Peer to Peer
Host/Application/Service/Phone	/Host/Application/Service/Phone

☐ Display Parameter Name

Figure 5-56 Setting Predefined Values for a Combo Input Parameter

- 1 In the **Display Name** field, specify the value to be displayed at run time. (The value the user will see.)
- 2 In the **Value** field, specify the value to pass (as a filter).
- 3 Click (Add) to add the display name to the list.

(To delete an option from the list, select the value and click .)
- 4 Repeat these steps for each option.

Select the check box **Display Parameter** if you want to provide the user with the option of adding the parameter as a control on a report. In **Name**, specify a name for the parameter.



The **Display Parameter** and **Name** settings have no effect when the Parameter Object is used in an ad hoc report.

Setting Multiple Default Values

If you selected Combo Input Type (as described in [“Defining Input Type” on page 166](#)), you need to define the following settings in the Parameter editor:

- *Maximum Selectable Values*—Specify the maximum number of values that can be selected or provided for a parameter.
- *Enclosed By*—Specify the character to use to enclose the set of values. This will depend on the database.
- *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

- 1 On the **Reports** right panel menu, click **Parameters** to bring up the Parameter Object List.
- 2 In the **Parameters** list, select the name of parameter that you want to modify.
- 3 Edit the parameter as needed (via the settings described in [“Creating New Parameters” on page 164](#)) and click **Save**.

Only custom parameters can be modified, not supplied parameters, since supplied parameters are required for use in Foundation Reports and Solution pack add-ons.

Deleting a Parameter

- 1 On the **Reports** right panel menu, click **Parameters** 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) next to the parameter you want to remove.
- 4 Click the **Save** button to make the change permanent. (If you do not click **Save**, the query object will not be deleted.)

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 menu bar.

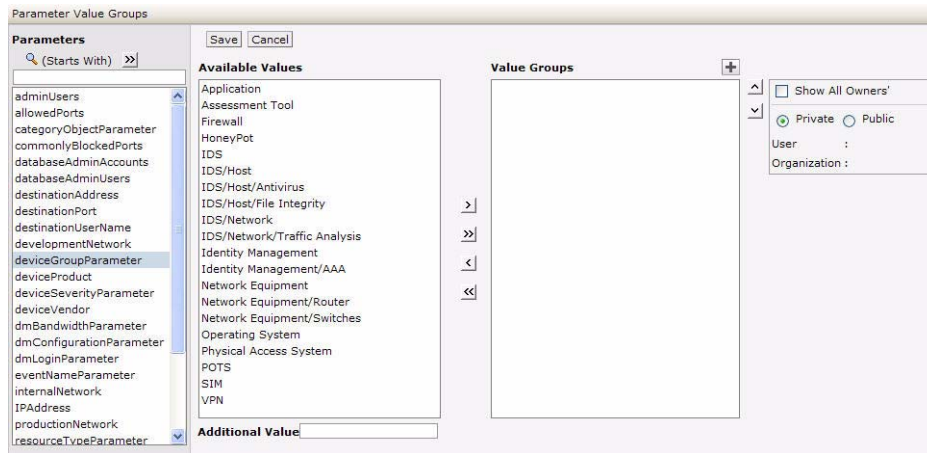



Figure 5-57 Parameter Value Groups



The following table describes the options on the Parameter Value Groups page. In addition, you can search for an existing parameter value group. To do so, either

- Enter the first few letters with which the parameter value group name begins (if the "Starts With" search criteria is selected) in the text box above the list of existing group names, OR
- Enter a word or part of a word that the parameter value group name contains (if the "Contains" search criteria is selected) in the text box above the list of existing group names.



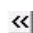
Table 5-19 Parameter Value Groups





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. Such groups will have  icon in the group title.
Option buttons Private and 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).

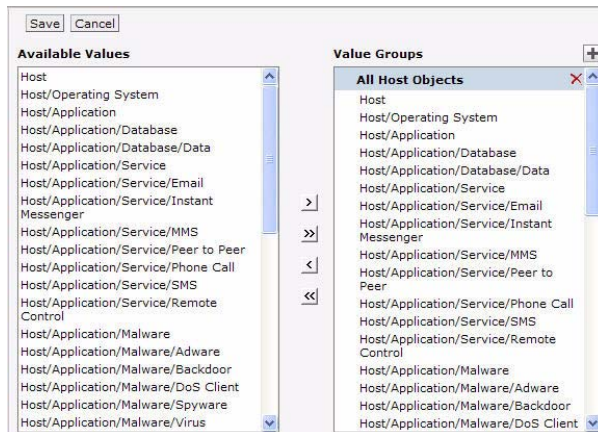
Click **Save** button to save the work.




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.

Click **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. Starting with Logger v4.0, you can define your own layout files. See ["Defining a New Template" on page 173](#) for more information.

To view and work with Logger Report template styles, click Design | **Template Styles** on the Reports left menu bar.

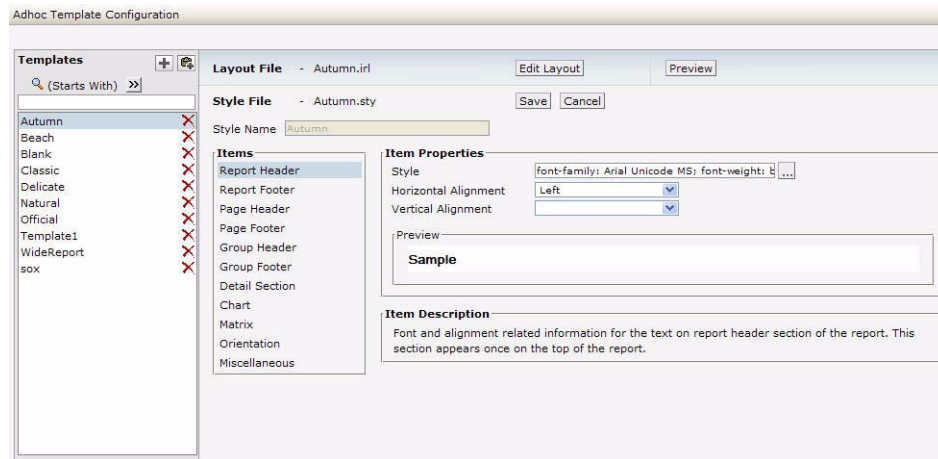


Figure 5-58 Report Template Styles Configuration


Defining a New Template



You can search for an existing template. To do so, either

- 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, OR
- 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.

ArcSight 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 114](#).

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 258](#) 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 menu to view a list of currently scheduled jobs.

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 5-59 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 174](#).



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 174](#).

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 258](#) 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 5-59](#).)

- Click **Add** to bring up the Add Report Job page.

Add Report Job

Name:

Schedule: 23 Hours

Report Name * 

Delivery Operations

Select Delivery Options

☐ Email ☒ Publish


File Name: ☒ Suffix Timestamp Format:

☒ Public ☐ Private

Valid Upto

☒ 1 Months after generation

☐ End of this Month

☐ Date: 

Report Parameters

No Parameters

Start: ☒ Dynamic

End: ☒ Dynamic

Device Groups

No Device Groups

Storage Groups

Default Storage Group
Internal Event Storage Group

Devices

kvuont43-wifi.sv.arcsight.com [LoggerReplay-psh]
NOT kvuont43-wifi.sv.arcsight.com [LoggerReplay-psh]

- 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.



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.

- Choose a Delivery Option (**Email** or **Publish**).

Depending on which delivery option you choose, the associated parameters are displayed. Click to enable (check) or disable (uncheck) these options.

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 122.](#)

Delivery Operations

Select Delivery Options

☒ Email ☐ Publish

Report Format ACROBAT PDF

Send Report As ☒ Link ☐ Attachment

To

Cc Bcc

Subject

Message

Report Untitled has been generated.
Please click the following link to view the report.
<%LINK%>
- System Administrator

- ◆ **Publish** - For details on setting publishing options, see [“Publishing Reports” on page 121.](#)

Delivery Operations

Select Delivery Options

☐ Email ☒ Publish

File Name SANS_Top5_Logs_by_Host ☒ Suffix Timestamp Format MM-dd-yyyy

☒ Public ☐ Private

Valid Up to

☒ 1 Months after generation

☐ End of this Month

☐ Date 11/12/2007

- 5 Fill in the rest of the fields based on the report you chose, as described in [“Add Report Job Settings” on page 176.](#)
- 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.

Option	Description
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 122. • Publish - For details on setting publishing options, see "Publishing Reports" on page 121.
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 120.</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 / Run In Background Report Parameters" on page 117.</p>

Deploying a Report Package

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 Package page to load and deploy packages of new reports onto your Logger system.

On the Reports page left panel menu, click **Deploy Report Package** to get started.

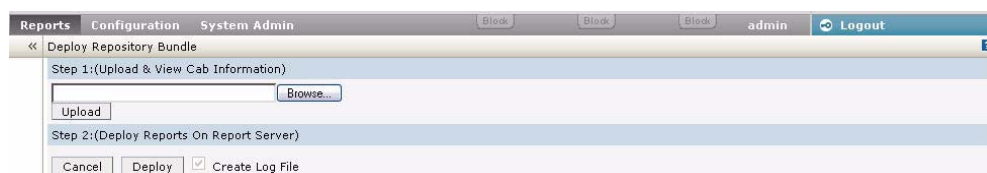


Figure 5-60 Deploy Report Package

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 contents 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 97](#).

Report Server Administration

ArcSight 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.

To view or modify the report server and client configuration, click **Report Administration** on the Reports page left panel menu.

Report Configuration

Save Cancel

DATABASE CONNECTION TIMEOUT	14400
LOG LEVEL	ERROR
DATA SOURCE FETCH SIZE	50
EMAIL FROM ADDRESS	
HOST URL	https://<logger_hostname>
SMTP SERVER	127.0.0.1

Save Cancel

Figure 5-61 Reports Server Configuration

The following table describes the report server configuration settings.

Table 5-21 Reports Server 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 is 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>
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>

Option	Description
SMTP Server	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. Example: SMTP_SERVER=127.0.0.1

Using Report Category Filters

A Search Group Filter can be optionally assigned to each report category, for example:

- Foundation Report categories:
 - ◆ Configuration Monitoring
 - ◆ Intrusion Monitoring
 - ◆ SANS Top 5
- User Report category:
 - ◆ Default Reports

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” on page 222](#) 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 pulldown 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 pulldown 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

Starting with Logger v3.0, you can backup and restore report content. For more information about this feature, see [“Configuration Backup and Restore” on page 235](#).

Chapter 6

Configuration

This chapter describes the Configuration tab, in which you create and manage Receivers, Forwarders, Devices, Device Groups, and Filters.

In this chapter:

- ["Devices" on page 181](#)
- ["Event Archives" on page 184](#)
- ["Storage" on page 187](#)
- ["Event Input/Output" on page 193](#)
- ["Alerts" on page 207](#)
- ["Scheduled Tasks" on page 221](#)
- ["Filters" on page 222](#)
- ["Saved Searches" on page 225](#)
- ["Search Optimization" on page 229](#)
- ["Peer Loggers" on page 231](#)
- ["Configuration Backup and Restore" on page 235](#)
- ["System Maintenance" on page 238](#)
- ["License Information" on page 245](#)
- ["Retrieve Logs" on page 246](#)
- ["Exporting and Importing Content" on page 247](#)

Configuration

The Configuration tab provides access to basic Logger functions, such as creating a Receiver or disabling an existing Forwarder.



Receivers, Devices, and other resources created by one user are visible to all other users, subject to user group privileges. Resources are shared by all sessions.

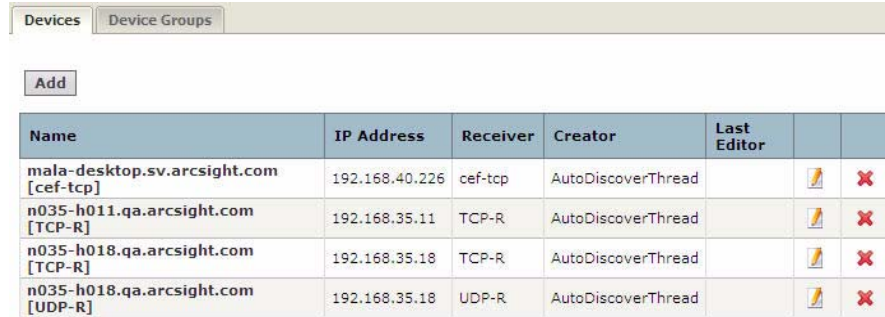
Devices

The Devices section manages both Devices and named collections of Devices called Device Groups.

Devices

A Device is a named event source, comprising an IP address (or hostname) and Receiver name. Two Receivers can receive events from the same IP address, so IP address alone is insufficient to identify a Device. 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.



Name	IP Address	Receiver	Creator	Last Editor		
mala-desktop.sv.arcsight.com [cef-tcp]	192.168.40.226	cef-tcp	AutoDiscoverThread			
n035-h011.qa.arcsight.com [TCP-R]	192.168.35.11	TCP-R	AutoDiscoverThread			
n035-h018.qa.arcsight.com [TCP-R]	192.168.35.18	TCP-R	AutoDiscoverThread			
n035-h018.qa.arcsight.com [UDP-R]	192.168.35.18	UDP-R	AutoDiscoverThread			

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 the **Configuration > Settings** tab, click **Devices** in the sub-menu, then click the **Devices** tab. A display similar to that shown in Figure 6-1 appears.
- 2 Click **Add**.
- 3 Enter a name, an IP address, and select a Receiver for the new Device.
- 4 Click **Save** to add the new Device, or **Cancel** to abandon it.

To edit a Device

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.

- 1 Click the **Configuration > Settings** tab, click **Devices** in the sub-menu, then click the **Devices** tab. A display similar to that shown in Figure 6-1 appears.
- 2 Locate the Device to be edited and click the edit icon () on that row.
- 3 Change the Name or IP address for the Device.
- 4 Click **Save** to update the Device Group, or **Cancel** to abandon your changes.

To delete a Device

- 1 Click the **Configuration > Settings** tab, click **Devices** in the sub-menu, then click the **Devices** tab. A display similar to that shown in Figure 6-1 appears.

- 2 Locate the Device to be deleted 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.
- 3 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, shown in [Figure 6-2](#), 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 device group are stored. Doing so enables you to retain event data from different sources for a 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 189](#).

Maximum number of device groups that can be created on Logger: No limit.

To create a Device Group

- 1 Click the **Configuration > Settings** tab, click **Devices** in the sub-menu, then click the **Device Groups** tab. A display similar to that shown in [Figure 6-2](#) appears.
- 2 Click **Add**.
- 3 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.
- 4 Click **Save** to create the new Device Group, or **Cancel** to abandon it.

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 unselect devices, ctrl-click each device name.


Name:

Devices:


- mala-desktop.sv.arcsight.com [cef-tcp]
- n035-h011.qa.arcsight.com [TCP-R]
- n035-h018.qa.arcsight.com [TCP-R]
- n035-h018.qa.arcsight.com [UDP-R]

Figure 6-2 Device Groups page

To edit a Device Group

- 1 Click the **Configuration > Settings** tab, click **Devices** in the sub-menu, then click the **Device Groups** tab. A display similar to that shown in [Figure 6-2](#) appears.
- 2 Locate the Device Group to be edited and click the edit icon () on that row.
- 3 Change the Name or 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 the **Configuration > Settings** tab, click **Devices** in the sub-menu, then click the **Device Groups** tab. A display similar to that shown in [Figure 6-2](#) appears.
- 2 Locate the Device Group to be deleted and click the delete icon () on that row. Deleting a Device Group does not affect the set of Devices.
- 3 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Device Group.

Event Archives

Event Archives let you 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.

For **Logger appliances**, the location needs to be an NFS mount, CIFS mount, or SAN, which is configured using the Logger user interface. For **the software version Loggers**, the location can be a local directory or a mount point for that you have already established on the machine on which the Logger software is installed.

You can also schedule a daily archive of the events. (See [“Scheduled Event Archive” on page 186](#) for information.)

Once events are archived, they are no longer included in search operations. To include those events in search operation, you must load the archive 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 a query that includes indexed fields is run on archived events, it runs slower than when the data was not archived because the index data on Logger is not archived with events. Therefore, when event archives are loaded, indices are not available.

When an Event Archive is unloaded, it is available for loading, but its events are not included in searches.

If you need to archive a large number of events (in the order of tens of GB), ArcSight recommends that you archive during the off-peak hours to prevent impacting the performance of your Logger.

Event Archives

To save events for a particular day, add an Event Archive. The table in the Event Archives tab shows the current archives and their status.

Name	Day	Status	Creator
test_alias	8/4/09	Archived	admin
archive [2009-08-03]	8/3/09	Archived	scheduledArchivor
archive [2009-08-02]	8/2/09	Archived	scheduledArchivor
archive [2009-08-01]	8/1/09	Archived	scheduledArchivor
archive [2009-07-31]	7/31/09	Archived	scheduledArchivor
archive [2009-07-30]	7/30/09	Archived	scheduledArchivor
archive [2009-07-29]	7/29/09	Archived	scheduledArchivor

To add an Event Archive



An archive storage location must be established on the Logger before you can archive its events. See [“Archive Storage Settings” on page 186](#) for more information.

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Archives** in the left panel.
- 3 Click **Add** in the Event Archives tab, in the right panel.
- 4 Enter a meaningful name in the Name field for the new Event Archive and specify the day 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.

The Event Archives table (under the Event Archives tab) lists the archives by Alias.

- 5 Click **Save** to start archiving events, or **Cancel** to quit.

To load or unload an Event Archive

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Archives** in the left panel.
- 3 Locate the Event Archive to be loaded or unloaded and click the edit icon (🔧) on that row. 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.



Loading an archive with events that are still current will result in duplicates.

To delete an Event Archive

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Archives** in the left panel.

- 3 Locate the Event Archive to be deleted and click the delete icon (✖) on that row.
- 4 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 appear on the archive list on the Event Archives tab.

Make sure you are familiar with the information in ["Impact of Daylight Savings Time Change on Logger Operations" on page 258](#) before you schedule an event archive.

To schedule a daily event archive

- 1 Click **Configuration** 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 For Daily Archive to Start" from the pulldown list. Scheduled archives must start on the hour, and midnight and 1:00 AM are not on the list.
- 5 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 version of 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.



Once the output location for Event Archives has been set up, it cannot be changed.

To perform one-time Archive Storage Setting setup

- 1 If you are using the Logger appliance, create the NFS or CIFS mount point. (See ["Storage" on page 263](#) and ["CIFS Settings" on page 263](#).)

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** 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 For the Logger appliance, select the NFS or CIFS mount point, or SAN for Event Archive output and enter a file path. Click **Save**.

For the software version of Logger, enter the directory path in the Archive Directory field. Click **Save**.



Deleting the NFS or CIFS Mount or detaching the SAN associated with Event Archive is not recoverable. If the mount point is deleted, the Event Archive command will no longer function. Scheduled Event Archive jobs will create daily errors.

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 [“Devices participate in Retention Policy” on page 23](#). The Storage section has three tabs: Storage Groups, Storage Rules, and Storage Volume.



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. As a result, now you have five storage groups available for event storage and one for Logger’s internal events.

ArcSight recommends that you create the maximum allowed four additional Storage Groups (in addition to the two that preexist—Default Storage Group and the Internal Storage Group) during Logger Initialization (as discussed in [“3 Storage Groups” on page 26](#)) even if you do not need all of them because **you cannot add storage groups after the Logger has been initialized**, although you can decrease or increase the size of a Storage Group later.

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.

A Default Storage Group and an Internal Storage Group are created automatically during the Logger initialization phase.



Once a Storage Group is created, it cannot be deleted however its size can be increased or decreased any time. **Storage Groups can only be created during the Logger initialization phase**, described in [“Initialization Sequence \(for all Loggers\)” on page 24](#). (See [“To edit \(including resizing\) a Storage Group” on page 188](#) to change the size of a Storage group.)


Storage Groups					
Name	Maximum Age (Days)	Maximum Size (GB)	Creator	Last Editor	
Default Storage Group	60	645	admin	admin	

Figure 6-3 Storage Groups page

To add a Storage Group

The Add button is not visible after Logger has been initialized.


- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Storage** in the left panel.
- 3 Click **Add** in the Storage Groups tab in the right panel.
- 4 Enter the following values:

Parameter	Description
Name	Choose a name for the Storage Group
Maximum Age	Specify the number of days to retain events. Events older than this number of days will be deleted.
Maximum Size	Enter a maximum event data size, in GB.

- 5 Click **Save** to store the changes, or **Cancel** to quit.

Once the Logger has been set up, the Storage Groups page, as shown in [Figure 6-3](#), does not allow adding or deleting Storage Groups.

To edit (including resizing) a Storage Group

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Storage** in the left panel.
- 3 Identify the Storage Group you want to modify and click the edit icon () for it.
- 4 Change the desired parameters such as the name of the storage group, or increase or decrease Maximum Age or Maximum size.

Note: The name of the 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 the following message, indicating that reducing storage group size in this situation will require you to delete existing data.

ArcSight Logger

Monitor Analyze Reports Configuration System Admin

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.

Setting Maximum Size to [15] GB will delete at least [9] GB of data. If you want to proceed, set the Maximum Age value to [2]. Doing so will automatically trigger data deletion. After the data has been deleted, return to this page and set the Maximum Size to [15] GB.

Note: Data deletion takes time. If new events are added to the storage group while data deletion is in progress, the final data size might be larger than [15] GB. If that occurs, repeat this operation.

Name: SGC

Maximum Age (Days): 30

Maximum Size (GB): 50

Current Size (GB): 24

Save Cancel

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.
- Note:** The “Current Size” value changes as data is 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.

To add a Storage Rule

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. You cannot create additional storage groups after a Logger has been initialized. However, you can create additional device groups, as described in [“Device Groups” on page 183](#).


- 1 Click the **Configuration > Settings** tab, click **Storage** in the sub-menu, and click the **Storage Rules** tab.
- 2 Click **Add** and enter the following parameters: The page shown in [Figure 6-4](#) is displayed.

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.


- 3 Click **Save** to add the new Storage Rule, or **Cancel** to quit.

Figure 6-4 Storage Rules page

To edit or reorder a Storage Rule

- 1 Click the **Configuration > Settings** tab, click **Storage** in the sub-menu, and click the **Storage Rules** tab.
- 2 Find the Storage Rule to be edited in the table.
- 3 Click the Edit icon (). 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 the **Configuration > Settings** tab, click **Storage** in the sub-menu, and click the **Storage Rules** tab.
- 2 Find the Storage Rule to be deleted in the table.
- 3 Click the Delete icon (). Confirm the delete.

Storage Volume

Storage Volume settings allows you to specify where Logger will store events. Logger can store events locally (on the storage provided with Logger), on a Network File System, or a Storage Area Network (SAN). This decision must be made when the Logger is first initialized. Although a Network File System (NFS) can be used as primary storage for events on a Logger, this configuration is not recommended because the performance is sub-optimal.

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.

For more information, see [“Initialization Sequence \(for all Loggers\)” on page 24](#).

**Note**

- A storage volume is automatically established with default values for a “Typical” installation of the software version of Logger. For a “Custom” installation, you need to establish the storage volume using the procedure described in this section.
- Storage volume can be extended after initialization, however its size cannot be reduced. For more information, see [“Storage Volume Size Increase” on page 244](#).

To specify storage volume settings

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Storage** in the left panel.
- 3 Click **Storage Volume** in the right panel. If a storage volume is not established already on the Logger, the following page is displayed.

**Note**

The options displayed on the following page depend on the type of Logger—appliance or software. Additionally, the following page is only displayed on the software version of Logger if you chose to Custom install it.

The following screen shot is from a Logger appliance. The table that describes the options on this page categorizes the options displayed for Logger appliance and the software version of Logger.

Storage Groups
Storage Rules
Storage Volume

Important

Before you may setup any receivers to receive events or storage groups to store events, you must perform a one-time setup of the single storage volume where all event data will be stored.

If you want the events to be stored on a remote file system you must first add the [remote file system mount](#).

These settings cannot be changed once saved, so be certain they are correct before you click **Save**.

It is highly recommended that you preallocate some or all of the space on the storage volume to ensure maximum performance.

Be aware that on a large remote volume, preallocation may take a very long time to complete. During this process no storage groups may be created or events received or stored. This process may not be stopped until it has fully completed.

Mount Location

Local

Path

Maximum Size (GB)

1291

Preallocation Amount (%)

0

Save

4 Enter the following values:

Parameter	Description
Only for the Logger appliance:	
Mount Location	Choose Local if you want to store events on Logger or the mount name of a remote file system. (To set up a remote file system mount, see "Storage" on page 263 .) Note: Use of a Network File System (NFS) as primary storage for Logger events is not recommended. However, an NFS system can be used for archiving Logger data.
Path	For Logger appliance , if mount Location is not Local, specify the root folder on the remote file system in which to store event data. If mount location is Local, event data is stored to the <code>/opt/data/logger</code> directory on the appliance.
Maximum Size	Enter the storage volume size, in GB.
Pre-allocation Amount	The percentage of the volume to pre-allocate (0-100). ArcSight recommends 100% for both local and remote volumes. Note: Even though 100% pre-allocation can take a long time on remote volumes, doing so improves performance on your Logger and protects it from running out of disk space. Allocating lesser than 100% space may result in sub-optimal performance on the Logger.
Only for the software version of Logger:	
Path	For the software version of Logger , the path is pre-configured to the <code>/opt/data/logger</code> directory and cannot be changed.

Parameter	Description
Maximum Size	<p>Enter the storage volume size, in GB.</p> <p>For the software version <i>Logger</i>, the maximum size is determined by the lower value of the following:</p> <ul style="list-style-type: none"> Limit specified in the <i>Logger</i> license. (See “License Information” on page 245 for this information.) The storage volume partition size on the system on which <i>Logger</i> software is installed.

5 Click Save.

To increase the size of a storage volume:

See [“Storage Volume Size Increase” on page 244](#).

Event Input/Output

Use the Event Input/Output section to manage the Receivers and Forwarders that listen for and capture events and send them to other destinations, including ArcSight ESM.

Receivers

Receivers are created to receive events from files and on the network. Receiver types include UDP, TCP, SmartMessage, and two types of file follower, File Transfer and File Receiver:

- **UDP.** UDP receivers listen for User Datagram Protocol messages, such as Syslog format datagrams.
- **TCP.** TCP receivers listen for Transmission Control Protocol messages. Syslog messages can also be sent using TCP.
- **CEF UDP.** UDP receiver that receives events in Common Event Format.
- **CEF TCP.** TCP receiver that receives events in Common Event Format.
- **File Transfer.** File Transfer receivers read remote log files using scp, sftp or ftp protocol.
- **File Receiver.** File Receiver-type receivers read log files from a network file system (NFS), CIFS, or Storage Area Network (SAN).
- **SmartMessage Receiver.** SmartMessage receivers listen for encrypted messages from ArcSight SmartConnectors.

Creating a receiver is a three-step process:

- 1 Create a named receiver of a certain type. Receiver type cannot be changed after the receiver is created. New receivers are initially disabled. See [“To create a receiver” on page 194](#) for more information.
- 2 Add type-specific parameters. Receiver parameters are documented in [Table 6-1, “Receiver Parameters,” on page 195](#).
- 3 Enable the new receiver.

Maximum number of receivers that can be created on Logger: The number is limited by system resources—memory, CPU, disk input/output and possibly network bandwidth.



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.

Receivers Forwarders ESM Destinations Certificates						
Add						
Name	Type	IP Address	Port			
TCP Receiver 1	TCP Receiver	All	6001			
TCP Receiver 2	TCP Receiver	All	6002			
TCP Receiver 3	TCP Receiver	All	6003			
TCP Receiver 4	TCP Receiver	All	6004			
TCP Receiver 5	TCP Receiver	All	6005			
Udp1	CEF UDP Receiver	All	514			
Udp2	CEF UDP Receiver	All	527			
udp3	CEF UDP Receiver	All	552			
udp4_syslog	UDP Receiver	All	1143			
udp5_syslog	UDP Receiver	All	1216			
Add						

Figure 6-5 Receivers page



TCP 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.

To create a receiver



Before creating a Receiver of type File Receiver:


- For the Logger appliance, set up a Network File System mount. See [“Storage” on page 263](#).
- 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.




Create a Receiver of type **SmartMessage** before configuring the SmartConnector that will send to it. Once the Receiver is created, configure the SmartConnector as described in [“Installing SmartConnectors to Send Events to Logger” on page 29](#) and specify:

- Logger IP or hostname
- Port 443 (port must be 443)
- Receiver name

If the Receiver name changes on the Logger, it must be changed in the SmartConnector.

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input/Output** (left panel) > **Receivers** tab (right panel).
- 3 Click **Add**.
- 4 Enter a name for the new receiver and choose UDP, TCP, File Transfer, File Receiver, SmartMessage, CEF UDP, or CEF TCP type.
- 5 Click **Next** to edit receiver parameters listed in [Table 6-1 on page 195](#).
- 6 Click **Save**.
- 7 New receivers are initially disabled. Click the disabled icon () to enable the new receiver.

To edit a receiver

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input/Output** (left panel) > **Receivers** tab (right panel).
- 3 Find the receiver to be edited in the table.
- 4 Click the Edit icon (). Change the information in the form and click **Save**.

To delete a receiver


- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input/Output** (left panel) > **Receivers** tab (right panel).
- 3 Find the receiver to be deleted in the table.
- 4 Click the Delete icon (). Confirm the delete.

Table 6-1 Receiver Parameters

Parameter	Receiver Types	Description
Name	All	The name of the Receiver, used for reporting and status monitoring. SmartMessage receiver names are used to configure the ArcSight SmartConnector.
Type		<p>The Type of a Receiver cannot be changed after the receiver is created.</p> <p>UDP</p> <p>TCP</p> <p>CEF UDP</p> <p>CEF TCP</p> <p>File Transfer (scp/sftp/ftp)</p> <p>File Receiver (Remote File System)</p> <p>SmartMessage</p>
Protocol	File Transfer	Select SCP, SFTP or FTP protocol.

Parameter	Receiver Types	Description
Ip/Host	All except File Receiver and SmartMessage	<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 254.</p>
Character Encoding	All except File Transfer	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.
Port	UDP, TCP, File Transfer	The default port is 514. (For SmartMessage receivers, configure the SmartConnector for port 443.)
User	File Transfer	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	File Transfer	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.)
FilePath	File Transfer	<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>
Schedule	File Transfer	<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 258 before you schedule a file transfer.</p>

Parameter	Receiver Types	Description
Zip Format	File Transfer	Choose gzip, zip, or none.
RFS Names	File Receiver	<p>For the Logger appliance:</p> <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 263. To mount CIFS shares, see "CIFS Settings" on page 263.</p> <p>For more information about SAN, see "SAN" on page 267.</p> <p>For the software version of Logger:</p> <p>You can only 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	File Receiver, File Transfer	<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
Wildcard	File Receiver	<p>Regular expression describing the log files to read. Note: This is a regular expression, not a typical file wildcard like <code>"*.txt"</code>.</p> <p>Example: <code>.*\.process</code> (all files ending with <code>.process</code>). The wildcard for Symantec Anti-Virus log files would be <code>\d{8}.log</code>.</p> <p>The default is <code>.*</code>, meaning all files.</p>
Mode	File Receiver	<p>Mode is one of:</p> <ul style="list-style-type: none"> Delete - delete the log file once it has been processed Rename - rename the log file once it has been processed. The file is named by appending the Rename Extension. Persist - Logger remembers which files have been processed and only processes them once.
Rename Extension	File Receiver (Mode=Rename)	The suffix to append to log files that have been processed.
Character Encoding	File Receiver	Select the type of character encoding from the drop-down list.
Delay after seen	File Receiver or File Transfer	<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>

Parameter	Receiver Types	Description
Date/time locale	File Receiver or File Transfer	Select a locale from the pulldown list, such as English (United States), Chinese (Hong Kong), Chinese (Taiwan), and so on.
Date/time format	File Receiver or File Transfer	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). See Step Table 6-3 for a list of format specifiers. The default is "" (no timestamp in log file).
Date/time zone	File Receiver or Transfer	Required if the timestamp in the log file does not specify a time zone. This parameter is ignored if either Date/time format or Date/time location regex are blank. The default is the time zone configured on the Logger (System Admin > Settings > Platform > Time/NTP) .
Date/time location regular expression	File Receiver or Transfer	A regular expression describing which characters represent the timestamp in the log file. For example: .*\[(.*)\].* 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. The default is "" (no timestamp in log file).

Date and Time Specification

To specify the date and time format so that it can be parsed from a file (File Receiver or File Transfer receivers), refer to [Table 6-3 on page 199](#).

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 6-2 on page 198](#) 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 6-2 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

Source	Date and Time Pattern
Wed, 4 Jul 2001 12:08:56 -0700	EEE, d MMM yyyy HH:mm:ss Z
010704120856-0700	yyMMddHHmmssZ
2001-07-04T12:08:56.235-0700	yyyy-MM-dd'T'HH:mm:ss.SSSZ

Table 6-3 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)	(Number)	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
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)

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 ESM, Logger might be used to forward events to a number of ESM Managers. Forwarder filters make it possible to split the flow between the Managers, using one Forwarder for each Manager.

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 will no longer forward events. The [\\$now](#) value, if specified in a time range, is **not** treated as a variable. Instead, the time when the forwarder was created or updated is assigned to [\\$now](#). For example, if the time when forwarder was created was 1:45 p.m. and

the time range specified in the forwarder is 10 a.m. to \$now, all events between 10 a.m. and 1:45 p.m. will be forwarded. After events within that time range have been forwarded, no additional events will be forwarded.

A forwarder's operation can be paused and resumed at any point in time. Additionally, if a forwarder fails during a forwarding operation and is restarted, event forwarding resumes from the point at which the failure had occurred.

Forwarder types include UDP, TCP, Connector Forwarder, and ArcSight ESM Forwarder:

- **UDP.** UDP Forwarders forward events as User Datagram Protocol messages, such as Syslog format datagrams.
- **TCP.** TCP Forwarders forward events as Transmission Control Protocol messages.
- **Connector Forwarder.** The Connector Forwarder sends events to the ArcSight Logger Streaming Connector.
- **ArcSight ESM.** The ArcSight ESM Forwarder sends Common Event Format (CEF) events to an ESM Destination.

Maximum number of forwarders that can be created on Logger: Although there is no limit on the number of forwarders you can configure, ArcSight recommends that you enable a maximum of 15 forwarders per Logger.



Tip

Wait a few minutes after enabling a forwarder before disabling it. 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.

Figure 6-6 Forwarders page

To create a forwarder

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input/Output** in the left panel.
- 3 In the Forwarders tab, click **Add**. The page shown in [Figure 6-6](#) is displayed.
- 4 Enter a name for the new forwarder and choose UDP Forwarder, TCP Forwarder, Connector Forwarder, or ESM Forwarder type.
- 5 Click **Next**.
- 6 Enter additional, type-specific information as described in [Table 6-4, "Forwarder Parameters," on page 201](#). Click **Save**.


- 7 New forwarders are initially disabled. Click the disabled icon (🔒) to enable the new forwarder.

Table 6-4 Forwarder Parameters

Parameter	Forwarder Types	Description
Name	All	The name of the Forwarder, used for reporting and status monitoring.
Type		<p>The Type of a Forwarder cannot be changed after the forwarder is created.</p> <p>UDP</p> <p>TCP</p> <p>Connector Forwarder</p> <p>ArcSight ESM (CEF)</p>
Query Terms	All	Specify the events to be forwarded. See “Searching for Events on Logger” on page 72 . Forwarder queries can be constrained by Device Groups and Storage Groups, but not by Peers. See Figure 6-7 .
Filter	All	A filter that specifies which events to forward. (See “Filters” on page 222 .) ESM forwarders always filter out non-CEF events.
Filter by time range	All	<p>Check this box to specify a time range of events to be sent by the forwarder. When this box is checked, Start and End date and time fields appear.</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> <p>Once a forwarder has forwarded all events within a time range, it will no longer forward events. The <code>\$now</code> value, if specified in a time range, is not treated as a variable. Instead, the time when the forwarder was created or updated is assigned to <code>\$now</code>. For example, if the time when forwarder was created was 1:45 p.m. and the time range specified in the forwarder is 10 a.m. to <code>\$now</code>, all events between 10 a.m. and 1:45 p.m. will be forwarded. After events within that time range have been forwarded, no additional events will be forwarded.</p>
Source Type	Connector	<p>Select from the pulldown list of log file types, including:</p> <p>Apache HTTP Server Access</p> <p>Apache HTTP Server Error</p> <p>Juniper Steel-Belted Radius</p> <p>Microsoft DHCP Log</p> <p>IBM DB2 Audit</p> <p>Note: Source Type must be the same in Receiver, Forwarder, and SmartConnector. See “Forwarding Log File Events to ESM” on page 207.</p>

Parameter	Forwarder Types	Description
Preserve Syslog Timestamp	UDP, TCP, ESM	Set to true to preserve the syslog timestamp. The default is true--the timestamp is the original receipt time of the event. If set to false, original timestamp is replaced with Logger's receipt time.
Preserve Original Syslog Sender	UDP, TCP, ESM	Set to true to preserve the original sender. The default is true--the sender is the original sender. If set to false, the original sender information is replaced with Logger's information.
IP/Host	UDP, TCP, Connector	The destination to receive forwarded events
Port	UDP, TCP, Connector	The destination port to receive forwarded events
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 203.)

To edit a forwarder

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input/Output** in the left panel.
- 3 In the Forwarders tab, locate the forwarder you want to edit and click the **Edit** icon (). The screen shown in [Figure 6-7](#) is displayed.

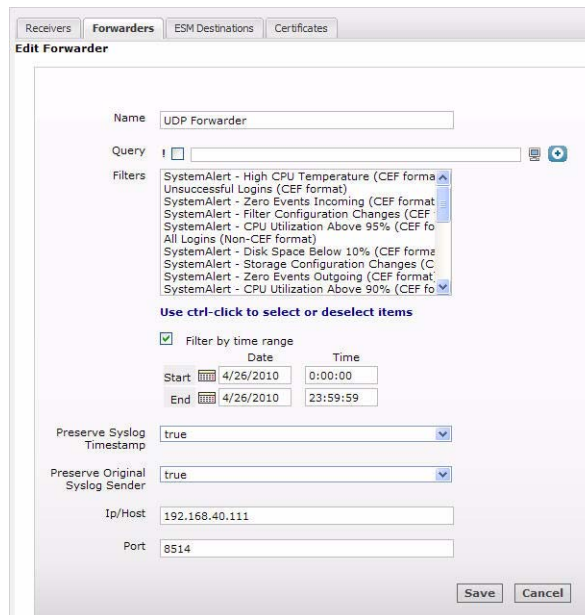




Figure 6-7 Specifying Query Terms, Filters, and other Forwarder parameters.


- 4 Edit the information in the form, as described in [Table 6-4 on page 201](#), and click **Save**.

- 5 If the forwarder is enabled, click to disable it. Then, click the disabled icon () to re-enable the forwarder and commit the changes.


To delete a forwarder

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input/Output** in the left panel.
- 3 Locate the forwarder to be deleted in the table.
- 4 Click the Delete icon () . Confirm the delete.

To pause a forwarder

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input/Output** in the left panel.
- 3 Locate the forwarder to be paused from the list of forwarders.
- 4 Click the Pause icon () .

To resume a forwarder

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input/Output** in the left panel.
- 3 Locate the forwarder whose operation you want to resume.
- 4 Click the Resume icon () .

ESM Destinations

ESM Destinations establish a connection to an ArcSight ESM 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 (see ["Common Event Format" on page 373](#)) that are not normalized or categorized.

Maximum number of ESM destinations that can be configured: As many allowable on the SmartConnectors you are using.

To setup Logger to forward events to an ArcSight ESM Manager

- 1 Copy the server SSL certificate file from an ESM 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 205](#).

If your Logger operates in FIPS mode, a valid and current (non-expired) server SSL certificate file from the ESM Manager is required on the Logger; otherwise, the forwarder will not forward events to it.

Note: Starting with Logger v4.0, 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 204](#).

- 3 Create an ESM Forwarder that refers to this ESM Destination. (See [“Forwarders” on page 199](#)).

The screenshot shows the 'Add ESM Destination' form with the following fields and values:

- ESM Destination Name: n111-h248
- Connector Name: n111-h248
- Connector Location: /All Connectors/Devices
- Logger Location: QA Lab
- IP/Host: n111-h248
- Port: 8443
- User Name: admin
- Password: [masked]

Buttons: Save, Cancel

Figure 6-8 ESM Destinations page

To create an ESM Destination

Note: Make sure you have loaded the certificate file for ESM Manager as described in [“Uploading a Certificate to the Logger” on page 205](#) 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** from the top-level menu bar.
- 2 Click **Event Input/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 Input/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 6-8](#) is displayed.
- 4 Enter the following parameters:

Parameter	Description
Destination Name	The name for this ESM Destination. Note: Make sure the name or IP address you specify in this field is exactly the name or IP address configured on the ESM Manager. If the two names or IP addresses do not match, you will not be able to set up an ESM destination successfully.
Connector Name	The SmartConnector name.

Parameter	Description
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."
IP or Host	The ESM Manager to which the Forwarder will direct events.
Port	Typically 8443.
Login	The name of an existing User of the ESM Manager with administrator privileges.
Password	The password for the Login user.

5 Click **Save**.



If you receive the following error when adding a new ESM destination, make sure the host name you specified in the Destination Name field exactly matches the name configured on the ESM Manager.

There was a problem: Failed to add destination

To delete an ESM Destination

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input/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 Input/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 to be deleted and click the delete icon (✖) on that row.
- 4 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the ESM Destination.

Uploading a Certificate to the Logger

You need to upload a valid server SSL certificate file for the ESM Manager that you are establishing as a Logger destination.

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 ESM 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, ESM 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** from the top-level menu bar.
- 3 Click **Event Input/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 Input/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.

- 4 In the **ESM Destinations** 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 an ArcSight Logger Streaming SmartConnector that sends the events on to ArcSight ESM. Unlike other events that Logger receives, such as syslog, SmartMessage, or CEF, log file events must be parsed to determine event timestamp.

To forward log file events to ESM, configure the Receiver, Forwarder, and SmartConnector to accept the same Source Type (as described in “Receiver Parameters” on page 195).

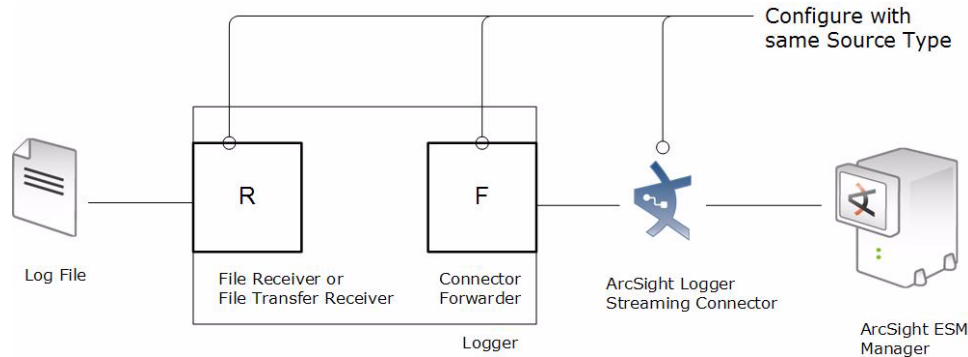


Figure 6-9 Configure the Receiver, Forwarder, and SmartConnector with the same Source Type to use Logger to forward log file events to ArcSight ESM.

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.

An alert is triggered if a specified number of matches occur within the specified threshold (time interval in seconds). When an alert triggers, an alert event is logged on the Logger and a notification is sent through previously configured destinations—e-mail addresses, SNMP server, Syslog server, and ESM Manager.

Logger provides two types of alerts:

- Real time alerts
- Saved Search Alerts

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.

Real Time Alerts

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.

Alerts are triggered in real time. That is, when specified number of matches occur within the specified threshold, an alert is **immediately** triggered.

To define a real time alert, you specify a query, match count, threshold, and one or more destinations.

A time range is not associated with the queries defined for these alerts. Therefore, whenever the specified number of matches occur within the specified threshold, an alert is triggered.

Saved Search Alerts

No limit on the number of configured e-mail destinations; however, you can only set one SNMP, one Syslog, and one ESM destination.

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.

These alerts are triggered at scheduled intervals. That is, when a specified number of matches occur within the specified threshold, an alert is triggered **at the next scheduled time interval**.

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.

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).

For example, if a Saved Search query has these start and end times:

Start Time: 5/11/2010 10:38:04

End Time: 5/12/2010 10:38:04

And, the number of matches and threshold are the following:

Match Count: 5

Threshold: 3600

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.

When an alert is triggered, Logger creates an alert event containing the trigger event. This alert event is also sent to the specified destinations if any are configured.

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 ArcSight Customer Support for instructions.

An e-mail message for an alert contains:

- The trigger alert information
- The matched 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]
```

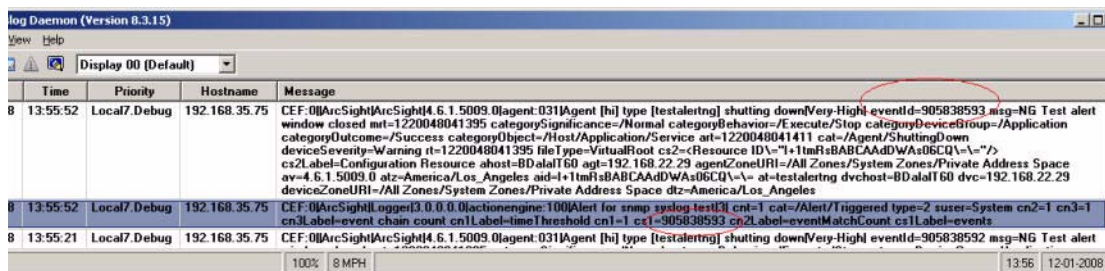
```
Event Receipt Time [Tue Nov 11 16:46:50 PST 2008]
```

```
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]
```

If you configure your Logger to include matched events for alerts sent to SNMP and Syslog destinations, make sure you are familiar with this information:

- Unlike an e-mail alert, a trigger alert is sent separately from the alert that contains the matched (base) events that triggered the alert. The trigger event includes the event IDs of all the matched events, as shown in the following example:



Time	Priority	Hostname	Message
8 13:55:52	Local7.Debug	192.168.35.75	CEF:0 ArcSight ArcSight 4.6.1.5009.0 agent:031 Agent [hi] type [testalertng] shutting down Very-High eventId=905838593 msg=NG Test alert window closed rt=1220048041395 categorySignificance=Normal categoryBehavior=/Execute/Stop categoryDeviceGroup=/Application categoryOutcome=/Success categoryObject=/Host/Application/Service art=1220048041411 cat=/Agent/ShuttingDown deviceSeverity=Warning rt=1220048041395 fileType=VirtualRoot cs2=Resource ID=1+1tmRsB8BCAAdDWAs06CQ\=-> cs2Label=Configuration Resource shost=BDalaIT60 agt=192.168.22.29 agentZoneURI=/All Zones/System Zones/Private Address Space av=4.6.1.5009.0 at=America/Los Angeles aid=1+1tmRsB8BCAAdDWAs06CQ\=-> al=testalertng dvchost=BDalaIT60 dvc=192.168.22.29 deviceZoneURI=/All Zones/System Zones/Private Address Space dtz=America/Los Angeles
8 13:55:52	Local7.Debug	192.168.35.75	CEF:0 ArcSight Logger 3.0.0.0 actionengine:100 Alert for snmp syslog test 3 cat=1 cat=/Alert/Triggered type=2 user=System cn2=1 cn3=1 cn3Label=event chain count cn1Label=timeThreshold cn1=1 cs1=905838593 cn2Label=eventMatchCount cs1Label=events
8 13:55:21	Local7.Debug	192.168.35.75	CEF:0 ArcSight ArcSight 4.6.1.5009.0 agent:031 Agent [hi] type [testalertng] shutting down Very-High eventId=905838592 msg=NG Test alert

- 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.
- 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.
- SNMP uses UDP to send packets. As a result, the order in which alerts arrive at an SNMP destination is not guaranteed. Use the event IDs in the trigger alert to identify its associated base events.

Similarly, 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.

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 [“SMTP Settings” on page 258](#). See also [“SNMP Destinations” on page 218](#), [“Syslog Destinations” on page 219](#), and [“ESM Destinations” on page 203](#).

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 211](#).

- 1 Click **Configuration** 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 211](#) 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.

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 43](#).



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.

- 6 Enter notification destinations. Enter any combination of:
 - ◆ One or more e-mail addresses, separated by commas
 - ◆ An SNMP Destination
 - ◆ A Syslog Destination
 - ◆ An ESM Manager

7 Click **Save**.

The screenshot shows the ArcSight Logger Configuration - System Admin interface. The 'Alert' tab is selected, and the 'Add Alert' dialog is open. The dialog contains the following fields and options:

- Name:** A text input field.
- Query:** A text input field with a search icon and a plus icon.
- Filters:** A list box containing the following items:
 - All Logins (CEF format)
 - All Logins (Non-CEF format)
 - CEF
 - High and Very High CEF Events
 - Malicious Code (CEF format)
 - Successful Logins (CEF format)
 - Successful Logins (Non-CEF format)
 - SystemAlert - CPU Utilization Above 90% (CEF fo
 - SystemAlert - CPU Utilization Above 95% (CEF fo
 - SystemAlert - Device Configuration Changes (CE
- Match Count:** A text input field.
- Threshold (sec):** A text input field.
- Email Address(es):** A text input field.
- SNMP Destination:** A dropdown menu with 'NONE' selected.
- Syslog Destination:** A dropdown menu with 'NONE' selected.
- ESM Destination:** A dropdown menu with 'NONE' selected.
- Buttons:** 'Save' and 'Cancel' buttons at the bottom right.

Figure 6-10 Add Alert dialog**To Enable or Disable a Real Time Alert**

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Locate the Alert to be disabled or enabled. Click the associated icon (🚫 or ✅) to enable or disable the Alert.



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** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Locate the Alert that you want to edit.
- 4 Click the Edit icon (✎). A screen similar to that shown in [Figure 6-10 on page 211](#) appears. Remember that only alphanumeric characters can be used in an Alert name.

To Remove a Real Time Alert

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Locate the Alert to be removed 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** from the top-level menu bar.
- 2 Click **Alerts** in the left panel. The Alerts list is displayed, as shown in Figure 6-11.

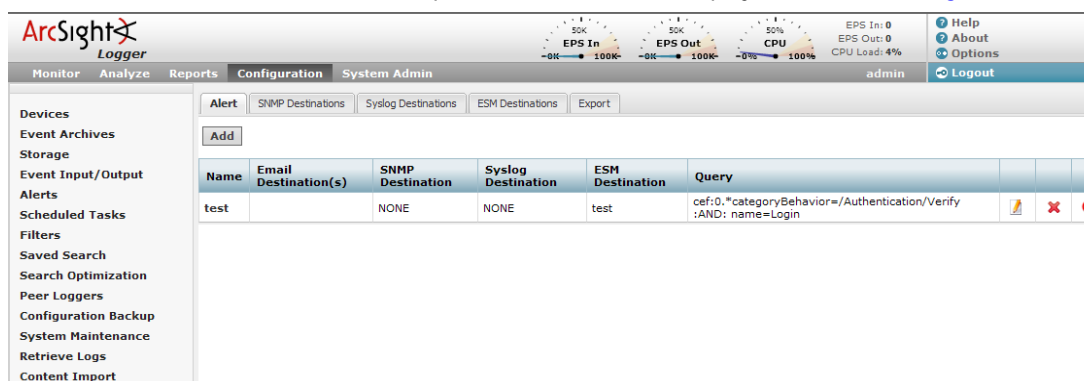


Figure 6-11 Alert list

Creating and Managing Saved Search Alerts

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 225](#). 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 occur within the specified threshold).

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 > Saved Search > Scheduled Searches)

To create a Saved Search Alert from the search results page

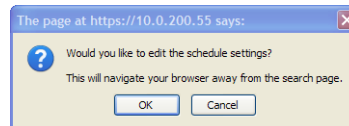
- 1 Run a search, as described in [“Searching for Events on Logger” on page 72](#).
- 2 Click the Save icon (💾) and enter the following settings.

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”.

Parameter	Description
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 prespecified location. Scheduled alerts run a search on a predetermined schedule and generate an alert if the specified number of events within the specified threshold are 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.

Parameter	Description
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>

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 225 .
----------------	--

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
------------------	--

Syslog destination	(Optional) A syslog server address to which the alert will be sent
--------------------	--

ESM destination	(Optional) An ESM Manager address to which the alert will be sent
-----------------	---

To create a Saved Search Alert from the Scheduled Searches page

- 1** Click **Configuration** 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**.

5 Enter the following information.

Saved Searches **Scheduled Searches** Saved Search Files (Logger)

Name:

Schedule:

Saved Searches:

- Invasion
- Scrutiny
- test
- Top Source Addresses - IDS

Use ctrl-click to select or deselect items

Job type:

Alert Options

Match count:

Threshold (sec):

Email address(es):

SNMP destination:

Syslog destination:

ESM destination:

Save Cancel

Parameter	Description
-----------	-------------



Name	A name for the Saved Search you are saving.
------	---

Parameter	Description
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>
Saved Searches	<p>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 225.</p> <p>Note: You can only select one Saved Search for each Alert you configure.</p>
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
Syslog destination	(Optional) A syslog server address to which the alert will be sent
ESM destination	(Optional) An ESM Manager address to which the alert will be sent


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 218](#).


To Enable or Disable a Saved Search Alert

- 1 Click **Configuration** 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** 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 215](#).
- 6 Click **Save**.

To remove a Saved Search Alert

- 1 Click **Configuration** 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** 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.

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.

To Add an SNMP Destination

- 1 Click **Configuration** 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** 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 to be removed and click the remove icon (✖) on that row.

5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the SNMP Destination.

Syslog Destinations

Syslog Destinations describe how Alert notifications should be sent using the comparatively simple Syslog protocol. Set up Syslog Destinations before creating Alerts that will use them.

To Add a Syslog Destination

1 Click **Configuration** 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
Send Syslog Timestamp	True or false (default is false). If false, the syslog message will have the current Logger time.
Send Original Syslog Sender	True or false (default is false). If false, the syslog message will be sent from Logger.
Ip/Host	Host name or IP address
Port	Port (default is 514)

- 7 Click **Save** to create the new Syslog Destination

To Edit a Syslog Destination

- 1 Click **Configuration** 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** 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 to be removed and click the remove icon () on that row.
- 5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Syslog Destination.

ESM Destinations

ESM Destinations describe how Alert notifications should be sent to an ArcSight ESM Manager. Set up ESM Destinations before creating Alerts that will use them. If an ESM Manager uses a signed SSL certificate, you will need to load it on the Logger.

To setup Logger to send alerts to an ArcSight ESM Manager

- 1 If the ESM Manager uses a certificate, copy the server SSL certificate file from an ESM 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 205](#).

Note: Starting with Logger v4.0, 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 204](#).

Export

See [“Exporting and Importing Content” on page 247](#).

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 258](#) 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 > Scheduled Tasks**.
- 2 Filter the list by selecting a specific type of Scheduled Task from the drop-down list, or select **All**.



Task	Type	Schedule	Next Run Time		
job_local	ScheduledSearch	Every hour	Mon Jun 25 12:00:00 PDT 2007		
job_remote	ScheduledSearch	Every hour	Mon Jun 25 12:00:00 PDT 2007		


Figure 6-12 Scheduled Tasks page

To add a Scheduled Task

Scheduled Tasks can be created for:

- Saved Search (See [“Scheduled Saved Search” on page 226](#))
- File Receivers and File Transfer Receivers (See [“Receivers” on page 193](#))
- Event Archive (See [“Event Archives” on page 185](#))
- Configuration Backup (See [“Configuration Backup and Restore” on page 235](#))

To delete a Scheduled Task

- 1 Click the **Configuration > Settings** tab, click **Scheduled Tasks** in the sub-menu, then click the **Scheduled Tasks** tab.
- 2 Locate the Scheduled Task to be deleted and click the delete icon () on that row.
- 3 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 the **Configuration > Settings** tab, click **Scheduled Tasks** in the sub-menu, then click the **Currently Running Tasks** tab.
- 2 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

- 1 Click the **Configuration > Settings** tab, click **Scheduled Tasks** in the sub-menu, then click the **Finished Tasks** tab.
- 2 Filter the list by selecting a specific type of Scheduled Task from the drop-down list, or select **All**.

Filters

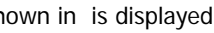
The Filters section has three tabs: Filters, Search Group Filters, and Export.

Filters

The Filters page provides a convenient place to manage filters. There are two types of filters: shared and Search Group. You can also create and delete shared filters on the Analyze page. Shared filters are optional; they provide a way to focus on events of interest. Search Group filters are not optional—they limit the events that users in a particular user group are able to see as an access control mechanism, those users can not opt to remove the filter. Search Group filters can also be used to limit the events processed by a category of reports (see [“Using Report Category Filters” on page 180](#)).

A set of predefined 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 86](#).

To create a filter

- 1 Click the **Configuration > Settings** tab, click **Filters** in the sub-menu, then click the **Filters** tab.
- 2 Click **Add**. The page shown in  is displayed.
- 3 Enter a name for the new filter in the Name field.

Filter names are case-sensitive.
- 4 If you are creating a shared filter, select **Unified**.

For Search Group filters, select **Search Group**. Additionally, non-administrator users cannot create Search Group filters.
- 5 Click **Next**.
- 6 If you selected Unified method in the previous step, enter the query for the new filter.

Click Advanced Search to use the Search Builder Tool to create the query. For details about using the Search Builder Tool, see [“Using the Search Builder Tool” on page 64](#).

For instructions on creating a query, see [“Searching for Events on Logger” on page 72](#).

The new filter will eliminate events that do not match the query.

- 7 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” on page 223](#).

To create a filter by copying an existing filter

- 1 Click the **Configuration** tab, click **Filters** in the sub-menu, then click the **Filters** tab.
- 2 Locate the filter to copy from the list of filters on the Filters page. 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, then click the **Filters** tab.
- 2 Find the filter to be edited in the table.
- 3 Click the Edit icon (). 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 to be deleted in the table.
- 3 Click the Delete icon (). Confirm the delete.

Search Group Filters

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 page 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 180](#).

- **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.

The Search Group Filters page is used to manage the association of User Groups with Search Group Filters.

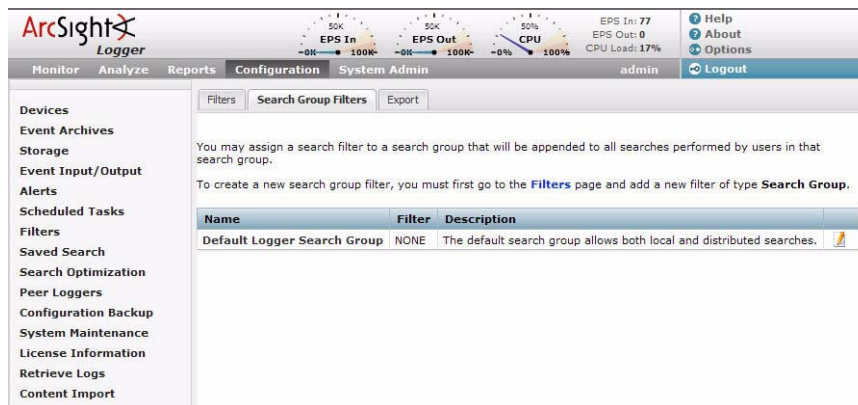


Figure 6-13 Search Group Filters Page




In the Search Group Filters page (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.

To create, edit, or delete Search Group Filters, see [“Filters” on page 222](#). To create, edit, or delete User Groups, see [“User Groups” on page 284](#).



Only users that are members of a System Admin group can assign Search Group Filters. For more information, see [“User Groups” on page 284](#).

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 [“User Groups” on page 284](#).
- 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 222](#). 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, 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**.

Export

See [“Exporting and Importing Content” on page 247](#).

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 207](#).

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 258](#) 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 > 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. Or 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.
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.

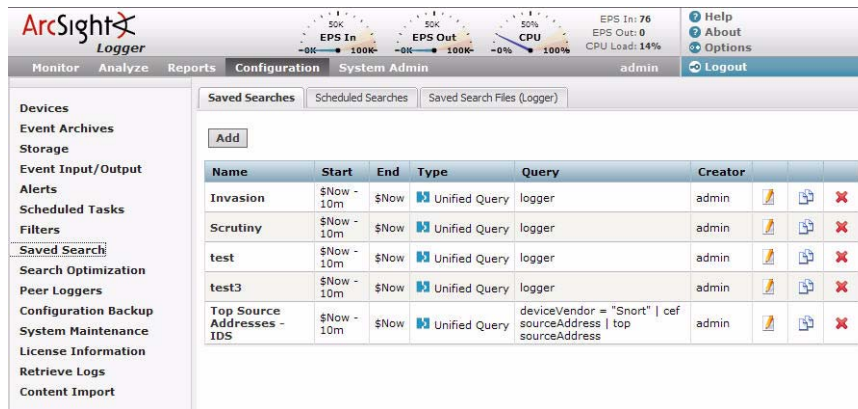


Figure 6-14 Saved Search page

To edit a Saved Search

- 1 Click the **Configuration > Saved Search**.
- 2 Find the Saved Search to be edited in the table.
- 3 Click the Edit icon (). Change the information in the form and click **Save**.

To delete a Saved Search

- 1 Click the **Configuration > Saved Search**.
- 2 Find the Saved Search to be deleted 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 212](#).

To add a scheduled Saved Search

- 1 Click **Configuration** 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.

Saved Searches

Scheduled Searches

Saved Search Files (Logger)

Add Scheduled Search

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

Export directory name

Fields

Event Time, Receipt Time, Device, Logger, Name, Version, Device Vendor, Device Product, Device Version, Signature ID, Severity

☒All fields

Clear

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.


Parameter	Description
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>
Saved Searches	<p>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 225.</p>
Job Type	<p>Select Search for a scheduled Saved Search.</p>
Export Options	<p>For the Logger appliance:</p> <p>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 229) feature to access those files. <p>For the software version of Logger:</p> <p>The only applicable option is “Save to Logger”, which is preselected for you.</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>
Fields	<p>Edit the list of fields desired for output or check the All Fields box. Click the Clear link to clear the text box.</p>
Include Summary	<p>Check this box to include an event count with the Saved Search, or a total when more than one Saved Search is specified.</p>
Include Non-CEF Events	<p>Check this box to include all events. Uncheck the box to include only CEF (see “Common Event Format” on page 373) events in the output.</p>

- 6 Click **Save** to add the new scheduled Saved Search, or **Cancel** to quit.

To edit a scheduled Saved Search

- 1 Click **Configuration** 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 to be edited 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** 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 to be deleted 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.








Saved Search Files						
Saved Search files. Under /opt/arc4sight/logger/user/logger/data/savedsearch.						
Name	Last Modified	Size	State	Error Message		
job_local_2007-06-22 17-00-00.csv	Fri Jun 22 17:00:02 PDT 2007	202 bytes	Exported		Retrieve	
job_local_2007-06-24 13-00-04.csv	Sun Jun 24 13:06:49 PDT 2007	202 bytes	Exported		Retrieve	
job_local_2007-06-22 07-00-00.csv	Fri Jun 22 07:00:02 PDT 2007	202 bytes	Exported		Retrieve	
job_local_2007-06-25 01-00-00.csv	Mon Jun 25 01:17:10 PDT 2007	202 bytes	Exported		Retrieve	
job_local_2007-06-24 11-00-00.csv	Sun Jun 24 11:12:35 PDT 2007	202 bytes	Exported		Retrieve	
job_local_2007-06-22 02-00-00.csv	Fri Jun 22 02:00:02 PDT 2007	202 bytes	Exported		Retrieve	
job_local_2007-06-24 23-00-01.csv	Sun Jun 24 23:17:38 PDT 2007	202 bytes	Exported		Retrieve	

Figure 6-16 Saved Search Files page

Search Optimization

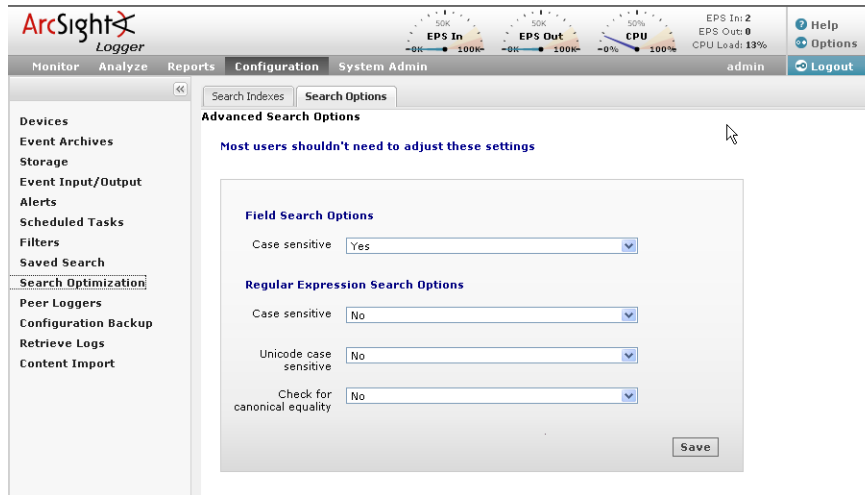
The search optimization option enable you to:

- Add search indexes for field query search operations
- Tune advanced search options
- Delete custom field sets

Add Search Indexes

See ["Indexing" on page 80](#) for more information.

Tuning Advanced Search Options



The following table lists the advanced search options you can view and configure. These options support i18n choices. If you change any of these options, you will need to reboot your Logger for those changes to take affect.

Option	Description
Case sensitive	<p>Defaults: Yes, for field query; No, for regular expression.</p> <p>Full-text search (keyword search) is case insensitive. You cannot change its case sensitivity.</p> <p>When this option is set to No, searching for “login” will find “login,” “Login,” and “LOGIN”.</p> <p>Note: Case-sensitive search only applies to the local Logger. Peer loggers will continue to use case-insensitive search.</p> <p>Set this option to Yes to increase local query performance.</p>
Unicode case sensitive	<p>Default: No</p> <p>Set to Yes if non-English events should be compared in a case-sensitive way.</p> <p>This option only applies to the regular expression search method.</p> <p>Note: ArcSight strongly recommends that you do not change this option.</p>
Check for canonical equality	<p>Default: No</p> <p>Set to Yes if non-English events should be compared using locale-specific algorithms.</p> <p>This option only applies to the regular expression search method.</p> <p>Note: ArcSight strongly recommends that you do not change this option.</p>

To change any of the above options, click **Configuration > Search Optimization > Search Options** tab (selected by default).

Deleting Custom Field Sets



You need to have the “Edit, save, and remove fieldsets” privilege to delete a custom field set.

Note

To delete a custom field set:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Search Optimization** from the left panel.
- 3 In the Fieldsets tab, identify the field set you want to delete and click the delete (✕) icon.
- 4 Confirm the deletion.

Peer Loggers

A Logger can establish peer relationships with one or more Loggers to enable distributed event searches.

When two Loggers peer with each other, one Logger initiates the relationship. The initiator Logger sends the credentials to authenticate itself to the other Logger, called the remote Logger from hereon. If the authentication succeeds, a peer relationship is established between the two Loggers. (The remote Logger must have the credentials for the initiator Logger configured on it for authentication to succeed.)

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.

Peer Loggers can authenticate using any of these methods:



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.

- 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.



If the peer-relationship initiating Logger has version 3.0.x installed and the other Logger is running v4.0 GA with SSL Client Authentication (CAC) enabled, you can still use this authentication method. Enter the generated Authorization ID in the User Name field and the Code in the Password field on the v3.0.x Logger.

Guidelines

You should be aware of these guidelines when peering Loggers:

- You can peer a Logger to one or more remote Loggers.
- A Logger appliance can peer with a software version of Logger.
- Peer Loggers can run different versions. However, these are the only supported paths for running a search across peers:
 - ◆ A search from a v4.0.x Logger to v4.5
 - ◆ A search from v4.5 Logger to v4.0.x
 - ◆ A search from v4.5 Logger to v4.5
- A search that is run across peers cannot contain pipeline operators, discussed at [“Search Operators” on page 47](#).
- Currently, report generation across peer Loggers is not supported.
- If the remote Logger is configured for SSL Client authentication (CAC), you must configure an authorization ID and code on the initiator Logger. If the initiator Logger is v3.0.x, enter the authorization ID and code in the User Name and Password fields.
- 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 readd a peer.
- A user must belong to the Logger Search User Group with “Search for events on remote peers” privilege set to Yes.
- 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, UserA is restricted by a search group filter to only search for events in which deviceVendor is set to “Cisco”. When UserA performs a search operation across LoggerA's peers, the same constraint (to search events where deviceVendor = “Cisco”) is applied on all peers.

- 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. However, if you delete the peering relationship or it breaks for other reasons, you will need to enter the updated credentials to 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" on page 289](#).
- 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 234](#).
- 5 Add Logger B's information, as described in ["To add a peer Logger" on page 233](#):

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](#).

The screenshot shows the ArcSight Logger web interface. The top navigation bar includes 'Monitor', 'Analyze', 'Reports', 'Configuration', and 'System Admin'. The 'Configuration' tab is active, and the 'Peer Loggers' sub-tab is selected. On the left, a sidebar lists various system components like 'Devices', 'Event Archives', 'Storage', etc. The main content area displays the 'Add Peer Logger' dialog box. This dialog has fields for 'Peer Host Name', 'Peer Port' (set to 443), 'Peer Login Credentials' (selected with a radio button), 'Peer User Name', and 'Peer Password'. Below these fields, a note states: 'Following fields are for local (currently connected) logger and are optional. This needs to be changed only seldomly.' There are also fields for 'External IP Address' (192.168.36.42) and 'Local Port' (443). 'Save' and 'Cancel' buttons are at the bottom right of the dialog.

To add a peer Logger

- 1 Click **Configuration** 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.
Peer Login Credentials	Select Peer Login Credentials for password-based authentication with the remote Logger.
Peer Authorization Credentials	Select Peer Authorization Credentials for SSL client authentication with the remote Logger. (See "SSL Client Authentication (CAC Authentication)" on page 273.) If the peer-relationship initiating Logger has version 3.0.x installed and the other Logger is running v4.0 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 v3.0.x Logger.

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 234 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 234 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	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. 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.
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 generate Authorization ID and Code for configuring a peer relationship


Use the following procedure to generate the authorization ID and code on the Logger to which you are establishing a peer relationship. (Logger B in the example described earlier in this section.) This ID and Code is then configured on the Logger that initiates the peer relationship. (Logger A in the earlier example.)

- 1 Click **Configuration** 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. Cut and paste this information when adding a peer Logger that is configured to use SSL client authentication.

To delete a peer Logger

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Peer Loggers** from the left panel.
- 3 Locate the Peer to be deleted 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** > **Peer Loggers**).

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.

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	The following Report content only:
Logs	<ul style="list-style-type: none"> • Queries, Reports, Parameters, Parameter Value Groups, Dashboard
Global Settings	<ul style="list-style-type: none"> • Templates
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, Dashboard • 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.


Running a Configuration Backup (Ad-hoc or Scheduled)

The screenshot shows the ArcSight Logger interface with the 'Configuration Backup' dialog box open. The dialog has a left sidebar with a tree view containing: Devices, Event Archives, Storage, Event Input/Output, Alerts, Scheduled Tasks, Filters, Saved Search, Search Optimization, Peer Loggers, Configuration Backup (selected), Retrieve Logs, and Content Import. The main area is titled 'Edit Configuration Backup' and contains the following fields:

- Protocol: SCP (dropdown)
- Port: 22 (text box)
- Ip/Host: (empty text box)
- User: (empty text box)
- Password: (empty text box)
- Remote directory: (empty text box)
- Backup content: All (dropdown)
- Schedule: ☒ One time only

At the bottom right are 'Save' and 'Cancel' buttons.

To run a configuration backup or to edit the configuration backup settings:

- 1 Click the **Configuration > 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
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.

Parameter	Description
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 258.</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 settings specific to your environment 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 can only restore from configuration or report content using a backup file if the Logger appliance model and the version running on it is the same as the one used to create the backup file.

For the software version of Logger, the operating system and Logger version running on the machine to which you are restoring should be the same as the one used to create the backup file.

To restore from a configuration backup:

- 1 Click the **Configuration > 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

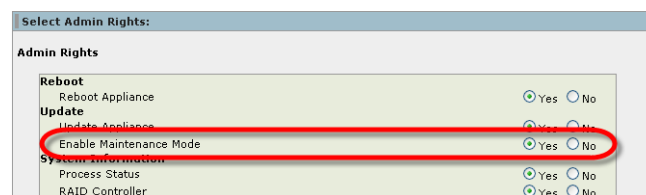
See [“Running a Configuration Backup \(Ad-hoc or Scheduled\)”](#) on page 236.

System Maintenance

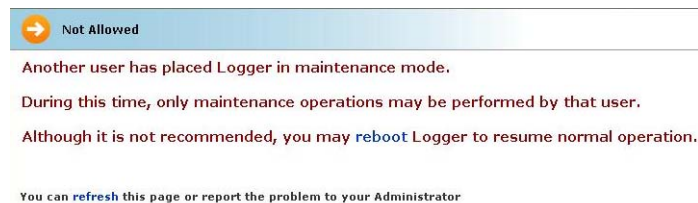
Certain operations on Logger, such as database defragmentation and extending the storage volume size, require that Logger be in a maintenance state—a state in which operations related to data on the Logger are not running. Maintenance mode enables you to place the Logger in such a state. When a Logger is in maintenance mode:

- Events are not processed
- Reports are not generated
- Search cannot run
- Scheduled jobs do not run

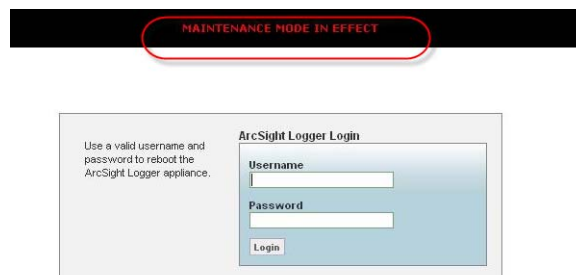
Logger users who will be performing operations that require it to be in maintenance mode must have the “Enable Maintenance Mode” privilege set to Yes (System Admin > User/Groups > Manage Groups > System Admin Group).



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:



Copyright © 2009 ArcSight Inc. All rights reserved.
ArcSight and the ArcSight logo are trademarks of ArcSight, Inc.

Entering Maintenance Mode

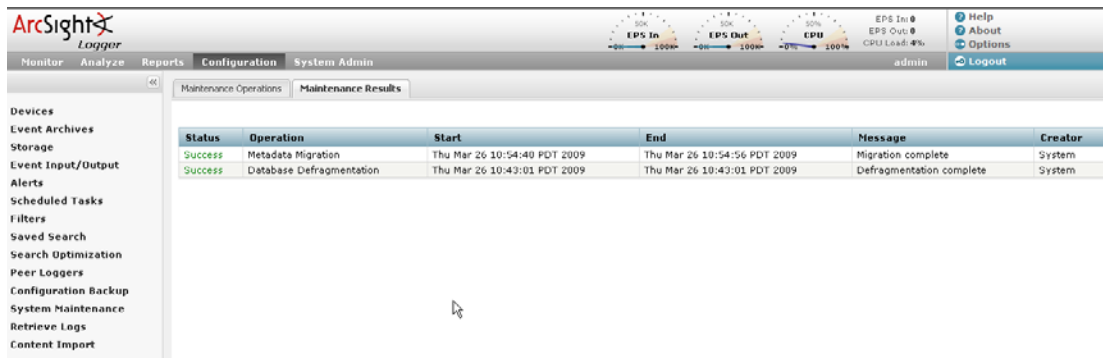
You cannot place a Logger in maintenance mode directly. A Logger can enter maintenance mode only when you perform an operation that requires it to be in that mode. For example, when defragmenting database, the user interface prompts you to enter Logger in maintenance mode, as illustrated in “Database Defragmentation” on page 239.

Exiting Maintenance Mode

To exit maintenance mode, reboot the 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 > System Maintenance > Maintenance Results**.



Status	Operation	Start	End	Message	Creator
Success	Metadata Migration	Thu Mar 26 10:54:40 PDT 2009	Thu Mar 26 10:54:56 PDT 2009	Migration complete	System
Success	Database Defragmentation	Thu Mar 26 10:43:01 PDT 2009	Thu Mar 26 10:43:01 PDT 2009	Defragmentation complete	System

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 is fragmented:

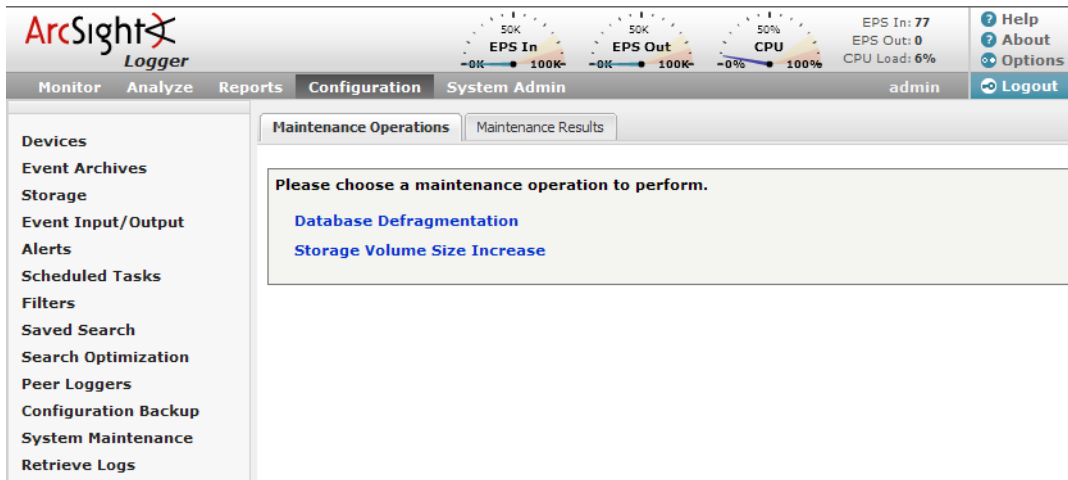
- Slow search and reporting
For example, even a search operation over the last two minutes of data is slow.
- Long pauses in the receiver and forwarder operations

Starting with this release, 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 238](#).
- 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 ArcSight 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 and restart the process from the beginning.

Defragmenting a Logger



To defragment a Logger:



- 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).
- If the defragmentation process fails at any point, reboot the Logger and restart the process from the beginning.

- 1 Click **Configuration > System Maintenance**.
- 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 238](#).



Database Defragmentation

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 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 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.

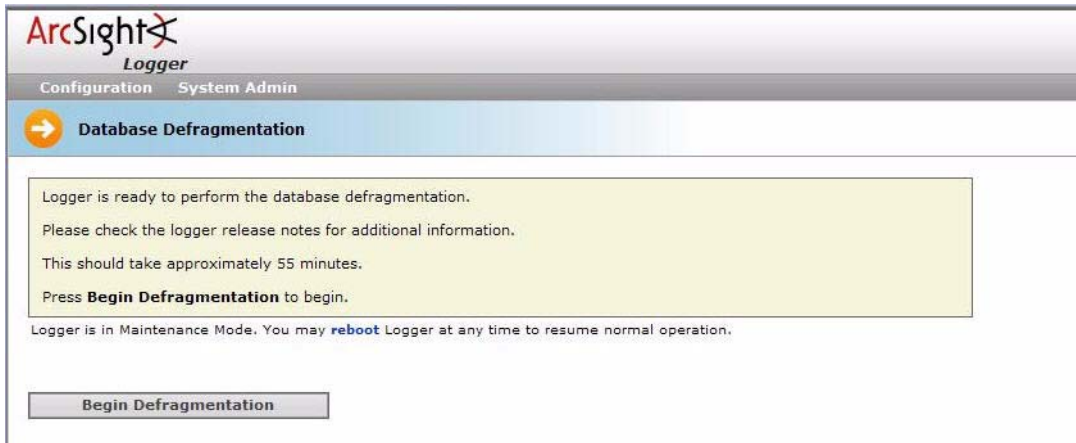


Figure 6-17 Begin Database Defragmentation

The defragmentation process starts. A progress indicator shows the status of defragmentation, as shown in the example below. ArcSight recommends that you do not attempt any operation on the Logger until defragmentation has completed.

Once defragmentation is complete, the Logger reboots automatically, thus exiting maintenance mode.



☀ 60.16% Defragmenting... (3 hours and 3 minutes elapsed)

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 displayed on L7100 Loggers as it is not applicable to those platforms.

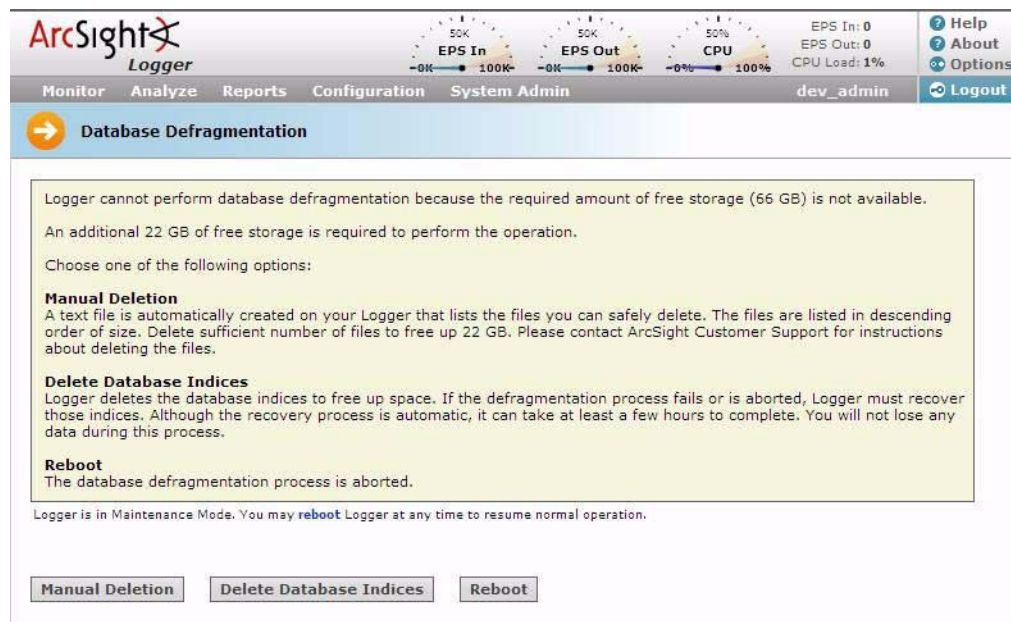
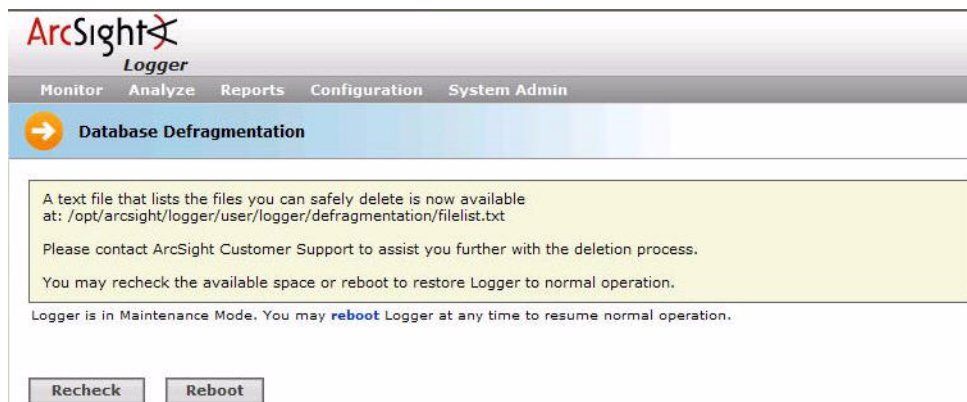


Figure 6-18 Required storage for Database Defragmentation is not available

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, as shown in the following figure.



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 ArcSight Customer support for instructions and guidance.

Follow these steps to proceed:

- i Leave the message screen without taking any action.
- ii Contact ArcSight 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 242](#). 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 241](#) is displayed. Click **Begin Defragmentation** to proceed further.

If sufficient storage is still not found, the screen in [Figure 6-18 on page 242](#) is displayed. Choose from the listed options to create additional space. See [“You can choose from one of the following options:” on page 242](#) for more information.



Note

If you need to exit the defragmentation process without creating sufficient storage, click **Reboot**.

◆ **Delete Database Indices**

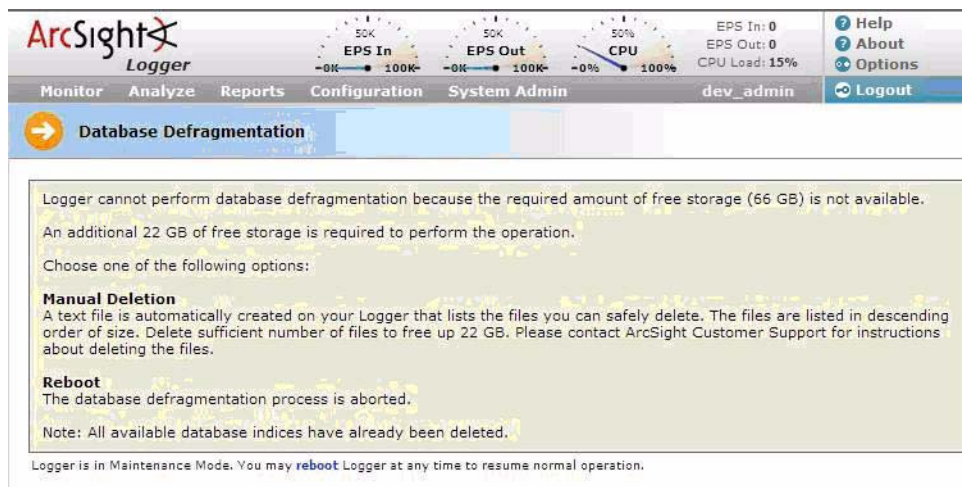
Logger automatically deletes a sufficient number of database indices, starting with the largest index, to free up the required amount of storage. If sufficient space becomes available after deleting database indices, defragmentation proceeds further automatically.

However, if sufficient storage is not available even after dropping database indices, the following screen is displayed.



Note

The Manual Deletion option (shown in the following figure) is not displayed on L7100 Loggers as it is not applicable to those platforms.



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 ArcSight 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 240](#).



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**

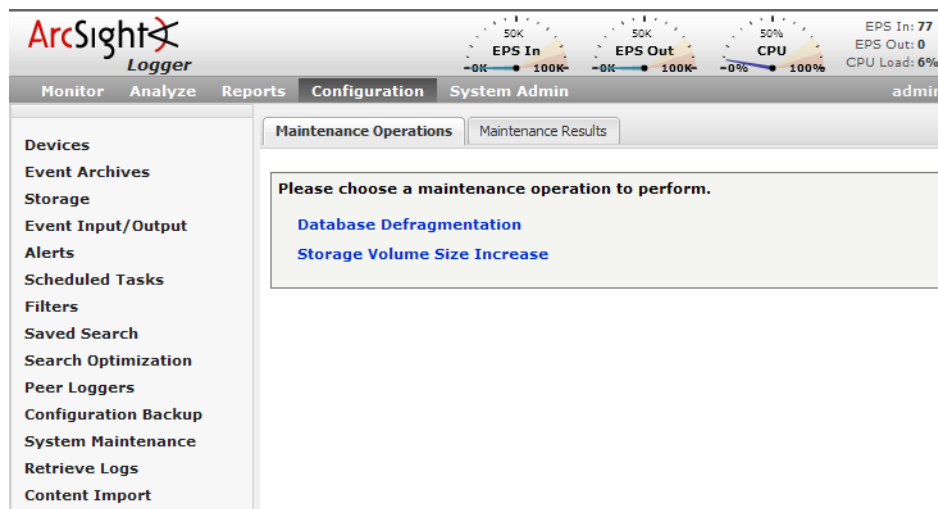
The database defragmentation process is aborted and Logger returns to the state it was in before you started the defragmentation utility.

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.



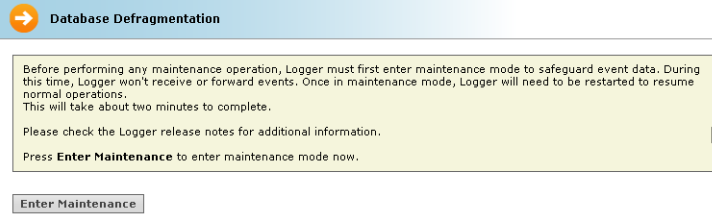
You can perform this process only if you have the “Enable Maintenance Mode” privilege set to Yes (System Admin > User/Groups > Manage Groups > System Admin Group).



To increase the size of a storage volume:

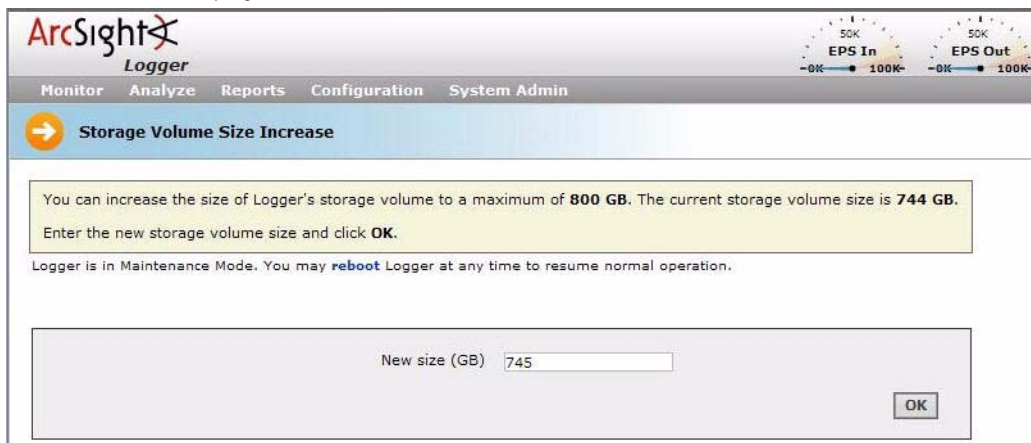
- 1 Click **Configuration > System Maintenance**.
- 2 Click **Storage Volume Size Increase**.
- 3 Click **Enter Maintenance** so that the Logger can enter maintenance mode.

For more information about maintenance mode, see [“System Maintenance” on page 238](#).

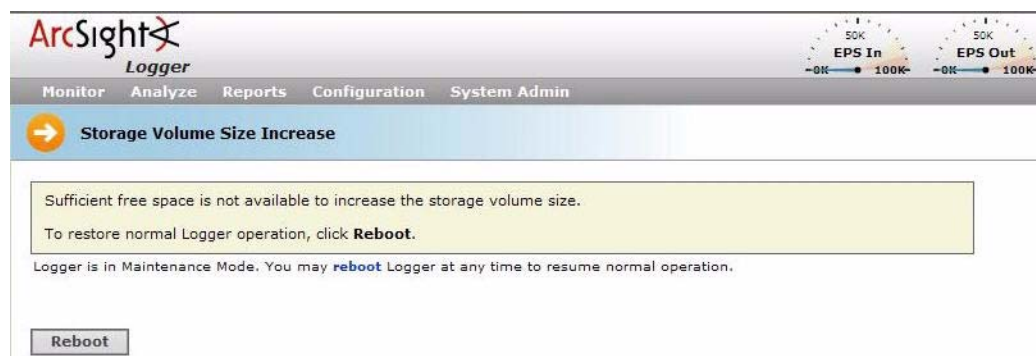


- 4 While entering the maintenance mode, Logger performs a check to determine if the storage volume size can be increased and by what amount.

If the storage volume can be increased, a message similar to the following is displayed. Enter the new size and click **OK**.



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.

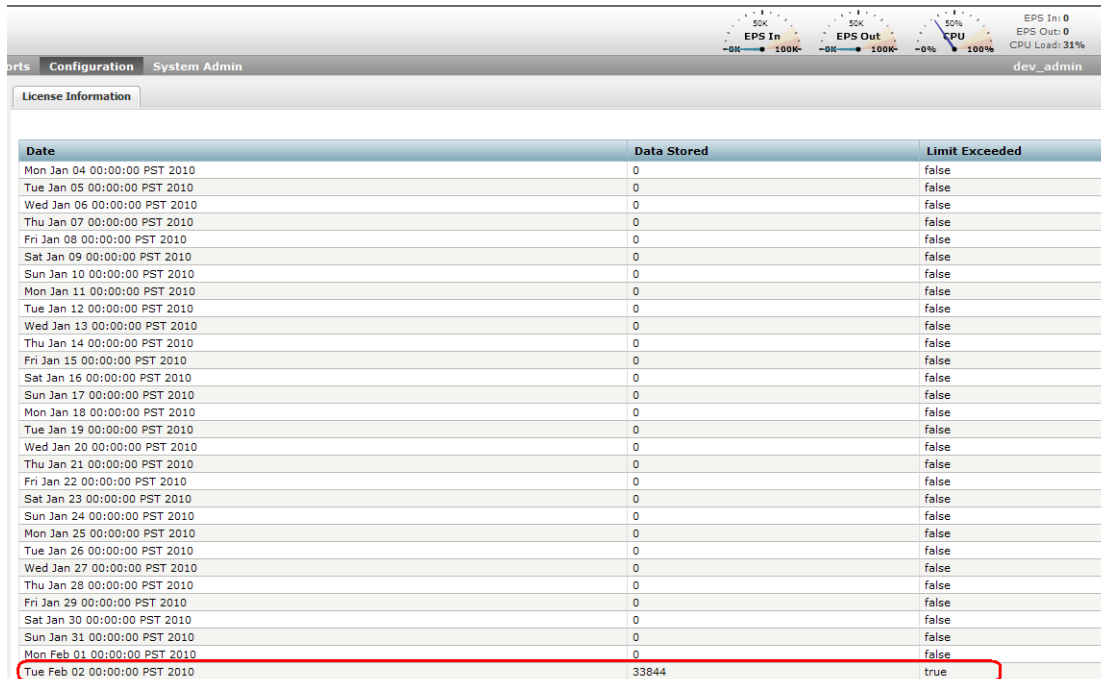


License Information

This user interface page is only available on the software version of Loggers and not on Logger appliances because on appliances, generally, a data storage limit is not imposed.

The License Information page (**Configuration > License Information**) 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.



Date	Data Stored	Limit Exceeded
Mon Jan 04 00:00:00 PST 2010	0	false
Tue Jan 05 00:00:00 PST 2010	0	false
Wed Jan 06 00:00:00 PST 2010	0	false
Thu Jan 07 00:00:00 PST 2010	0	false
Fri Jan 08 00:00:00 PST 2010	0	false
Sat Jan 09 00:00:00 PST 2010	0	false
Sun Jan 10 00:00:00 PST 2010	0	false
Mon Jan 11 00:00:00 PST 2010	0	false
Tue Jan 12 00:00:00 PST 2010	0	false
Wed Jan 13 00:00:00 PST 2010	0	false
Thu Jan 14 00:00:00 PST 2010	0	false
Fri Jan 15 00:00:00 PST 2010	0	false
Sat Jan 16 00:00:00 PST 2010	0	false
Sun Jan 17 00:00:00 PST 2010	0	false
Mon Jan 18 00:00:00 PST 2010	0	false
Tue Jan 19 00:00:00 PST 2010	0	false
Wed Jan 20 00:00:00 PST 2010	0	false
Thu Jan 21 00:00:00 PST 2010	0	false
Fri Jan 22 00:00:00 PST 2010	0	false
Sat Jan 23 00:00:00 PST 2010	0	false
Sun Jan 24 00:00:00 PST 2010	0	false
Mon Jan 25 00:00:00 PST 2010	0	false
Tue Jan 26 00:00:00 PST 2010	0	false
Wed Jan 27 00:00:00 PST 2010	0	false
Thu Jan 28 00:00:00 PST 2010	0	false
Fri Jan 29 00:00:00 PST 2010	0	false
Sat Jan 30 00:00:00 PST 2010	0	false
Sun Jan 31 00:00:00 PST 2010	0	false
Mon Feb 01 00:00:00 PST 2010	0	false
Tue Feb 02 00:00:00 PST 2010	33844	true

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. Figure 4-8 shows a typical example of a .zip archive of log files.

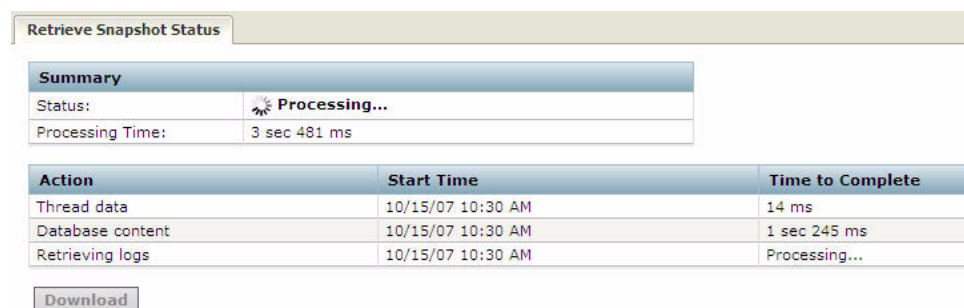
ArcSight Customer Support may ask you to retrieve logs as part of an incident investigation. If so, follow the steps below and upload the resulting .zip file to ArcSight Support.

To retrieve Logger system logs

- 1 Click the **Configuration > Retrieve Logs**.

The page shown in [Figure 6-19](#) 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		
Status:	Processing...	
Processing Time:	3 sec 481 ms	
Action	Start Time	Time to Complete
Thread data	10/15/07 10:30 AM	14 ms
Database content	10/15/07 10:30 AM	1 sec 245 ms
Retrieving logs	10/15/07 10:30 AM	Processing...
Download		

Figure 6-19 Retrieve Logs provides snapshot status.

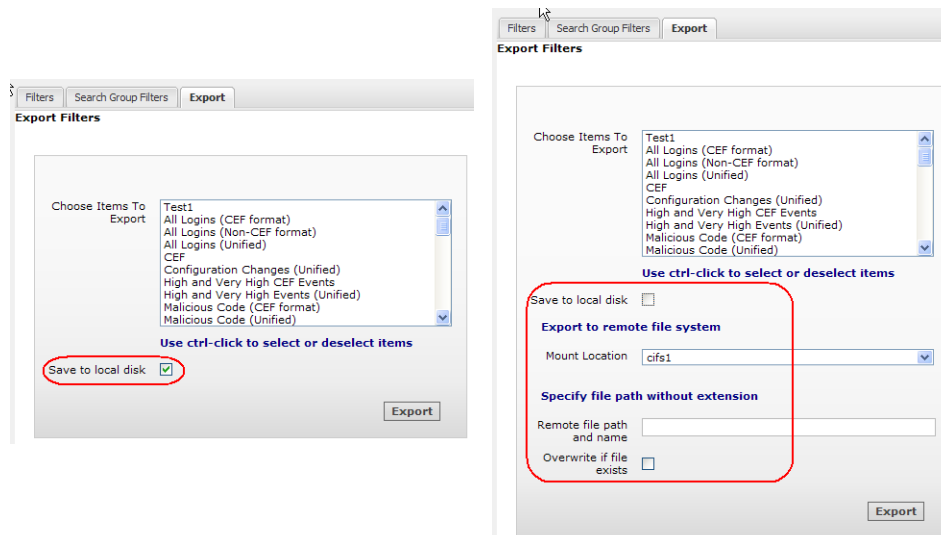
Exporting and Importing Content

You can export and import content (alerts and filters) from one Logger to another. Doing so is useful in these situations:

- The exported content serves as a backup for the 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 to the existing content on a Logger.

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.

Starting with Logger v4.0, saving content to the local disk is the default option. If you want to export to a remote location, you need to uncheck the “Save to local disk” option in the user interface to display the remote location options.



Guidelines for Exporting and Importing

Make sure you are familiar with these guidelines before exporting or importing content:

Exporting Guidelines

- 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.
- The information exported for an alert includes the query associated with the alert, match count, threshold, and status. It does not include e-mail, SNMP, and syslog destination information.
- The alert destinations (SNMP, Syslog, and SMTP servers) information is not exported; therefore, you will need to set this information for alerts you import.

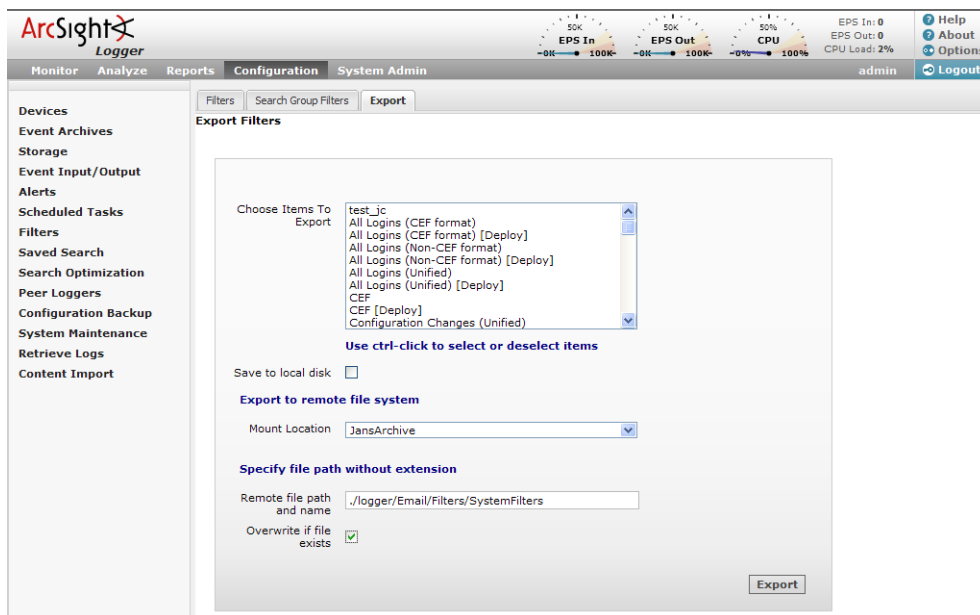
Importing Guidelines

- Existing content on a Logger is not deleted when new content is imported. The new content is added to the existing content.
- If an alert contains a filter, the filter is automatically created on the importing Logger. Such a filter is prefixed with "fwd" in the name. For example, "fwd-23456790".
- If an alert with the same name exists on the importing system, the alert being imported is named *AlertName[import]*. Similarly, an imported filter is named *FilterName[import]*.

If an alert with the name *AlertName[import]* exists on the importing Logger (from a previous import procedure), the alert being imported is named *AlertName[import][import]*. Similarly, a filter is named *FilterName[import][import]*.

- You will need to set the alert destinations (SNMP, Syslog, and SMTP servers) for alerts you import because this information is not included in the exported content.

Exporting Content



To export Alerts or Filters:

- 1 Click **Configuration** > **Alerts** (or **Filters**, for filters) > **Export** tab.
- 2 Select the Alerts or Filters to export in the Choose Items to Export field.

To select one alert (or filter), click its name.

To select multiple alerts (or filters), hold the **Ctrl** key down and click the names.

- 3 To save the exported content on the local disk of the computer from which you connect to the Logger, go to Step 7.

To export the content to a remote storage system, uncheck the "Save to local disk" field.

- 4 Select the location to which you want to export the content in the Mount Location field.

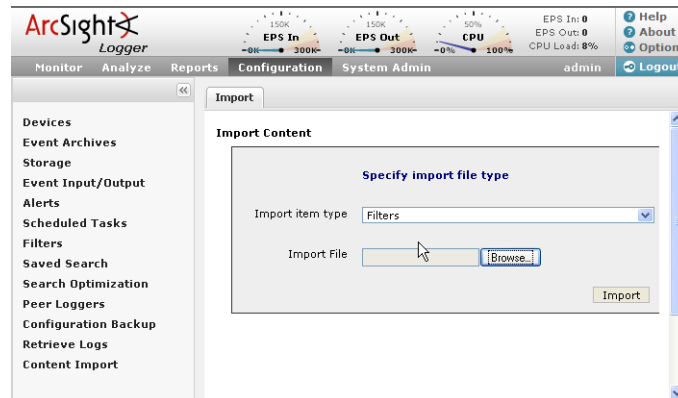
If the location you want is not in the drop-down list, you need to add it. For information about adding a network storage location, see ["Storage" on page 263](#).

- 5 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 needs to exist on the Mount Location. It is not created by the Logger.

- 6 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.
- 7 Click **Export**.

Importing Content



To import Alerts or Filters:

- 1 Click **Configuration > Content Import**.
- 2 Select the content type that you are importing in the “Import item type” field.

You can choose from Alerts or Filters.

- 3 Click **Browse** to locate the file.

The file needs to exist on a local or remote drive accessible to the system whose browser you are using to access Logger’s user interface.

- 4 Click **Import**.

Chapter 7

System Admin

This chapter describes the System Admin tab, which enables you to administer your Logger appliance and the software version of Logger. You create and manage Users in the System Admin tab, as well.

Not all System Admin settings are available on the software version of Logger, therefore this chapter is divided into two sections:

- [Section 1: Logger Appliance System Administration](#), for settings that apply to the Logger appliance.
- [Section 2: Software Version Logger Administration](#), for settings that apply to the software version of Logger.

Section 1: Logger Appliance System Administration

This section discusses the menu options available on a Logger appliance for system administration. On an appliance, you can configure network, storage, and security settings. In addition, the System Admin tab is where user accounts are managed.

This section contains the following topics:

["Reboot" on page 252](#)
["DNS Settings" on page 253](#)
["Hosts" on page 253](#)
["Network" on page 254](#)
["Time/NTP" on page 256](#)
["SMTP Settings" on page 258](#)
["Static Routes" on page 259](#)
["Static Routes" on page 259](#)
["License & Update" on page 259](#)
["Process Status" on page 260](#)
["Support Login" on page 261](#)
["Logs - Audit and Error" on page 262](#)
["Logs - Audit Forwarding" on page 262](#)
["Storage" on page 263](#)
["SAN" on page 267](#)
["Security" on page 270](#)
["Users/Groups" on page 279](#)

Reboot

There is no reason to reboot Logger during normal operations except for network configuration changes. If it becomes necessary to reboot the appliance, an administrator can perform this function using the browser UI.

To reboot Logger:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Reboot** from the System section.
- 3 Click **Start Reboot Now**.

Logger will reboot in about 60 seconds. The boot process normally takes 5-10 minutes, during which time the system is unavailable.



During reboot, Logger is not able to receive events. Events may be lost while the Logger reboots, unless SmartConnectors are used. SmartConnectors cache events when destinations like Logger are temporarily unavailable.

DNS Settings

ArcSight Platform Settings

DNS Hosts Network Time/NTP SMTP Static Routes

DNS Settings

Please enter DNS Servers

Primary IP Address
0.0.0.0

Secondary IP Address
0.0.0.0

Search Domains
localdomain

Update Settings

Figure 7-1 Domain Name Servers page

To change DNS settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.
- 3 In the **DNS** tab on the ArcSight Platform Settings page, enter new values for the IP address of the primary and secondary DNS servers, or edit the list of search domains.
- 4 Click **Update Settings** to make the changes, or click another tab or sub-menu to cancel. You must reboot Logger for the changes to become effective. See [“Reboot” on page 252](#).

Hosts

You can edit the Logger /etc/hosts file. The file will always contain an uneditable definition for localhost (127.0.0.1), used for static hostname mappings.

ArcSight Platform Settings

DNS **Hosts** Network Time/NTP SMTP Static Routes

Hosts Entries

System hosts file:

127.0.0.1 logger localhost.localdomain localhost

Update File

Figure 7-2 Hosts tab allows direct editing of etc/hosts file

To change Hosts file:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.
- 3 In the **Hosts** tab on the ArcSight Platform Settings page, edit the system's hosts file, adding one host per line. (The file will always contain a line for localhost.)
- 4 Click **Update File** to make the changes, or click another tab or sub-menu to cancel. Reboot the appliance for the changes to take effect. See ["Reboot" on page 252](#).

Network

Network settings, such as the Logger host name or the IP addresses for Logger's network interface cards (NICs), can be changed using the Network Settings page, shown in [Figure 7-3](#). Logger must be rebooted for the changes to take effect, however. (See ["Reboot" on page 252](#).)

ArcSight Platform Settings

DNS | Hosts | **Network** | Time/NTP | SMTP | Static Routes

Network Settings

Note: Settings take effect after reboot.

System Hostname
localhost

Default Gateway
192.168.35.1

☐ Automatically route outbound packets
(interface homing)

<p>NIC 'ETH0'</p> <p>IP Address 192.168.35.35</p> <p>Mask 255.255.255.0</p> <p>Speed/Duplex Auto (recommended) ▼</p>	<p>NIC 'ETH1'</p> <p>IP Address 192.168.36.35</p> <p>Mask 255.255.255.0</p> <p>Speed/Duplex Auto (recommended) ▼</p>
--	--

Update Settings

Figure 7-3 Network Settings page

To change network settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.

- 3 In the **Network** tab on the ArcSight Platform Settings page, enter new values for the following fields.

Parameter	Description
System Hostname	<p>The network host name for this Logger. A meaningful name will help, for example, when making a set of Loggers aware of each other.</p> <p>This name must be identical to the domain specified in the Certificate Signing Request, described in “Generating a Certificate Signing Request” on page 271.</p> <p>Note: If you use a CA-signed certificate on this Logger and you are changing its host name, you must generate a new CSR, obtain a new certificate for the Logger, and upload it to ensure that the connectors (in FIPS mode) that communicate with the Logger will be able to validate the host name. For more information about generating a CSR, see “Generating a Certificate Signing Request” on page 271.</p>
Default Gateway	The IP address of the default gateway.
Automatically route outbound packets (interface homing)	<p>When this feature is enabled (checked box), the response packets are sent back on the same Logger interface on which the request packets had arrived. Doing so 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 the Logger. If you have default gateway and static routes configured, they are ignored when this feature is enabled.</p> <p>When this feature is disabled (unchecked box), the default gateway and static routes (if configured) are used to determine the interface through which the response packets should leave the Logger.</p> <p>If you configure only one network interface, this setting does not provide any additional benefits.</p>
IP Address	The IP address for each Logger network interface card (NICs). These IP addresses should be on separate subnets to avoid confusion and to allow load balancing between receivers and forwarders.
Mask	Each Logger NIC has its own subnet mask, indicating which part of the IP address is local to its subnet.
Speed / Duplex	<p>Choose a speed and duplex mode, or let Logger automatically determine the network speed:</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 **Update Settings** to make the changes, or click another tab or sub-menu to cancel. The new settings will take effect after the next reboot.

**Note**

- Run the System Reboot command (see [“Reboot” on page 252](#)) to commit changes to network settings.
- It is important that the System hostname is resolvable by DNS and that it resolves to the Logger’s IP address. Performance is significantly affected if DNS cannot resolve the host name.

Time/NTP

The Time/NTP settings page enables you to configure system time, date, local timezone, and NTP servers. The times displayed for Logger operations such as searches, reports, and scheduled jobs are in the Logger’s local time zone.

**Tip**

Because precise time stamping of events is critical for accurate and reliable log management, ArcSight strongly recommends using an NTP server.

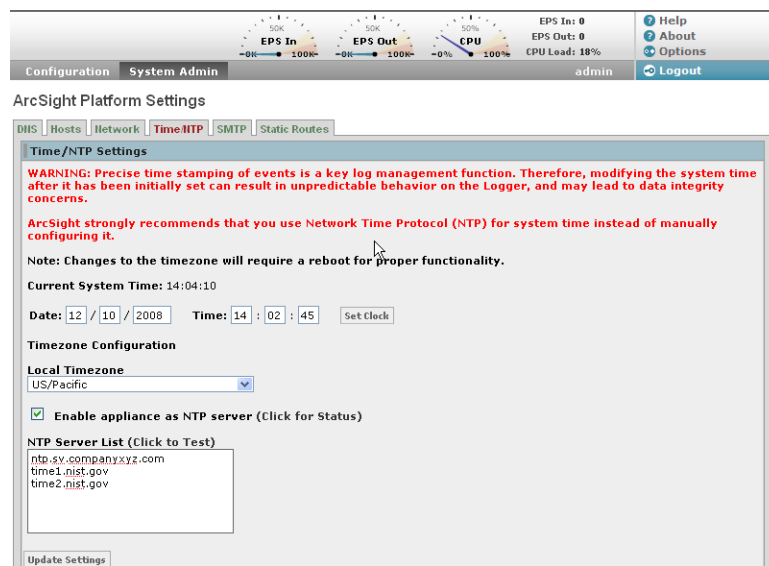


Figure 7-4 Time Settings page

To change the current Logger time:

**Caution**

Modifying the system time after it has been initially set can result in unpredictable behavior on the Logger, thus compromising data integrity.

ArcSight strongly recommends that you use Network Time Protocol (NTP) for system time instead of manually configuring it. However, if you need to change the system time manually, please contact ArcSight Customer Support for guidance.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.

- 3 In the **Time/NTP** tab on the ArcSight Platform Settings page, enter new values for hour, minute, second, month, day, or year.
- 4 Click **Set Clock** to set the Logger clock to the new values.

To change time configuration:

ArcSight Platform Settings

Time/NTP Settings

WARNING: Precise time stamping of events is a key log management function. Therefore, modifying the system time after it has been initially set can result in unpredictable behavior on the Logger, and may lead to data integrity concerns.

ArcSight strongly recommends that you use Network Time Protocol (NTP) for system time instead of manually configuring it.

Note: Changes to the timezone will require a reboot for proper functionality.

Current System Time: 14:04:10

Date: 12 / 10 / 2008 Time: 14 : 02 : 45 **Set Clock**

Timezone Configuration

Local Timezone: US/Pacific

☒ Enable appliance as NTP server (Click for Status)

NTP Server List (Click to Test)

ntp.sv.companyxyz.com
time1.nist.gov
time2.nist.gov

Update Settings

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.
- 3 In the **Time/NTP** tab on the ArcSight Platform Settings page, enter new values for the following fields.

Parameter	Description
Local timezone	Choose GMT or an appropriate timezone.
Enable appliance as NTP Server	Check this setting if this Logger appliance should be used as an NTP server.
NTP Server List	<p>Enter the host name of an NTP server. For example, time.nist.gov.</p> <p>ArcSight recommends using at least three NTP servers to ensure precise system time on Logger. 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 you added are reachable from this Logger appliance.</p> <p>Notes:</p> <ul style="list-style-type: none"> A Logger can serve as an NTP server for any Logger; not only its peers. If Logger A serves as an NTP server for Logger B, Logger B needs to list Logger A in its NTP Server List.

- 4 Click **Update Settings** to make the changes, or click another tab or sub-menu to cancel. You must reboot Logger for the changes to become effective. See ["Reboot" on page 252](#).

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 8, 2009) 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 1, 2009) are run at standard time instead of the DST time.

Examples:

- A report scheduled to run at 1 a.m. DST on November 1, 2009 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 1, 2009 will run at 2 a.m.; however, due to time adjustment, an hour later than it ran on the previous day (October 31, 2009).
- A report scheduled to run at 2 a.m. on March 8, 2009 will not run.

SMTP Settings

Alerts use Simple Mail Transfer Protocol (SMTP) to send e-mail.

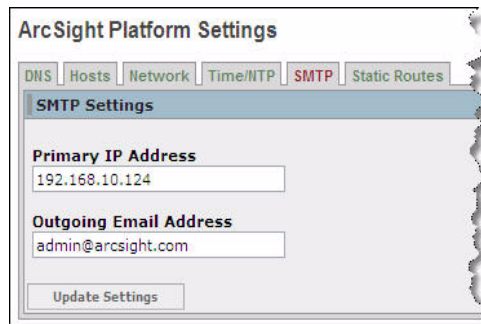


Figure 7-5 Simple Mail Transfer Protocol (SMTP) settings

To change SMTP configuration:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.
- 3 In the **SMTP** tab on the ArcSight Platform Settings page, enter new values for the following fields.

Parameter	Description
Primary SMTP Address	Enter the IP address of the SMTP server that will process outgoing e-mail.
Outgoing Email Address	The e-mail address that will appear in the From: field of outbound e-mail.

- 4 Click **Update Settings** to make the changes, or click another tab or sub-menu to cancel. Changes take effect immediately; reboot is not required.

Static Routes

Advanced users can specify static routes for either or both network adapters. The Static Routes page displays a table of all specified static routes.

ArcSight Platform Settings

DNS | Hosts | Network | Time/NTP | SMTP | **Static Routes**

Static Routes Settings

Add Static Route

Network Adapter:

Dest. Type:

Destination:

Subnet Mask:

Gateway:

Current Static Routes

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface	Last Update	Action
<input type="button" value="Add Static Route"/>									

Figure 7-6 Static Routes page

To add a static route:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.
- 3 In the **Static Routes** tab on the ArcSight Platform Settings page, click **Add Static Route**.
- 4 Enter new values for the following fields.

Parameter	Description
Network Adapter	Choose the network interface card (NIC).
Destination Type	Select Network or Host.
Destination	Specify the IP address for the static route destination.
Subnet Mask	Enter the subnet mask (for example, 255.255.255.0) for network only.
Gateway	Specify the IP address for the default gateway.

- 5 Click **Create Static Route** to add the new static route to the table, or click another tab or sub-menu to cancel. You must reboot Logger for the changes to become effective. See ["Reboot" on page 252](#).

License & Update

Updating system software requires uploading an upgrade file provided by ArcSight Customer Support using the System Update page. The System Update page also displays the elapsed time since the appliance was last rebooted, and the version of the Logger components. The Logger version and build number is found at 'arcsight-logger'.

To upload an upgrade file:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **License & Update** from the System section.
- 3 Click **Browse** to locate the file.
- 4 Click **Upload Update**.



Note

System Update will take effect after the next reboot. To update immediately, reboot the system after performing a System Update. See [“Reboot” on page 252](#).

Process Status

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Process Status** from the System section to display a page similar to the one shown in [Figure 7-7](#).

Process Status

System	Status	Load			CPU	Memory
logger	running	[0.30]	[0.28]	[0.25]	1.4%us 0.2%sy 0.1%wa	1909144 kB [7.7%]

Process	Status	Uptime	CPU	Memory	Memory (kB)
apache	running	23h 51m	0.0%	0.0%	7300
aps	running	23h 51m	0.0%	0.7%	191736
connector	running	23h 51m	0.0%	0.0%	568
insp	running	23h 49m	0.0%	0.1%	27604
mysqld	running	23h 51m	0.0%	0.0%	21460
nullmailer	running	23h 49m	0.0%	0.0%	816
postgresql	running	23h 51m	0.0%	0.0%	9244
processors	running	23h 49m	0.0%	0.6%	149160
receivers	running	23h 49m	0.0%	0.2%	50700
reportengine	running	23h 49m	0.0%	0.3%	94856
servers	running	23h 51m	0.0%	2.3%	587224
web	running	23h 49m	1.3%	1.6%	406388

Figure 7-7 Process Status page

(In the process list, processors refers to Forwarders.)

Additional system information, specifically the system uptime and component versions, is available on the System Update page. (See [“License & Update” on page 259](#).)

Each process is a hyperlink. Clicking on an individual process displays more detail about that process, as shown in [Figure 7-7](#).

Status detail for apache

Parameter	Value
children	15
cpu_percent	0.0%
cpu_percent_total	0.0%
data_collected	Wed Sep 2 13:58:02 2009
memory_kilobytes	7300
memory_kilobytes_total	250200
memory_percent	0.0%
memory_percent_total	1.0%
monitoring_status	monitored
parent_pid	1
pid	4279
status	running
uptime	23h 53m

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

[BACK](#) [RESTART](#)

Figure 7-8 Process Status detail for apache

Support Login

When Customer Support needs access to your appliance for troubleshooting and diagnostics, they work with you to assign a single-use password to the appliance. Doing so enables Support Login access to the appliance. This password is valid only for one support session and is automatically disabled after the session ends. (You can also explicitly disable Support Login access.)

ArcSight
Logger

Monitor Analyze Configuration **System Admin** admin Logout

System
Reboot
Network
License & Update
Process Status
Support Login

Logs
Audit
Error
Forwarding

Storage
CIFS
NFS
RAID Controller

Security
SSL Server Certificate
SSL Client Authentication
FIPS 140-2

Users/Groups
Authentication
Groups
Users
Change Password

Enable/Disable ArcSight Support Login

Set up a password for this appliance that ArcSight Customer Support can use to access this appliance for problem diagnosis. The password is automatically disabled after one support session. Or you can explicitly disable it.

ArcSight Support Login

Access to this appliance is currently disabled.

Request Code: PKA88

Activation Code:

Root Password:

Confirm Root Password:

[Enable Support Login](#)

When you report an issue to ArcSight Customer Support that requires them to access your appliance, they will direct you on enabling Support Login access.

The only circumstance in which you will need to explicitly disable Support Login access is if access was enabled but the support session never occurred.

To **disable** support login access:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Support Login** from the System section.
- 3 Click **Disable Support Login** in the right-side panel.

Logs - Audit and Error

Logger audit and error logs are available for viewing.

The screenshot shows the ArcSight Logger System Admin interface. The top navigation bar includes Monitor, Analyze, Reports, Configuration, and System Admin. The System Admin section is active, showing a sidebar with System, Logs, Storage, Security, and Users/Groups. The main content area is titled 'Search Audit Logs' and contains the following sections:

- Select Audit Type:** A dropdown menu labeled 'Select an audit type:'.
- Select Date Range:** Fields for Start Date (Jun 11, 2010) and End Date (Jun 11, 2010).
- Select User (optional):** A table with columns Login, First Name, Last Name, Email, and Phone Groups. The table contains one entry: adminDefault, Admin, admin@arcsight.com, and a list of default groups.

A 'View Audit Logs' button is located at the bottom of the table.

To view Audit or Error logs:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Audit** (for audit logs) or **Error** (for Error logs) from the Logs section.
- 3 Select the type of log—Application or Platform.
- 4 Select the date range for which you want to obtain the log.
- 5 Click **View Error Logs**.



To search again after clicking **View Audit Logs**, use the browser's Back button.

Audit logs, as Common Event Format (CEF) audit events, can be sent to ArcSight ESM directly for analysis and correlation because the Logger Forwarder supports ESM Manager's event protocol.

For more information about audit event forwarding, see [“Logs - Audit Forwarding” on page 262](#). For information about audit events that you can forward, see [Appendix E, Logger Audit Events, on page 401](#).

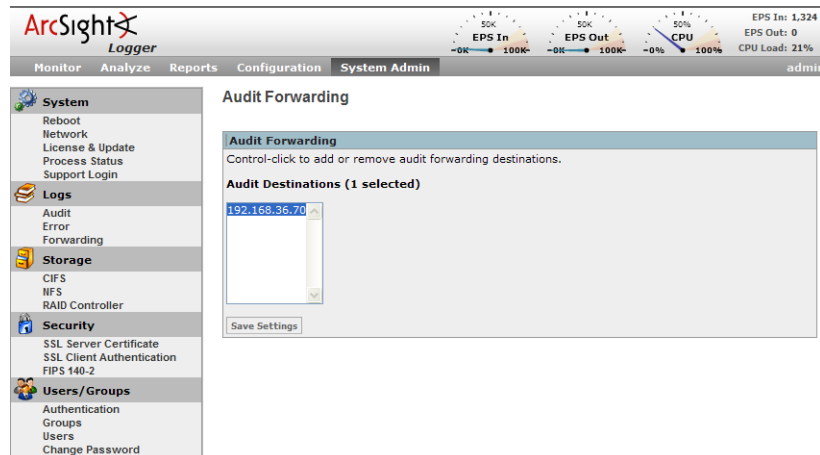
Logs - Audit Forwarding

For information about audit events that you can forward, see [Appendix E, Logger Audit Events, on page 401](#).

To forward audit events to specific destinations:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Forwarding** from the Logs section.
- 3 Select destinations from the Audit Destinations list, as shown in the following figure. Click on a destination to select a single destination, or Ctrl+click to select or de-select

multiple destinations. The destinations are ESM destinations that you configure on the ESM Destinations page (**Configuration > Event Input/Output > ESM Destinations**).



- 4 Click **Save Settings**.

Storage

Logger can mount NFS and CIFS 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. Loggers with Storage Area Network (SAN) capability can also interface with a SAN.

The Storage tab includes the ability to configure NFS and CIFS mounts for archiving data and configure LUNs (on systems that support SAN).

In addition, this tab provides status of the hard disk array (RAID) controller and specific system processes.

CIFS Settings

Logger can mount a CIFS remote file system (Windows share) to archive data such as events, exported filters and alerts, and Saved Searches. A CIFS file system cannot be used as the primary storage device for Logger.

Before you mount a Windows share to a Logger, make sure

- A user account with read-write privileges to the share exists on the Windows system.
- The folder to which you are establishing the mount point is configured for sharing.

ArcSight Logger

Monitor Analyze Reports Configuration **System Admin** admin Logout

System Admin

- System
 - Reboot
 - Network
 - License & Update
 - Process Status
 - Support Login
- Logs
 - Audit
 - Error
 - Forwarding
- Storage
 - CIFS
 - NFS
 - RAID Controller
- Security
 - SSL Server Certificate
 - FIPS 140-2
- Users/Groups
 - Authentication
 - Groups
 - Users
 - Change Password

CIFS Mount Administration

Add Remote Mount Point

Name: Hurricane

File System Mount Options: rw

Remote Hostname/IP Address: 192.0.2.11

Username: admin

Password:

Share Name: CIFS1

Description: CIFS archival for Logger appliance.

Save CIFS Mount Cancel

To add a CIFS mount:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **CIFS** under the Storage section in the left panel.
- 3 Click **Add CIFS Mount** in the right panel.
- 4 Enter values for the following fields.

Parameter	Description
Name	A meaningful name for the Windows share. The name cannot contain spaces. This name is used locally on your Logger to refer to the mount point and needs to be specified when configuring archive settings for data that will be stored on the share.
File System Mount Options	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. 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.
Remote Hostname / IP Address	Host name or IP address of the host to which you are creating the CIFS mount.
Username	Name of the user account with read-write privileges to the Windows share. Make sure the username is prefixed with the domain information. For example, <code>tahoe/arcsight</code> .
Password	Password for the user name specified above.

Parameter	Description
Share Name	<p>The folder on the Windows host to which you are creating the CIFS mount. For example, <code>logger_logs</code>.</p> <p>This folder needs to be configured for sharing. (Typically, to configure a Windows folder for sharing, right click on the folder name > Properties > Sharing.)</p> <p>Note: If you cannot mount successfully, try specifying a leading slash (\) in the remote path. For example, <code>\connector_logs</code>.</p>
Description	A meaningful description of the mount point.

- 5 Click **Save CIFS Mount**.
- 6 (Optional) Click **test** in the Action column of the mount point you added to test connectivity to the Windows share.

To edit a CIFS mount:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **CIFS** under the Storage section in the left panel.
- 3 Click **edit** in the Action column for the CIFS mount that you want to edit. Change field values as needed.
- 4 Click **Save CIFS Mount**.

To delete a CIFS mount:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **CIFS** under the Storage section in the left panel.
- 3 Click **delete** in the Action column for the CIFS mount that you want to delete.
- 4 Confirm the deletion.



Caution

Deleting the CIFS mount (or detaching the SAN) used for Event Archive (see [“Archive Storage Settings” on page 186](#)) will permanently disable the Event Archive feature.

Network File System (NFS) Settings

An NFS mounted system can be used to archive data such as events, exported filters and alerts, and Saved Searches. Use of a Network File System (NFS) as primary storage for Logger events is not recommended.

Before you mount an NFS share of a remote system, make sure you grant Logger read and write permission on that system. The account name is 'arcsight', but use numeric ids instead: 1500 for uid, or 750 for gid.

Logger supports only NFS v3.0.



Tip

ArcSight recommends creating a Configuration Backup whenever NFS settings are changed. A current backup is useful for disaster recovery. For more information, see [“Configuration Backup and Restore” on page 235](#).

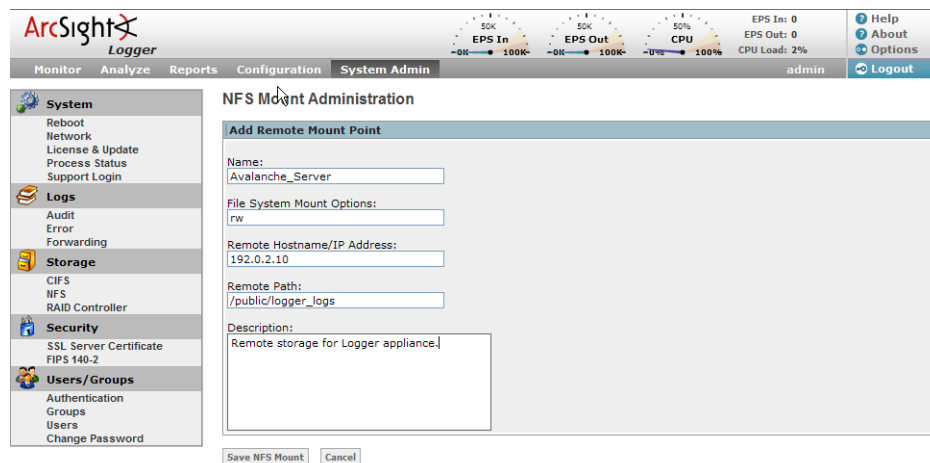


Figure 7-9 NFS Mount Administration page

To add an NFS mount:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **NFS** under the Storage section in the left panel.
- 3 Click **Add NFS Mount** in the right panel.
- 4 Enter new values for the following fields:

Parameter	Description
Name	A name for the network file system mount. The name cannot contain spaces.
File System Mount Options	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. 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.
Remote Hostname / IP Address	Host name or IP address of the host to which you are creating the NFS mount.
Remote Path	The folder on the remote host that will act as the root of the network file system mount. For example, <code>/public/logger_logs</code> . Make sure that only this Logger can write to the location you specify in this field. If multiple Loggers (or other systems) mount this location and write to it, data on this location will be corrupted.
Description	A meaningful description of the mount point.

- 5 Click **Save NFS Mount**.
- 6 (Optional) Click **test** in the Action column of the mount point you added to test the network file system connectivity.

To edit an NFS mount:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **NFS** under the Storage section in the left panel.
- 3 Click **edit** in the Action column for the NFS mount that you want to edit. Change field values as needed.
- 4 Click **Save NFS Mount** to make the changes, or click **Cancel** to quit.

To delete an NFS mount:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **NFS** under the Storage section in the left panel.
- 3 Click **delete** in the Action column for the CIFS mount that you want to delete.
- 4 Confirm the deletion.



Deleting the NFS mount (or detaching the SAN) used for Event Archive (see [“Archive Storage Settings”](#) on page 186) will permanently disable the Event Archive feature.

SAN

Some models of Logger appliance include the ability to connect to a Storage Area Network (SAN) for various purposes. SANs contain Logical Units (LUNs), identified by their World Wide Name. As shown in [Figure 7-10](#), a LUN's Attachment Status can be 'available,' 'attached,' or 'detached.' LUNs in a SAN are in one state at a time. Actions such as “attach” change from one state to another.

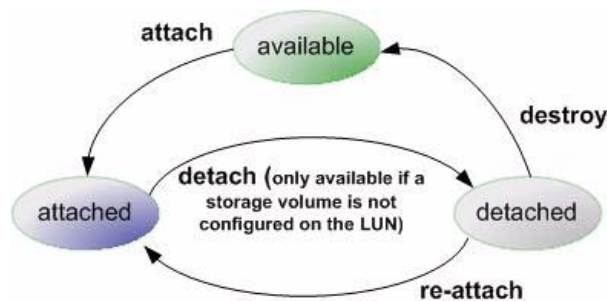


Figure 7-10 SAN Logical Unit state diagram.



Destroying a Logical Unit that has been detached puts that LUN into a state in which a subsequent attach will erase any data stored on the Logical Unit. If a LUN is accidentally destroyed, ArcSight Customer Support may be able to recover the data, provided the LUN is not attached.

The following table summarizes the states and possible actions:

Attachment Status	Actions	Description
available	attach	Logical Units detected on a SAN are initially available for attachment.
attached	detach	Attached Logical Units 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 D, Restoring Factory Settings, on page 393 .
detached	re-attach destroy	When an attached Logical Unit 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 wipes out the data and releases the Logical Unit back to the available state. Note: When you detach, the only action available immediately is re-attach . The destroy state takes a few minutes to display because it takes a few minutes for the LUN to detach on the system.



Note

Changes to the SAN (adding or removing LUNs, for example) will not be reflected in Logger until Logger is rebooted.

To attach a LUN:



Note

- Logger can attach to only one LUN (on SAN) at a time for primary storage. You can add more LUNs for event archival, configuration backup, and export.
- Although the HBA card on the Logger contains two physical interfaces, only a single interface can be enabled. Therefore multi-path support is not available currently.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN Storage** under the Storage section in the left panel.
- 3 Find the LUN in the SAN Logical Unit List.
- 4 In the Action column, click **attach** for that row.
- 5 The LUN's Attachment Status will change to 'attached' when the LUN is ready for use.

To detach a LUN:



Note

LUNs that are used for primary storage may not be detached.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN Storage** under the Storage section in the left panel.

- 3 In the SAN Logical Unit List, locate the LUN to be detached. In the action column, click **detach** for that row. Change field values as needed.
- 4 The LUN's Attachment Status will change to 'detached.'

To re-attach a LUN:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN Storage** under the Storage section in the left panel.
- 3 In the SAN Logical Unit List, locate the LUN to be reattached. The LUN must be in the 'detached' state. In the action column, click **re-attach** for that row.
- 4 The LUN's Attachment Status will change to 'attached.'

To release a LUN:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN Storage** under the Storage section in the left panel.
- 3 In the SAN Logical Unit List, locate the LUN to be released. The LUN must be in the 'detached' state. In the action column, click **destroy** for that row.
- 4 The LUN's Attachment Status will change to "available".

SAN Storage Administration

SAN Logical Unit List						
Name	Local Device Name	World Wide Name	Size	Attachment Status	Action	
SAN1	/dev/sdb	5006016830224f88:0000000000000000	49.95 GB	attached	detach	
san1	/dev/sdc	5006016830224f88:0001000000000000	49.95 GB	available	destroy re-attach	

Figure 7-11 SAN Storage Administration page

Restoring a SAN

To restore a SAN to either the Logger to which it was formerly attached or a new Logger (in the case of disaster recovery), follow these steps:

- 1 With Logger powered off, attach the SAN physically. Turn on Logger.
- 2 Restore the configuration to Logger. ArcSight 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 SAN has been restored.
- 3 Enable one-time Support Login (see ["Support Login" on page 261](#)). Contact ArcSight Customer Support.
- 4 ArcSight Customer Support will login remotely, stop all Logger processes by issuing the command

```
/opt/local/monit/bin/monit stop all
```

and migrate the internal database to the SAN by creating a symbolic link with the command

```
ln -s <remote storage path> /opt/local/pgsqldata
```

When Customer Support has finished these tasks, reboot Logger. This will disable future Support Logins.

RAID Controller

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **RAID Controller** under the Storage section in the left panel to display a page similar to the one shown in [Figure 7-12](#).



Note

Logger hardware models use different RAID controllers, which display information differently.

Status of RAID Controller

General Controller Information	
Type:	RAID-5
State:	Optimal
Versions:	
Product Name	: PERC 6/i, Integrated
Serial No	: 1122334455667788
FW Package Build:	6.1.1-0047
Image Versions In Flash:	
FW Version	: 1.21.02-0528
BIOS Version	: 2.01.00
WebBIOS Version	: 1.1-46-e_15-Rel
Ctrl-R Version	: 1.02-014B
Boot Block Version	: 1.00.00.01-0011
HW Configuration:	
SAS Address	: 50024e805edb8600
BBU	: Present
Alarm	: Absent
NVRAM	: Present
Serial Debugger	: Present
Memory	: Present
Flash	: Present
Memory Size	: 256MB
Device Present:	
Virtual Drives	: 2
Degraded	: 0
Offline	: 0
Physical Devices	: 7
Disks	: 6
Critical Disks	: 0
Failed Disks	: 0
Error Counters:	
Memory Correctable Errors	: 0
Memory Uncorrectable Errors	: 0
Drive states:	
0:	Online
1:	Online
2:	Online

Figure 7-12 RAID Controller Information page

Obviously, this information is highly technical. It is not needed during normal Logger operations, but it can be helpful for diagnosing specific hardware issues. Due to the redundant nature of RAID storage, unit failure does not disable Logger. Instead, performance degrades. Use this report to determine whether a performance issue is caused by a disk failure. ArcSight Customer Support can use this information to better diagnose problems, as well.

Security

Security settings enable you to configure SSL Server certificates, enable and disable FIPS (Federal Information Processing Standards) mode on the Logger, and configure SSL client authentication for CAC support.

SSL Server Certificate

Logger uses Secure Sockets Layer (SSL) technology to communicate securely over an encrypted channel with its clients—users, SmartConnectors when using the SmartMessaging technology, and peer Loggers. To establish a typical SSL session, an SSL certificate is required on the server (Logger) side and a truststore is required on the client side. The truststore contains a list of Certificate Authorities (CA) that the client trusts.

When a client initiates communication with Logger, the Logger sends its SSL certificate to the client to authenticate itself. The client checks its truststore to validate the certificate. (In addition, the client verifies whether the hostname in the certificate matches the one with which it initiated communication, and the current time on the client machine is within the validity range specified in the certificate.) If the certificate is validated, a session key is exchanged between the client and the Logger. This key is used to encrypt and decrypt data exchanged between the Logger and the client.

Logger ships with a self-signed certificate. Although you can use this certificate, ArcSight strongly recommends using a CA-signed certificate.

Even if FIPS is not enabled on a Logger, 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 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*.

To facilitate obtaining a CA-signed certificate, Logger can generate a Certificate Signing Request. Once a signed certificate file is available from the CA, it can be uploaded to Logger for use in subsequent authentication.

Generating a Certificate Signing Request

The first step in configuring an SSL certificate is to generate a Certificate Signing Request (CSR). The CSR must be generated on the Logger appliance for which you are requesting a certificate. That is, you cannot generate a CSR for Logger A on Logger B or use a third-party utility to generate it.

The resulting CSR should be sent to a CA, such as VeriSign, which responds with a signed certificate file.

The screenshot shows the ArcSight Logger web interface. At the top, there's a status bar with system metrics: EPS In (1,324), EPS Out (0), CPU (72%), and CPU Load (72%). Below this is a navigation menu with tabs: Monitor, Analyze, Reports, Configuration, System Admin, and a user dropdown (admin). The 'System Admin' tab is selected, and the 'Security' sub-tab is active. The main content area is titled 'ArcSight SSL Settings' and contains a 'Generate CSR' tab. The 'Generate CSR' tab is active, showing a form titled 'Generate Certificate Signing Request'. The form has a section 'Please enter the Certificate Settings' with the following fields: Country (2-letter code) (US), State/Province (California), City/Locality (Cupertino), Organization Name (ArcSight, Inc.), Organizational Unit (Support Team), Hostname (loggerA.arcsight.com), Email Address (support@arcsight.com), and Private Key Password (empty). A 'Generate CSR' button is at the bottom of the form. The left sidebar shows a navigation menu with categories: System (Reboot, Network, License & Update, Process Status, Support Login), Logs (Audit, Error, Forwarding), Storage (CIFS, NFS, RAID Controller), Security (SSL Server Certificate, SSL Client Authentication, FIPS 140-2), and Users/Groups (Authentication, Groups, Users, Change Password).

Figure 7-13 Certificate Signing Request page

To generate a certificate signing request:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Server Certificate** under the Security section in the left panel to display the Generate Certificate Signing Request page, as shown in [Figure 7-13 on page 271](#).
- 3 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 'Cupertino.'
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 Logger. When specifying the host name, make sure that this name matches the name registered in the Domain Name Service (DNS) server for the Logger. Additionally, this name must be identical to the host name specified in "Network" on page 254 . Note: If the host name or IP address of this Logger appliance changes in future, you must generate a new CSR, obtain a new certificate for the Logger, and upload it to ensure that the connectors (in FIPS mode) that communicate with the Logger will be able to validate the host name.
Email Address	The e-mail address of the administrator or contact person with regard to this CSR.
Private key password	The password to secure the private key on the appliance. This password is not included in the generated CSR. It is stored locally on your Logger.

- 4 Click **Generate CSR** to generate a Certificate Signing Request for download, or click another tab or sub-menu to cancel.

Installing a Signed Certificate

ArcSight SSL Settings

Generate CSR | **Install Cert** | View Results

Upload Signed Certificate

Note: After uploading the new certificate, close and re-open the browser.

Please select the signed Certificate file

Figure 7-14 Install Certificate page

To install a signed certificate:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Server Certificate** under the Security section in the left panel.
- 3 On the **Install Cert** tab (as shown in [Figure 7-14 on page 272](#)), click **Browse** to find the signed certificate file on your local file system.
- 4 Click **Upload and Install** to install the specified certificate, or click another tab or sub-menu to cancel.

Certain browsers require that you close your current browser and restart it for the new certificate to take affect. If you are aware of this requirement for your browser or are unsure of it, restart your browser.

View Results of Certificate Installation

The **View Results** tab displays the results of the most recent certificate installation.

SSL Client Authentication (CAC Authentication)

Logger 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 (user name and password) and RADIUS authentication. As a result, Logger 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.

Configuring Logger to Support SSL Client Authentication (CAC)

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 [“SSL Server Certificate” on page 271](#) to load the certificate.
- 2 Enable client certificate authentication, as described in [“Client Certificate Authentication” on page 282](#).



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.

- 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 274](#).

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 Logger user name for each user who will be connecting to the Logger using a client certificate, as described in [“Users” on page 289](#).
- 5 (Optional) Upload a certificate revocation list (CRL), as described in [“Uploading a Certificate Revocation List” on page 274](#).

- 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 peer with it. For more information, see ["Peer Loggers" on page 231](#).

On the Client (Web browser)

Configure your browser to provide the SSL client certificate when accessing Logger. That is, upload the private key in PKCS 12 format in your web browser.

Uploading Trusted Certificates

A trusted certificate is used to authenticate users that log in to the Logger. 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** from 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.

Viewing Details of a Trusted Certificate

To view details of a trusted certificate:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Client Authentication** from the Security section in the left panel.
- 3 On the Trusted Certificates tab, click the certificate whose details you want to view in the "Certificates in Repository" list.

Deleting a Trusted Certificate

To delete a trusted certificate:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Client Authentication** from the Security section in the left panel.
- 3 On the Trusted Certificates tab, select the certificate from the "Certificates in Repository" list and click the **Delete** button.

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 the Logger. 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** from 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.

Viewing Details of a CRL file

To view details of a CRL file:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Client Authentication** from the Security section in the left panel.
- 3 In the **Certificate Revocation List** tab, click the link displayed in the Issuer Name column.

Deleting a CRL File

To delete a CRL file:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Client Authentication** from the Security section in the left panel.
- 3 In the **Certificate Revocation List** tab, select it and click the **Delete** button.

FIPS 140-2

Logger 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 Logger needs to be FIPS 140-2 compliant, you can enable FIPS on it. Once you do so, the Logger 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 277](#) to ensure that your connector is FIPS compliant.
- Logger forwarders, such as ESM 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 ESM Manager, it needs to run ESM v4.0 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 203](#) to complete configuration of this setup.

- Peer Loggers

Loggers running v4.0 automatically use FIPS 140-2 compliant algorithms. Therefore, no action is required on a peer Logger running version 4.0. A FIPS-enabled version 4.0 Logger can communicate with a non-FIPS enabled Logger running v4.0. Additionally, a Logger running v3.0 SP1 Patch1 can be peered with a Logger running v4.0.

- 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.

Before you enable FIPS mode on your Logger, make sure:

- Your Logger is set up with a CA-signed SSL certificate. For more information, see ["SSL Server Certificate" on page 271](#).
- 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, load it on the container by following instructions in the "Managing Certificates on a Container" section in the *Connector Appliance Administrator's Guide*.
- 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 and later.

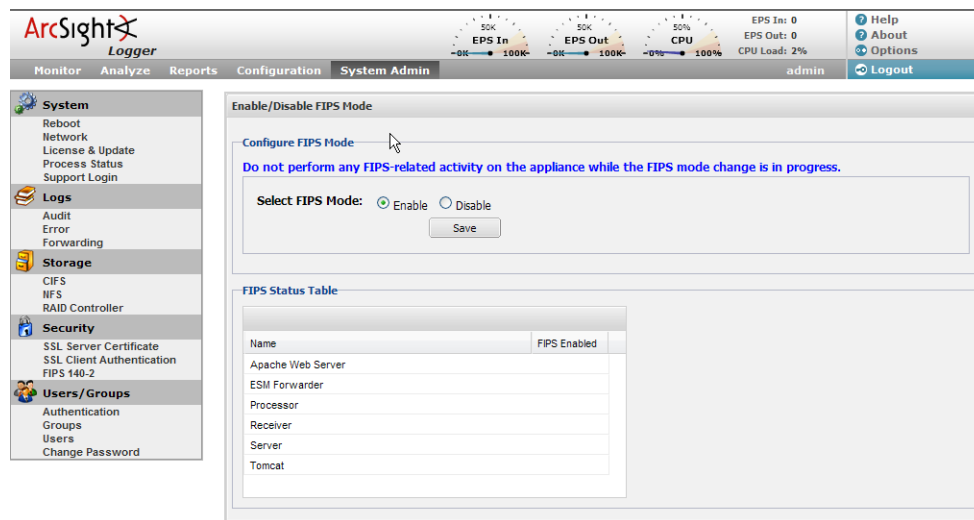
To enable or disable FIPS mode on Logger:



Make sure you are familiar with the configuration requirements on your Logger as described in ["Before you enable FIPS mode on your Logger, make sure:" on page 276](#).

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **FIPS 140-2** from the Security section in the left panel.
- 3 Click **Enable** or **Disable** for the Select FIPS Mode option.
- 4 Click the **Save** button.
- 5 If the System Reboot Required message displays, click the **System Reboot** link.

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 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 (version 4.7.5.5372 or later) from the ArcSight Customer Support site. 2 Go to Step 1 on page 277.
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 (version 4.7.5.5372 or later). Follow instructions in the <i>SmartConnector User's Guide</i> to upgrade the SmartConnector. 2 Only perform Step 2a on page 278.
Updating a SmartConnector to be FIPS-compliant and the SmartConnector is running version 4.7.5.5372 or later	Only perform Step 2a on page 278 .

- 1 Follow device configuration steps provided in the SmartConnector's configuration guide (available from the ArcSight Customer Support site at <https://support.arcsight.com>), then follow the installation procedure through installation of the core connector software (SmartConnector Installation step 2).

At Step 3 of the Connector setup, as shown below, click **Cancel** to exit the setup to configure the NSS DB, which is necessary for installing the connector in FIPS-compliant mode.

Step 3: When the installation of ArcSight SmartConnector core component software is finished, the following window is displayed:



- 2 Click **Cancel** to exit the configuration wizard. You will return to this wizard and resume SmartConnector configuration, after

- ◆ Enabling FIPS mode on it, and
- ◆ Importing Logger's certificate

Enable FIPS Mode on the SmartConnector

- a Create an `agent.properties` file at the following location:

```
$ARCSIGHT_HOME\current\user\agent
```

- b Enter the following property, then save and close the file.

```
fips.enabled=true
```

Import Logger's Certificate on the SmartConnector

- c 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
```

- d 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 3.0.x, click **Tools > Options > Encryption > View Certificates > Servers > Select your Logger appliance > Export**. Save the certificate file with a `.cert` or `.cer` extension.
- ii Copy the certificate file you exported in the previous step (in this example, `loggercert.cert`) to the `$ARCSIGHT_HOME\current\bin` directory.

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.cert
```

- e Enter the following command to re-enable FIPS mode that you turned off in Step 1:

```
arcsight runmodutil -fips true -dbdir0
user/agent/nssdb.client
```

- f 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.
- g *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**.

If you are installing a new SmartConnector, go to the next step.

- 3 To return to the SmartConnector configuration wizard, enter the following from `$ARCSIGHT_HOME\current\bin`:

```
arcsight connectorsetup
```

- 4 When prompted whether you want to start in Wizard Mode, click **Yes**.
- 5 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

Authentication Settings

The Authentication settings enable you to specify settings and policies for login, password, and the authentication mechanism to use.

Login

The Authentication Settings page lets you specify the maximum number of simultaneous sessions for a single user account, which may impact system performance.

The form, shown in [Figure 7-15](#), also lets you specify how many seconds of inactivity to allow before automatically ending the current session. The default is 900 (15 minutes).

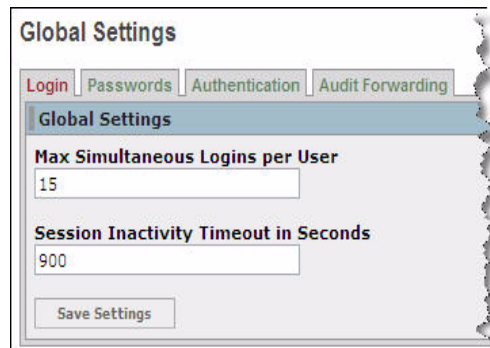
The image is a screenshot of a web-based configuration interface. At the top, there's a title bar that says "Global Settings". Below this, there are four tabs: "Login", "Passwords", "Authentication", and "Audit Forwarding". The "Login" tab is currently selected. Under the "Login" tab, there's a sub-header "Global Settings". Below this, there are two input fields. The first is labeled "Max Simultaneous Logins per User" and has the value "15" entered. The second is labeled "Session Inactivity Timeout in Seconds" and has the value "900" entered. At the bottom of the form, there is a button labeled "Save Settings".

Figure 7-15 Login Settings page, changing the Session Inactivity time-out to 3 minutes.

To change login settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** under the Users/Groups section in the left panel to display the Global Settings page in the Login tab, as shown in [Figure 7-15 on page 280](#).
- 3 Enter new values for the maximum simultaneous logins per user or the session inactivity time-out.
- 4 Click **Save Settings** to make the changes, or click another tab or sub-menu to cancel. You must reboot Logger for the changes to become effective. See ["Reboot" on page 252](#).

Password

Password policies include the minimum and maximum number of characters and other requirements for passwords. The Logger administrator can specify that an account should be locked out after an authentication failure under certain circumstances.

Authentication Settings

Login | **Passwords** | Authentication

Password Settings

Enable Password Lockout ☒ Yes ☐ No

3 Number of failed attempts before lockout

60 Maximum time between attempts (in seconds)

15 Lockout duration (in minutes)

Enable Password Expiration ☒ Yes ☐ No

90 Days until password expires

5 Days before expiration to notify user

Enable Password Validation ☒ Yes ☐ No

Password Length Limits

10 Minimum password length

20 Maximum password length

Minimum Requirements

2 Numeric characters [0-9]

0 Uppercase characters [A-Z]

0 Lowercase characters [a-z]

2 Non-alphanumeric characters [!\$^*...]

2 Number of characters different from old password

Save Settings

Figure 7-16 Password Policy Settings page

To change password policy settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** under the Users/Groups section in the left panel to display the Password Settings page in the Password tab, as shown in [Figure 7-16 on page 281](#).
- 3 Update any or all of the parameters listed in the following table:

Parameter	Description
Enable Password Lockout	Choose Yes to enforce the password policy. The default is No .
Number of failed attempts before lockout	Default is 3 .
Maximum time between attempts (in seconds)	Default is 60 , or one minute.
Lockout duration (in minutes)	Default is 15 .
Enable Password Expiration	Choose Yes to expire passwords automatically. The default is No .
Days until password expires	The default is 90 .

Parameter	Description
Days before expiration to notify user	The default is 5 .
Enable Password Validation	Choose Yes to enforce the length limits and other requirements for new passwords. The default is No .
Minimum password length	Enter the minimum number of characters in a password. The default is 10 .
Maximum password length	Enter the maximum number of characters in a password. The default is 20 .
Numeric characters	Enter the minimum number of numeric characters (0-9) in a valid password. The default is 2 .
Uppercase characters	Enter the minimum number of uppercase characters (A-Z) in a valid password. The default is 0 .
Lowercase characters	Enter the minimum number of lowercase characters (a-z in a valid password. The default is 0 .
Non-alphanumeric characters	Enter the minimum number of characters that are not digits or letters that are required in a valid password. The default is 2 .
Number of characters different from old password	The default is 2 .

- 4 Click **Save Settings** to make the changes, or click another tab or sub-menu to cancel. You must reboot Logger for the changes to become effective. See ["Reboot" on page 252](#).

Authentication

Logger supports optional RADIUS password and client certificate authentication. You can enable both authentication mechanisms at the same time. If both are enabled, client certificate authentication overrides RADIUS authentication unless the "Allow password fallback" setting is set to Yes. (For details about "Allow password fallback" setting, see [Step 3 on page 283](#)).

Client Certificate Authentication

Even if SSL client certificate authentication is enabled on the Logger, a user name must be defined on it for users to connect to it. See ["Users" on page 289](#) for specifics about setting up a user name for client certificate authentication.

The default 'admin' user is exempt and can log on without a certificate even if client certificate authentication is configured on a Logger.



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.

To configure client certificate authentication:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** under the Users/Groups section in the left panel to display the Authentication Settings page in the Authentication tab.

- 3 Update any or all of the parameters listed in the following table:

Parameter	Description
Use client certificate	Select Yes to enable client certificate authentication. Default: No
Require additional password	Select Yes to require a password, in addition to a client certificate, for authentication. This is the password configured for a user's name on Logger. (See "Users" on page 289 for more information.) Default: No
Allow password fallback	Select Yes if a user should be allowed to log in to Logger using only the RADIUS or local password when a certificate is not available or is invalid. Default: No

- 4 Click **Save Settings** to make the changes, or click another tab to cancel.

- 5 Click **Reboot** in the left panel to reboot the appliance.

RADIUS Authentication

If RADIUS authentication is enabled, only user names that are defined as Logger users (see ["Users" on page 289](#)) and are found on the RADIUS server will be able to log in. That is, RADIUS users also require user accounts on Logger. User names must match, but passwords may be different--users will use their RADIUS password to log in.

Whether or not RADIUS authentication is enabled, the default 'admin' user will be able to log in to Logger without having a matching user name on the RADIUS server.

To configure RADIUS authentication settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** under the Users/Groups section in the left panel to display the Authentication Settings page in the Authentication tab.
- 3 Update any or all of the parameters listed in the following table:

Parameter	Description
Use RADIUS authentication	Select Yes to enable RADIUS authentication. The default is No .
Allow local password fallback	Select Yes if a user should be allowed to log in to Logger using the local password when RADIUS authentication fails or is not available. Default: No
RADIUS server hostname[:port]	The host name and port of the RADIUS server.
Shared authentication secret	The RADIUS passphrase
NAS IP Address	The IP address of the Network Access Server (NAS).
Request timeout (in seconds)	How long to wait for a response from the RADIUS server (in seconds). Default is 10 .

Parameter	Description
Number of retries	Number of times to retry a RADIUS request. The default is 1 .

- 4 Click **Save Settings** to make the changes, or click another tab or sub-menu to cancel.

Groups

Logger users are granted permissions by membership in a user group. A user group is a set of permissions and a set of users. User groups have types, such as Logger user groups, or Filter user groups.

User Groups

Groups are organized by type, as shown in [Figure 7-17](#). Each user group is one of the following types: System Admin, Logger Rights, Logger Search, or Logger Report.

Each type has a default user group pre-defined, and the default user group has all privileges for its type enabled. To authorize a subset of the default user group's privileges, create a new User Group (as described below) and revoke some privileges. Then move restricted users from the default user group into the newly created group.

Table 7-1 System Admin Groups

Section	Privilege
Reboot	Reboot Logger. (See "Reboot" on page 252.)
Update	Update Logger. (See "License & Update" on page 259.)
	Enable Maintenance Mode (See "System Maintenance" on page 238.)
System Information	Process Status. (See "Process Status" on page 260.)
	RAID Controller. (See "RAID Controller" on page 270.)
SSL Certificates	Generate SSL Certificate Signing Request (CSR). (See "Generating a Certificate Signing Request" on page 271.)
	Install new SSL certificates. (See "Installing a Signed Certificate" on page 272.)
Platform Settings	Configure DNS settings. (See "DNS Settings" on page 253.)
	Configure network settings. (See "Network" on page 254.)
	Configure time settings. (See "Time/NTP" on page 256.)
	Configure SMTP settings. (See "SMTP Settings" on page 258.)
	Configure static routes. (See "Static Routes" on page 259.)
	Configure Hosts File. (See "Hosts" on page 253.)
	Configure Security Settings. (See "Security" on page 270.)
External File Systems	Configure NFS, CIFS, and SAN settings. (See "Storage" on page 263, "CIFS Settings" on page 263, and "SAN" on page 267.)

Section	Privilege
Global Settings	<p>Configure login settings. (See “Authentication Settings” on page 279.)</p> <p>Configure password settings. (See “Password” on page 281.)</p> <p>Configure password authentication. (See “Authentication” on page 282.)</p> <p>Configure audit forwarding destination. (See “Logs - Audit Forwarding” on page 262.)</p>
System Logs	<p>View Audit Logs. (See “Logs - Audit and Error” on page 262.)</p> <p>View Error Logs. (See “Logs - Audit and Error” on page 262.)</p>
User/Groups	<p>Manage users. (See “Users” on page 289.)</p> <p>Manage user groups. (See “User Groups” on page 284.)</p> <p>Run user entitlement reports.</p>
Console Access	<p>Allow console access. (See “Connecting to the Command Line Interface” on page 17.)</p> <p>Control support login access. (See “Support Login” on page 261.)</p>

Table 7-2 **Logger Rights Groups**

Section	Privilege
Monitor	<p>Monitor Logger throughput. (See “Monitor” on page 37.)</p> <p>Monitor Logger throughput on remote peers. (See “Monitor” on page 37 and “Peer Loggers” on page 231.)</p>
Application Options	<p>View options. (See “Options” on page 37.)</p> <p>Edit, save, and remove options. (See “Options” on page 37.)</p>
Filters	<p>Use and view shared filters. (See “Filters” on page 222.)</p> <p>Edit, save, and remove shared filters. (See “Filters” on page 222.) Also, import and export filters.</p>
Peers	<p>View registered peers. (See “Peer Loggers” on page 231.)</p> <p>Edit, save, and remove registered peers. (See “Peer Loggers” on page 231.)</p>
Devices and Device Groups	<p>View devices. (See “Devices” on page 182.)</p> <p>Edit, save, and remove devices. (See “Devices” on page 182.)</p> <p>View device groups. (See “Device Groups” on page 183.)</p> <p>Edit, save, and remove device groups. (See “Device Groups” on page 183.)</p>

Section	Privilege
Receivers	View receivers. (See "Receivers" on page 193.) Edit, save, and remove receivers. (See "Receivers" on page 193.)
Forwarders and Alerts	View forwarders and alerts. (See "Forwarders" on page 199 and "Alerts" on page 207.) Edit, save, and remove forwarders and alerts. (See "Forwarders" on page 199 and "Alerts" on page 207.) For alerts, this privilege enables you to import and export them.
ESM Connectors	View ESM connectors. (See "ESM Destinations" on page 203.) Edit, save, and remove ESM connectors. (See "ESM Destinations" on page 203.)
Search Filters	View search group filters (aka user group filters). (See "Search Group Filters" on page 223.) Edit, save, and remove search group filters. (See "Search Group Filters" on page 223.)
Configuration Backup	View backups. (See "Configuration Backup and Restore" on page 235.) Edit, save, and remove backups. (See "Configuration Backup and Restore" on page 235.)
Retrieve Logs	Download system logs. (See "Retrieve Logs" on page 246.)
Scheduling	View scheduled tasks. (See "Scheduled Tasks" on page 221.)
Storage Groups	View storage groups. (See "Storage Groups" on page 187.) Edit and add storage groups. (See "Storage Groups" on page 187.)
Event Archive/Restore	View event archives. (See "Event Archives" on page 185.) Edit, save, and remove event archives. (See "Event Archives" on page 185.)
Saved Search	View Saved Search. (See "Saved Searches" on page 225.) Edit, save, and remove Saved Search. (See "Scheduled Saved Search" on page 226.)

Table 7-3 Logger Search Groups

Section	Privilege
Search	Search for events. (See "The Need to Search Events" on page 43.) Search for events on remote peers. (See "Searching Peer Loggers (Distributed Search)" on page 73.)

Table 7-4 Logger Report Groups.

Section	Privilege
Report	<p>Global access to all report objects and permission to change reporting configuration. (See Chapter 5, Reporting, on page 93.)</p> <p>If this user right is set to Yes, it overrides all other rights. Therefore, to granularly control user rights for reports, set this right to No and then selectively set other rights to Yes.</p> <p>Edit, save, and delete report queries, parameters, and parameter values groups. (See information on queries, parameters, and parameter value groups in “Designing Reports” on page 124.)</p> <p>Edit and save report style. This overrides the corresponding permission on individual report groups. (See “Applying Report Template Styles” on page 172.)</p> <p>View all published reports. This overrides the corresponding permission on individual report groups. (See Chapter 5, Reporting, on page 93.)</p> <p>View, run, and schedule all reports. This overrides the corresponding permission on individual report groups. (See “Running, Viewing, and Publishing Reports” on page 113 and “Scheduling Reports” on page 173.)</p> <p>Edit and save reports. This overrides the corresponding permission on individual report groups. (See Chapter 5, Reporting, on page 93.)</p>

Each individual report group--Default Reports, Configuration Monitoring, Intrusion Monitoring, or SANS Top 5, for example--will have its own set of rights. Each report group will have privileges for View published reports, View, run, and schedule reports, and Edit and save reports.

Groups Administration

Add User Group

System Admin Groups			
Name	Description	Number Of Members	Action
Default System Admin Group	The default group allows all system admin operations e.g., reboot, install SSL, platform configuration, and so on.	1	edit
Logger Rights Groups			
Name	Description	Number Of Members	Action
Default Logger Rights Group	The default group allows all logger operations e.g., monitor, creating and editing filters, peers, devices, receivers, forwarders, and so on.	1	edit delete
Logger Search Groups			
Name	Description	Number Of Members	Action
Default Logger Search Group	The default search group allows both local and distributed searches.	1	edit delete
Logger Report Groups			
Name	Description	Number Of Members	Action
Default Logger Report Group	The default report group allows all report operations e.g., view published report, view, run, schedule, edit, delete all reports, and so on.	1	edit delete

Figure 7-17 Groups page

Maximum number of user groups that can be created on Logger: No limit.

To create a new user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Groups** under the Users/Groups section in the left panel to display the page shown in [Figure 7-17](#).
- 3 Click **Add User Group**.
- 4 Enter the definition of the new group.
 - a Define the group by choosing a type and entering a name and description.
 - b Define the group's rights and permissions.
 - c Optionally, add users to the new group.
- 5 Click **Save Group**.

To edit a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Groups** under the Users/Groups section in the left panel to display the page shown in [Figure 7-17](#).
- 3 Identify the group to be edited and click the **edit** link.
- 4 Update the user group information as necessary.
- 5 Click **Save Group**.

To delete a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Groups** under the Users/Groups section in the left panel to display the page shown in [Figure 7-17](#).
- 3 Identify the user group to be deleted and click the **delete** link.

Users

Add User

Add User Information:

Login:

First Name: (Use Client DN)

Last Name:

Password:

Confirm Password:

Email:

Phone Number:

System Admin Groups

Select	Name	Description
<input type="radio"/>	Default System Admin Group	The default group allows all system admin operations e.g., reboot, install SSL, platform configuration, and so on.

clear

Logger Rights Groups

Select	Name	Description
<input type="radio"/>	Default Logger Rights Group	The default group allows all logger operations e.g., monitor, creating and editing filters, peers, devices, receivers, forwarders, alerts, and so on.

clear

Logger Report Groups

Select	Name	Description
<input type="radio"/>	Default Logger Report Group	The default report group allows all report operations e.g., view published report, view, run, schedule, edit, delete all reports, and so on.

clear

Logger Search Groups

Select	Name	Description
<input type="radio"/>	Default Logger Search Group	The default search group allows both local and distributed searches.

clear

Save User

Figure 7-18 Add User page

Maximum number of users that can be created on Logger: No limit.

To create a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Users** under the Users/Groups section in the left panel.
- 3 Click **Add User** in the right panel. The form shown in [Figure 7-18](#) is displayed.
- 4 Enter the following parameters.

Parameter	Description
Login	A login name for the user
First Name	User's first name.
	<p>If you enabled SSL client authentication (see "SSL Client Authentication (CAC Authentication)" on page 273), click Use Client DN to enter the Distinguished Name (Certificate Subject) information for the user. Distinguished Name should be in this format:</p> <p>ST=California, C=US, L=Cupertino, O=ArcSight, Inc., OU=Engg Team, CN=UserA D/emailAddress=email@xyz.com</p> <p>Obtain the DN information for a user from the browser that the user will use to connect to the Logger. For example, on Firefox 3.0, click Tools > Options > Encryption > View Certificates > Your Certificates > Select the certificate > View.</p>

Parameter	Description
Last Name	User's last name. This information is not required when creating a user for SSL client authentication.
Password	A password for the user.
Confirm Password	Reenter the password.
Email	An e-mail address for the user.
Phone Number	User's phone number.
Select User Groups	Select the groups to which this user belongs. This setting controls the privileges a user has on this Logger.

5 Click **Save User**.

To edit a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Users** under the Users/Groups section in the left panel.
- 3 Identify the user to be edited and click the **edit** link. Update the user information as necessary.
- 4 Click **Save User**.

To delete a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Users** under the Users/Groups section in the left panel.
- 3 Identify the user to be deleted and click the **delete** link.
- 4 Confirm the delete operation.

Change Password

Password management is the responsibility of individual users. Users can choose their password, and they may change their password as often as desired.

Change Password For Default Admin

Figure 7-19 Change Password page

To change your password:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Change Password** under the Users/Groups section in the left panel to display the Change Password for <User Name> page, as shown in [Figure 7-19 on page 290](#).
- 3 Enter the old password, the new password, and enter the new password a second time to confirm.
- 4 Click **Set Password**.



Passwords are subject to the password policy specified by the Admin user. See ["Password" on page 281](#).

Section 2: Software Version Logger Administration

This section describes the System Administration settings that are applicable to the software version of Logger.

This section contains the following topics.

["Network - SMTP Settings" on page 291](#)

["Process Status" on page 292](#)

["Logs - Audit and Error" on page 293](#)

["Logs - Audit Forwarding" on page 293](#)

["Users/Groups - Groups" on page 294](#)

["Users/Groups - Change Password" on page 300](#)

["Using a CA-signed Certificate on Software Version of Logger" on page 301](#)

Network - SMTP Settings

Alerts use Simple Mail Transfer Protocol (SMTP) to send e-mail.

The screenshot shows the ArcSight Network Settings page. The top navigation bar includes 'Monitor', 'Analyze', 'Reports', 'Configuration', and 'System Admin'. The 'System Admin' section is active, showing a sidebar with 'System', 'Logs', and 'Users/Groups'. The main content area is titled 'ArcSight Network Settings' and contains the 'SMTP Settings' section. This section has two input fields: 'Primary SMTP Server' with the value '192.168.10.124' and 'Outgoing Email Address' with the value 'admin@arcsight.com'. There is an 'Update Settings' button at the bottom of the form. The top right of the page displays system metrics: EPS In: 0, EPS Out: 0, CPU Load: 8%, and a 'Logout' button.

Figure 7-20 Simple Mail Transfer Protocol (SMTP) settings

To change SMTP configuration:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.

- 3 In the **SMTP** tab on the ArcSight Platform Settings page, enter new values for the following fields.

Parameter	Description
Primary SMTP Address	Enter the IP address of the SMTP server that will process outgoing e-mail.
Outgoing Email Address	The e-mail address that will appear in the From: field of outbound e-mail.

- 4 Click **Update Settings** to make the changes, or click another tab or sub-menu to cancel. Changes take effect immediately; reboot is not required.

Process Status

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Process Status** from the System section to display a page similar to the one shown in [Figure 7-7](#).

Process Status

System	Status	Load	CPU	Memory
logger	running	[0.30] [0.28] [0.25] 1.4%us 0.2%sy 0.1%wa	1909144 kB [7.7%]	

Process	Status	Uptime	CPU	Memory	Memory (kB)
apache	running	23h 51m	0.0%	0.0%	7300
aps	running	23h 51m	0.0%	0.7%	191736
connector	running	23h 51m	0.0%	0.0%	568
insp	running	23h 49m	0.0%	0.1%	27604
mysqld	running	23h 51m	0.0%	0.0%	21460
nullmailer	running	23h 49m	0.0%	0.0%	816
postgresql	running	23h 51m	0.0%	0.0%	9244
processors	running	23h 49m	0.0%	0.6%	149160
receivers	running	23h 49m	0.0%	0.2%	50700
reportengine	running	23h 49m	0.0%	0.3%	94856
servers	running	23h 51m	0.0%	2.3%	587224
web	running	23h 49m	1.3%	1.6%	406388

Figure 7-21 Process Status page

In the process list, processors refers to Forwarders.

Each process is a hyperlink. Clicking on an individual process displays more detail about that process, as shown in [Figure 7-7](#).

Status detail for apache

Parameter	Value
children	15
cpu_percent	0.0%
cpu_percent_total	0.0%
data_collected	Wed Sep 2 13:58:02 2009
memory_kilobytes	7300
memory_kilobytes_total	250200
memory_percent	0.0%
memory_percent_total	1.0%
monitoring_status	monitored
parent_pid	1
pid	4279
status	running
uptime	23h 53m

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

[BACK](#)

[RESTART](#)

Figure 7-22 Process Status detail for apache

Logs - Audit and Error

Logger audit and error logs are available for viewing.

The screenshot shows the ArcSight Logger System Admin interface. The top navigation bar includes Monitor, Analyze, Reports, Configuration, and System Admin. The left sidebar lists various system and security options. The main content area is titled 'Search Audit Logs' and contains the following sections:

- Select Audit Type:** A dropdown menu labeled 'Select an audit type:'.
- Select Date Range:** Fields for Start Date (Jun 11 2010) and End Date (Jun 11 2010).
- Select User (optional):** A table with columns for Login, First Name, Last Name, Email, and Phone Groups. The table lists 'adminDefault' and 'Admin'.
- View Audit Logs:** A button at the bottom of the search section.

To view Audit or Error logs:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Audit** (for audit logs) or **Error** (for Error logs) from the Logs section.
- 3 Select the type of log—Application or Platform.
- 4 Select the date range for which you want to obtain the log.
- 5 Click **View Error Logs**.



To search again after clicking **View Audit Logs**, use the browser's Back button.

Audit logs, as Common Event Format (CEF) audit events, can be sent to ArcSight ESM directly for analysis and correlation because the Logger Forwarder supports ESM Manager's event protocol.

For more information about audit event forwarding, see [“Logs - Audit Forwarding” on page 262](#). For information about audit events that you can forward, see [Appendix E, Logger Audit Events, on page 401](#).

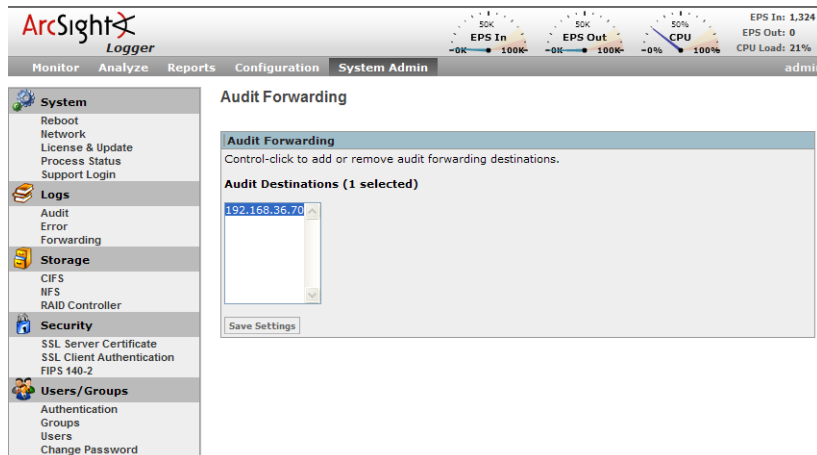
Logs - Audit Forwarding

For information about audit events that you can forward, see [Appendix E, Logger Audit Events, on page 401](#).

To forward audit events to specific destinations:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Forwarding** from the Logs section.

- 3 Select destinations from the Audit Destinations list, as shown in the following figure. Click on a destination to select a single destination, or Ctrl+click to select or de-select multiple destinations. The destinations are ESM destinations that you configure on the ESM Destinations page (**Configuration > Event Input/Output > ESM Destinations**).



- 4 Click **Save Settings**.

Users/Groups - Groups

Logger users are granted permissions by membership in a user group. A user group is a set of permissions and a set of users. User groups have types, such as Logger user groups, or Filter user groups.

User Groups

Groups are organized by type, as shown in [Figure 7-17](#). Each user group is one of the following types: System Admin, Logger Rights, Logger Search, or Logger Report.

Each type has a default user group pre-defined, and the default user group has all privileges for its type enabled. To authorize a subset of the default user group's privileges, create a new User Group (as described below) and revoke some privileges. Then move restricted users from the default user group into the newly created group.

Table 7-5 System Admin Groups

Section	Privilege
Reboot	Reboot Logger. (See "Reboot" on page 252.)
Update	Update Logger.
	Enable Maintenance Mode (See "System Maintenance" on page 238.)
System Information	Process Status. (See "Process Status" on page 260.)
	Hard Disk SMART Data.
SSL Certificates	Generate SSL Certificate Signing Request (CSR). (See "Generating a Certificate Signing Request" on page 271.)
	Install new SSL certificates. (See "Installing a Signed Certificate" on page 272.)

Section	Privilege
Platform Settings	<p>Configure DNS settings. (See "DNS Settings" on page 253.)</p> <p>Configure network settings. (See "Network" on page 254.)</p> <p>Configure time settings. (See "Time/NTP" on page 256.)</p> <p>Configure SMTP settings. (See "SMTP Settings" on page 258.)</p> <p>Configure static routes. (See "Static Routes" on page 259.)</p> <p>Configure Hosts File. (See "Hosts" on page 253.)</p> <p>Configure Security Settings. (See "Security" on page 270.)</p>
External File Systems	<p>Configure NFS, CIFS, and SAN settings. (See "Storage" on page 263, "CIFS Settings" on page 263, and "SAN" on page 267.)</p>
Global Settings	<p>Configure login settings. (See "Authentication Settings" on page 279.)</p> <p>Configure password settings. (See "Password" on page 281.)</p> <p>Configure password authentication. (See "Authentication" on page 282.)</p> <p>Configure audit forwarding destination. (See "Logs - Audit Forwarding" on page 262.)</p>
System Logs	<p>View Audit Logs. (See "Logs - Audit and Error" on page 262.)</p> <p>View Error Logs. (See "Logs - Audit and Error" on page 262.)</p>
User/Groups	<p>Manage users. (See "Users" on page 289.)</p> <p>Manage user groups. (See "User Groups" on page 284.)</p> <p>Run user entitlement reports.</p>
Console Access	<p>Allow console access. (See "Connecting to the Command Line Interface" on page 17.)</p> <p>Control support login access. (See "Support Login" on page 261.)</p>

Table 7-6 Logger Rights Groups

Section	Privilege
Monitor	<p>Monitor Logger throughput. (See "Monitor" on page 37.)</p> <p>Monitor Logger throughput on remote peers. (See "Monitor" on page 37 and "Peer Loggers" on page 231.)</p>
Application Options	<p>View options. (See "Options" on page 37.)</p> <p>Edit, save, and remove options. (See "Options" on page 37.)</p>

Section	Privilege
Filters	<p>Use and view shared filters. (See “Filters” on page 222.)</p> <p>Edit, save, and remove shared filters. (See “Filters” on page 222.) Also, import and export filters.</p>
Peers	<p>View registered peers. (See “Peer Loggers” on page 231.)</p> <p>Edit, save, and remove registered peers. (See “Peer Loggers” on page 231.)</p>
Devices and Device Groups	<p>View devices. (See “Devices” on page 182.)</p> <p>Edit, save, and remove devices. (See “Devices” on page 182.)</p> <p>View device groups. (See “Device Groups” on page 183.)</p> <p>Edit, save, and remove device groups. (See “Device Groups” on page 183.)</p>
Receivers	<p>View receivers. (See “Receivers” on page 193.)</p> <p>Edit, save, and remove receivers. (See “Receivers” on page 193.)</p>
Forwarders and Alerts	<p>View forwarders and alerts. (See “Forwarders” on page 199 and “Alerts” on page 207.)</p> <p>Edit, save, and remove forwarders and alerts. (See “Forwarders” on page 199 and “Alerts” on page 207.) For alerts, this privilege enables you to import and export them.</p>
ESM Connectors	<p>View ESM connectors. (See “ESM Destinations” on page 203.)</p> <p>Edit, save, and remove ESM connectors. (See “ESM Destinations” on page 203.)</p>
Search Filters	<p>View search group filters (aka user group filters). (See “Search Group Filters” on page 223.)</p> <p>Edit, save, and remove search group filters. (See “Search Group Filters” on page 223.)</p>
Configuration Backup	<p>View backups. (See “Configuration Backup and Restore” on page 235.)</p> <p>Edit, save, and remove backups. (See “Configuration Backup and Restore” on page 235.)</p>
Retrieve Logs	Download system logs. (See “Retrieve Logs” on page 246.)
Scheduling	View scheduled tasks. (See “Scheduled Tasks” on page 221.)
Storage Groups	<p>View storage groups. (See “Storage Groups” on page 187.)</p> <p>Edit and add storage groups. (See “Storage Groups” on page 187.)</p>
Event Archive/Restore	<p>View event archives. (See “Event Archives” on page 185.)</p> <p>Edit, save, and remove event archives. (See “Event Archives” on page 185.)</p>

Section	Privilege
Saved Search	View Saved Search. (See “Saved Searches” on page 225.) Edit, save, and remove Saved Search. (See “Scheduled Saved Search” on page 226.)

Table 7-7 Logger Search Groups

Section	Privilege
Search	Search for events. (See “The Need to Search Events” on page 43.) Search for events on remote peers. (See “Searching Peer Loggers (Distributed Search)” on page 73.)

Table 7-8 Logger Report Groups.

Section	Privilege
Report	Global access to all report objects and permission to change reporting configuration. (See Chapter 5, Reporting, on page 93.) If this user right is set to Yes, it overrides all other rights. Therefore, to granularly control user rights for reports, set this right to No and then selectively set other rights to Yes. Edit, save, and delete report queries, parameters, and parameter values groups. (See information on queries, parameters, and parameter value groups in “Designing Reports” on page 124.) Edit and save report style. This overrides the corresponding permission on individual report groups. (See “Applying Report Template Styles” on page 172.) View all published reports. This overrides the corresponding permission on individual report groups. (See Chapter 5, Reporting, on page 93.) View, run, and schedule all reports. This overrides the corresponding permission on individual report groups. (See “Running, Viewing, and Publishing Reports” on page 113 and “Scheduling Reports” on page 173.) Edit and save reports. This overrides the corresponding permission on individual report groups. (See Chapter 5, Reporting, on page 93.)

Each individual report group--Default Reports, Configuration Monitoring, Intrusion Monitoring, or SANS Top 5, for example--will have its own set of rights. Each report group

will have privileges for View published reports, View, run, and schedule reports, and Edit and save reports.

Groups Administration

Add User Group

System Admin Groups			
Name	Description	Number Of Members	Action
Default System Admin Group	The default group allows all system admin operations e.g., reboot, install SSL, platform configuration, and so on.	1	edit

Logger Rights Groups			
Name	Description	Number Of Members	Action
Default Logger Rights Group	The default group allows all logger operations e.g., monitor, creating and editing filters, peers, devices, receivers, forwarders, and so on.	1	edit delete

Logger Search Groups			
Name	Description	Number Of Members	Action
Default Logger Search Group	The default search group allows both local and distributed searches.	1	edit delete

Logger Report Groups			
Name	Description	Number Of Members	Action
Default Logger Report Group	The default report group allows all report operations e.g., view published report, view, run, schedule, edit, delete all reports, and so on.	1	edit delete

Figure 7-23 Groups page

Maximum number of user groups that can be created on Logger: No limit.

To create a new user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Groups** under the Users/Groups section in the left panel to display the page shown in [Figure 7-17](#).
- 3 Click **Add User Group**.
- 4 Enter the definition of the new group.
 - a Define the group by choosing a type and entering a name and description.
 - b Define the group's rights and permissions.
 - c Optionally, add users to the new group.
- 5 Click **Save Group**.

To edit a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Groups** under the Users/Groups section in the left panel to display the page shown in [Figure 7-17](#).
- 3 Identify the group to be edited and click the **edit** link.
- 4 Update the user group information as necessary.
- 5 Click **Save Group**.

To delete a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Groups** under the Users/Groups section in the left panel to display the page shown in [Figure 7-17](#).

- 3 Identify the user group to be deleted and click the **delete** link.

Users

Add User

Figure 7-24 Add User page

Maximum number of users that can be created on Logger: No limit.

To create a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Users** under the Users/Groups section in the left panel.
- 3 Click **Add User** in the right panel. The form shown in [Figure 7-18](#) is displayed.
- 4 Enter the following parameters.

Parameter	Description
Login	A login name for the user
First Name	User's first name. If you enabled SSL client authentication (see "SSL Client Authentication (CAC Authentication)" on page 273), click Use Client DN to enter the Distinguished Name (Certificate Subject) information for the user. Distinguished Name should be in this format: ST=California, C=US, L=Cupertino, O=ArcSight, Inc., OU=Engg Team, CN=UserA D/emailAddress=email@xyz.com Obtain the DN information for a user from the browser that the user will use to connect to the Logger. For example, on Firefox 3.0, click Tools > Options > Encryption > View Certificates > Your Certificates > Select the certificate > View .

Parameter	Description
Last Name	User's last name. This information is not required when creating a user for SSL client authentication.
Password	A password for the user.
Confirm Password	Reenter the password.
Email	An e-mail address for the user.
Phone Number	User's phone number.
Select User Groups	Select the groups to which this user belongs. This setting controls the privileges a user has on this Logger.

5 Click **Save User**.

To edit a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Users** under the Users/Groups section in the left panel.
- 3 Identify the user to be edited and click the **edit** link. Update the user information as necessary.
- 4 Click **Save User**.

To delete a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Users** under the Users/Groups section in the left panel.
- 3 Identify the user to be deleted and click the **delete** link.
- 4 Confirm the delete operation.

Users/Groups - Change Password

Password management is the responsibility of individual users. Users can choose their password, and they may change their password as often as desired.

Change Password For Default Admin

Figure 7-25 Change Password page

To change your password:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Change Password** under the Users/Groups section in the left panel to display the Change Password for <User Name> page, as shown in [Figure 7-19 on page 290](#).
- 3 Enter the old password, the new password, and enter the new password a second time to confirm.
- 4 Click **Set Password**.



Passwords are subject to the password policy specified by the Admin user. See ["Password" on page 281](#).

Using a CA-signed Certificate on Software Version of Logger

Logger ships with a self-signed certificate. Although you can use this certificate, ArcSight strongly recommends using a CA-signed certificate.

The first step in configuring an SSL certificate is to generate a Certificate Signing Request (CSR). The CSR must be generated on the Logger appliance for which you are requesting a certificate. The resulting CSR should be sent to a CA, such as VeriSign, which responds with a signed certificate file.

To generate a CSR to obtain a CA-signed certificate on software version of Logger:

- 1 Ensure that you are logged in as "root" on the system on which the software version of Logger is installed.

- 2 Run this command to create a new private key and a server certificate request:

```
OPENSSL_FIPS=1 /opt/local/openssl/bin/openssl req -text -newkey
rsa:1024 -keyout server.key -out newreq.csr
```

You are prompted to enter a password and certificate details. Make sure you note down the password and keep it safely. The bit size in the certificate details is set to 1024 by default. You can increase or decrease the size to suit your needs.

- 3 Send the generated CSR to a CA. Once you receive the signed certificate from the CA, go to the next step.

- 4 Run this command to archive any previously applied certificate files:

```
cd /opt/local/apache/conf/ssl.crt
mv server.pem server.pem.good
mv server.crt server.crt.good
```

- 5 Run this command to install the signed certificate file you obtained from CA in Step 3:

```
cp newreq.csr /opt/local/apache/conf/ssl.crt/server.crt
```

- 6 Run this command to copy the server key to the Apache directory:

```
/opt/local/openssl/bin/rsa -in server.key -out
/opt/local/apache/conf/ssl.crt/server.pem
```

- 7 Restart the Apache server for the new certificate and the server key to take effect:

```
/sbin/service apache stop  
/sbin/service apache start
```

Managing Connectors on Connector Appliance

The information in this chapter is applicable only to Logger **appliance platforms with integrated Connector Appliance**.

The chapter discusses the following topics.

- [“SmartConnector Overview” on page 304](#)
- [“Navigating the Manage Connectors Tab” on page 305](#)
- [“Locations” on page 307](#)
- [“Hosts” on page 310](#)
- [“Containers” on page 316](#)
- [“Connectors” on page 331](#)
- [“Configuration Suggestions for SmartConnector Types” on page 349](#)
- [“Troubleshooting Connector Communication Issues” on page 353](#)

SmartConnector Overview

You can manage the configuration of these kinds of SmartConnectors:

- **Local (on-board) SmartConnectors:** Pre-installed connectors on the Logger appliance running Connector Manager.
- **Remote Connector Appliance SmartConnectors:** Pre-installed connectors on a remotely-managed Connector Appliance
- **Software-based SmartConnectors:** Software-based SmartConnectors installed manually on a remote host

A SmartConnector configuration consists of properties such as name and type, and a set of *parameters* that customize how the SmartConnector works in a specific environment. Parameters vary based on the type of SmartConnector; 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), 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.



Note

The SmartConnectors you manage are configured automatically to run as *services* or *daemons*.

Individual software-based SmartConnectors are described in ArcSight documents specific to those connectors, including the SmartConnector-specific configuration guides available with each SmartConnector. You can also find general SmartConnector information in the *SmartConnector User's Guide*. All of these documents are available from the ArcSight Customer Support site.

Navigating the Manage Connectors Tab

The Manage Connectors tab enables you to configure and organize SmartConnectors. This section describes the user interface elements and explains how to use them effectively.

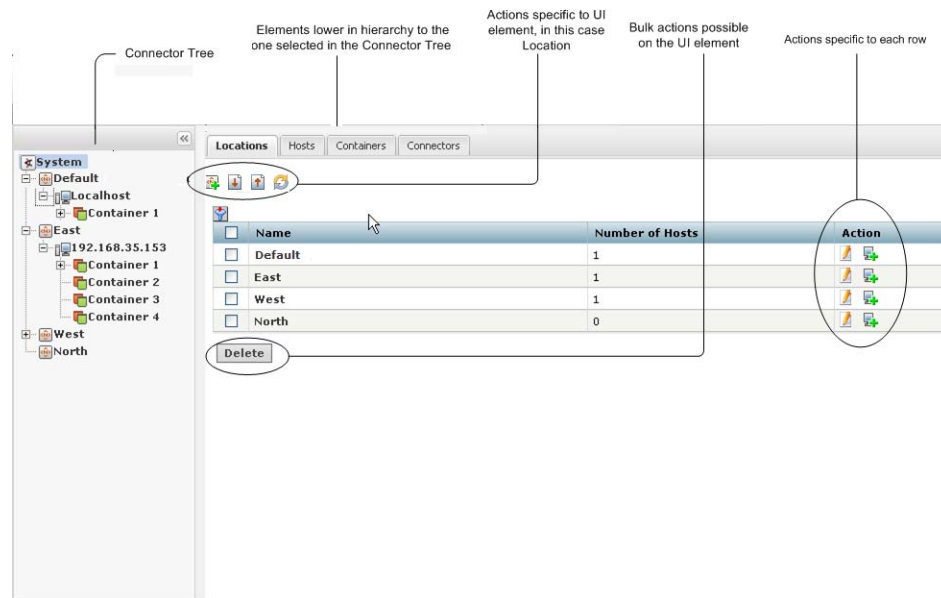
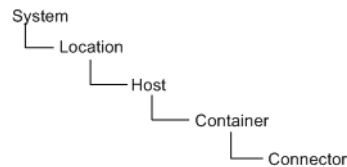


Figure 8-1 Managing Connectors

The Connector tree (the left panel of the window shown in [Figure 8-1](#)) organizes SmartConnectors into a hierarchy as follows:

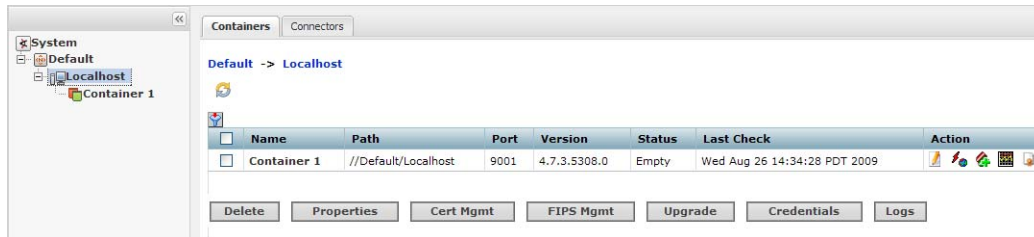


Each connector you manage belongs to a container; each container belongs to a host; each host belongs to a location; and, all locations belong to root of the System.

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.



Note



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.



Tip

- 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 check box in the table header to check or uncheck all check boxes in a single column.

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.



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.



Note

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**.

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.
- 4 Follow instructions in the Upload CSV 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 **I want to import the certificates to Connector Appliance from the containers**, then click **Next** to import the certificates and continue.
 - ◆ Select **I do not want to 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.



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. The new name needs to follow the naming guidelines described in [“Adding a Location” on page 308](#).


You can edit 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 **OK**.

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 311](#).

Hosts

A host is a computer on a network, associated with an IP address, on which connectors are installed. A host can be of two types:

- Localhost (the local Connector Appliance or Logger appliance running Connector Manager) or a remote Connector Appliance. A host can contain up to **eight** containers. 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.
- 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.

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 308](#).

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 313](#).

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**.

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.
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 SmartConnectors applicable to each model type and container specifics, see the <i>ArcSight Appliance Specifications</i> document. This document is available on the ArcSight Customer Support site at https://support.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 SmartConnectors 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 313.

- 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 **I want to import the certificates to Connector Appliance from the containers**, then click **Next** to import the certificates and add the host.

- ◆ Select **I do not want to 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.



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 310](#) 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 starting at which the scan operation will scan for containers.
Ending Port	The port number on the host at which the scan operation will end scanning for containers.
User	The user name that the system will use 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 **I want to import the certificates to Connector Appliance from the containers**, then click **Next** to import the certificates and continue.
- ◆ Select **I do not want to 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 309](#)) or move an existing host (as described in [“Moving a Host to a Different Location” on page 314](#)).

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.
- Container 1 on the remotely-managed Connector Appliance needs to be running the build specified in the Connector Appliance v5.5 Release Notes.

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 ArcSight Customer Support site to the AUP (Upgrade) repository.

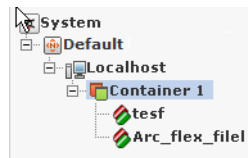
This step is only required if the version that you want to upgrade does not already exist in the repository.
- 2 Push the .AUP upgrade file to the remote Connector Appliances, as follows:
 - a Click **Configuration** > **Manage Connectors**.
 - b From the System-level page, click **System** (left panel) > **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 Hosts Upgrade wizard.

Adding a Container to a Host

See [“Adding a Container” on page 317](#).

Containers

A container is a single Java Virtual Machine (JVM) that can run up to four SmartConnectors. The following illustration depicts Container 1 and the connectors it runs.

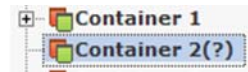


A default number of containers exist on each Connector Appliance. The number depends on the Connector Appliance hardware platform. Each container is identified with a label (Container Name) and an associated port number (9001 or higher).

The Connector Manager on a Logger appliance contains one default container in the default host **Localhost**. You cannot delete this container.

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 327](#).



Viewing All Containers

You can see all the containers you are managing.

To view all containers:

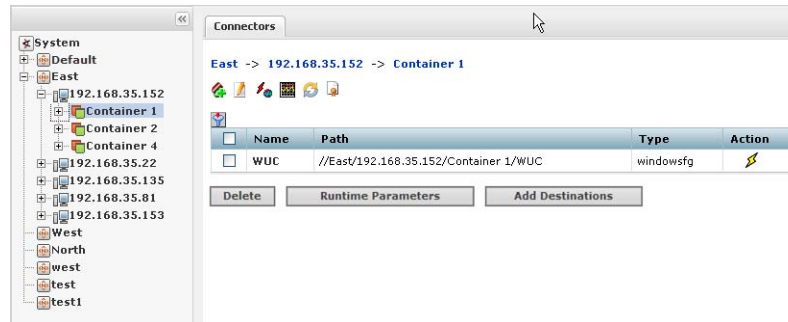
- 1 Click **Configuration** > **Manage Connectors**.
- 2 Click **System** in the left panel. All containers display on the Containers tab in the right panel.

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 313](#). 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 331](#).

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).

User Interface Options	Path
From the Containers page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel).

- 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.

- Enter the new name in the **Name** field and click **Next**.
- 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:

- 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).

- Select the container you want to delete. You can select multiple containers.
- Click **Delete**.

Updating Container Properties

You can update existing container properties (the `agent.properties` file), delete them, or add new ones.

To update container properties:

- 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).

User Interface Options	Path
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 Container Properties Update wizard to update connector properties.



When a property is removed, it is still visible until the container is restarted.

Note

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 Container Password Update wizard to update connector credentials.



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 ArcSight Customer Support.



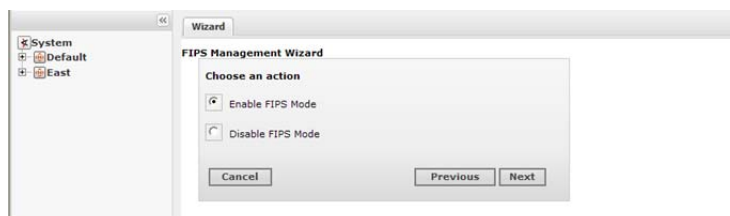
After you enable or disable FIPS mode on a container, check that the appropriate CA certificates are in the trust store of the connectors so that they can validate their configured destinations successfully. If the appropriate CA certificates are not present, you need to add them (refer to [“Managing Certificates on a Container” on page 321](#)).

To enable or disable FIPS mode on a container:

- 1 Click **Configuration** > **Manage Connectors**.
- 2 Use one of these navigation paths:

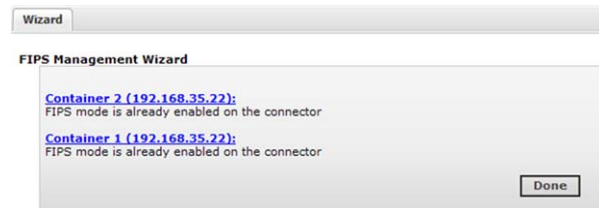
User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (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 Mgmt**, 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](#).

Managing Certificates on a Container

SmartConnectors 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 CA 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 CA 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 325](#).

For information about resolving invalid certificates, see [“Resolving Invalid Certificate Errors” on page 327](#).

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.



- 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.

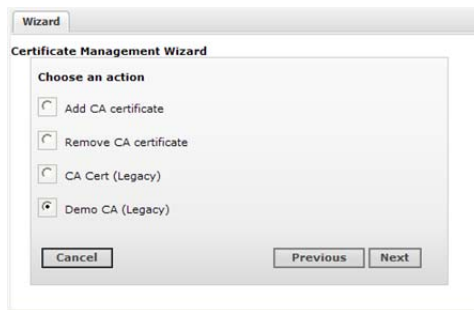
To enable or disable a demo certificate on a container:

- 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).

User Interface Options	Path
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 **Cert Mgmt**, 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 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.

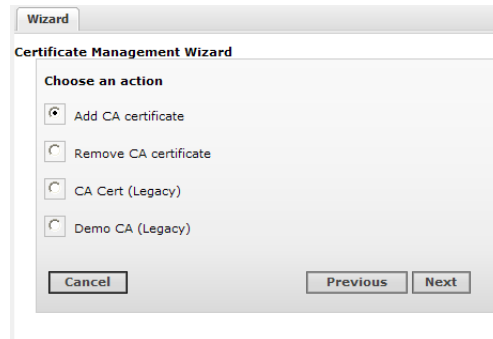
To apply a single CA certificate on a container:

- 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).

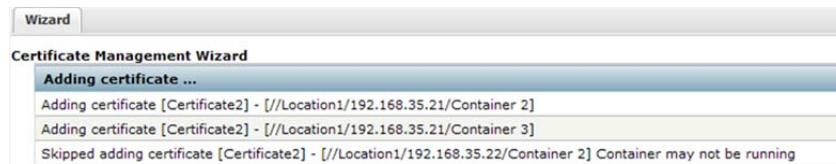
User Interface Options	Path
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Select the container to which you want to add the CA certificate. You can select multiple containers.
- 4 Click **Cert Mgmt**, then click **Next** to run the Certificate Management wizard.
- 5 Click **Add CA Certificate** to add a certificate.



- 6 Follow the instructions in the Certificate Management wizard.

If a container is down or a connector is running an older build, the Certificate Management wizard reports errors in the progress bar and on the Summary page.



Adding a CA Certs File on a Container

You can add a CA Certs file on any container that is in non-FIPS mode.



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.

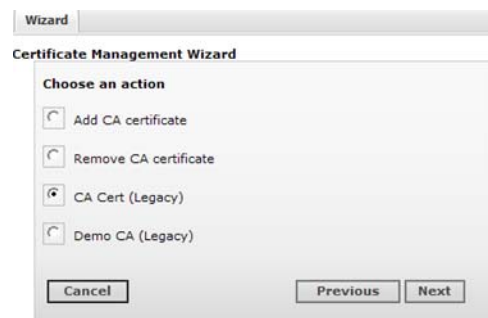
To add a CA Certs file on a container:

- 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).

User Interface Options	Path
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 **Cert Mgmt**, then click **Next** to run the Certificate Management 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 Certificate Management wizard.

After the CA Certs file has been added to a container, the container restarts automatically.

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:

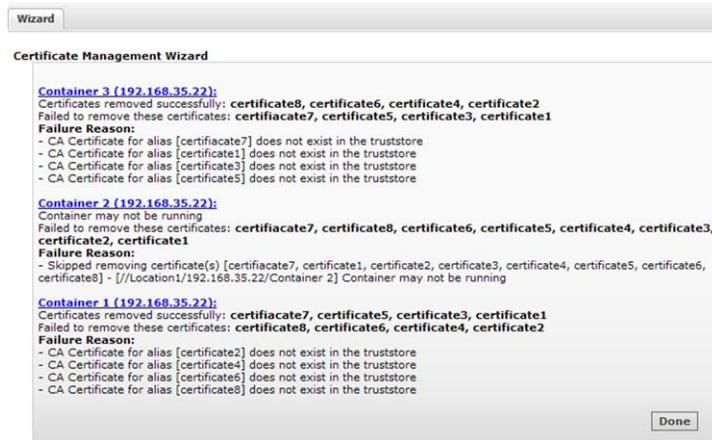
User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).

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 Select the container from which you want to remove the CA certificates. You can select multiple containers.
- 4 Click **Cert Mgmt**, then click **Next** to run the Certificate Management wizard.
- 5 Click **Remove CA 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.

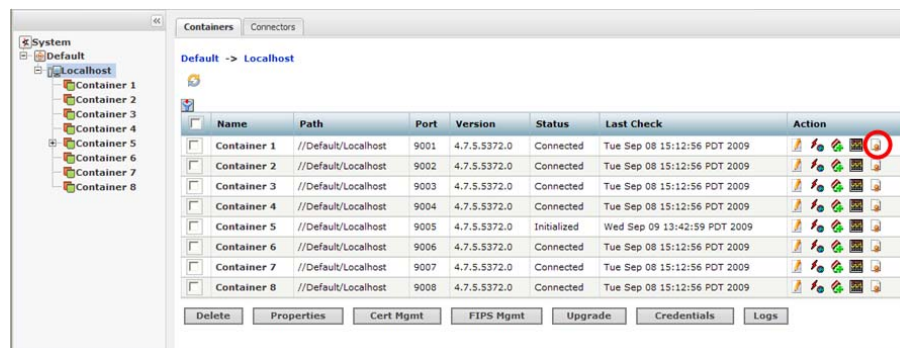
The Certificate Management wizard displays the certificates that are removed successfully in a comma-separated list. Certificates that cannot be removed are shown in a comma-separated list together with a reason why the certificate removal failed.




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.




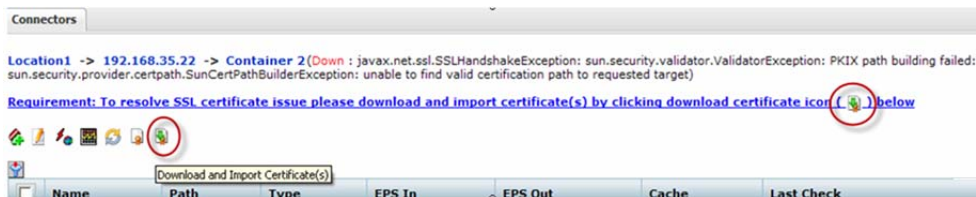
Resolving Invalid Certificate Errors

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 **Finish**.

Upgrading a Container to a Specific Connector Version

A container needs to comprise of at least one connector before you can upgrade it. All connectors in a container are upgraded to the version you select.

To upgrade a container to a specific connector version:

- 1 Upload a connector build AUP from the ArcSight Customer Support site 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.

Viewing Container Logs

You can retrieve and view the logs 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 the container logs, click **Configuration** > **Repositories** > **Logs** > from the top-level menu bar. In the right panel, click  next to the logs 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).
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. 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.

- 4 From the **Group** box, choose which type of data you would like to view. The Group box lists all SmartConnectors 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.

- 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**.

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 ArcSight Customer Support site.

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:

- 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 308](#), [“Adding a Host” on page 311](#), and [“Adding a Container” on page 317](#).
- Follow the configuration best practices described in [“Configuration Suggestions for SmartConnector Types” on page 349](#).

If you are configuring the Check Point OPSEC NG Connector, see [“Configuring the Check Point OPSEC NG Connector” on page 350](#).

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 352](#).

- File-based SmartConnectors 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 [“CIFS Settings” on page 263](#).

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 [“Network File System \(NFS\) Settings” on page 265](#).

- 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 `.sdrfilereader.properties` is appended automatically.


To add a Connector:




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 350](#).

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

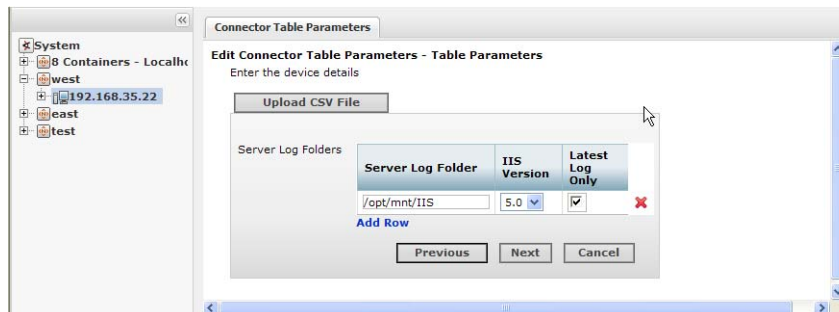
User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (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 SmartConnector. 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**.

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/<Name_of_CIFS_Share>`.)



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).



For detailed information about individual SmartConnector 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 SmartConnector.

- 6 Choose a primary destination for the SmartConnector and enter destination-specific parameters on the following page(s), then click **Next**. Destinations can be one of the following.

- ◆ ArcSight Logger SmartMessage (encrypted)
- ◆ ArcSight Manager (encrypted)
- ◆ CEF Syslog (cleartext, that is, unencrypted)

- 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 **Finish**.

Editing Connector Parameters

ArcSight supports a large number of SmartConnector 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 SmartConnector being configured.

You can edit parameters (simple and table) for a specific connector.

Updating Simple Parameters for a Specific Connector

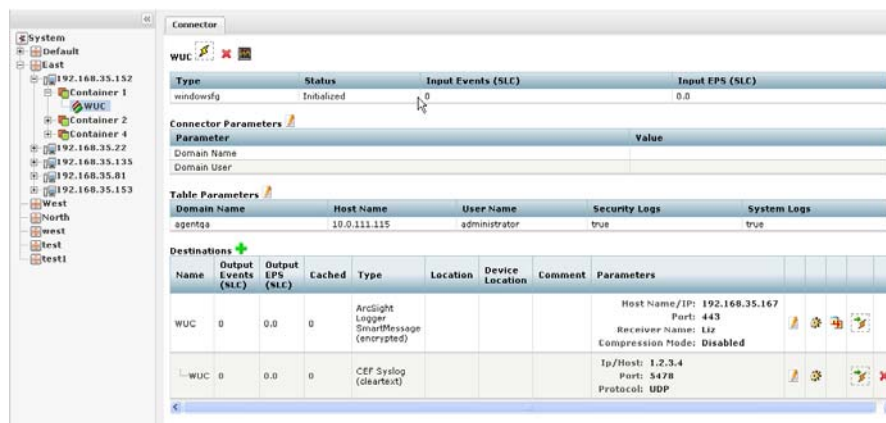
The following procedure describes how to update simple parameters for a specific connector. To update *table* parameters for a specific connector, see [“Updating Table Parameters for a Specific Connector” on page 336](#).

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.



Parameter	Value
Domain Name	
Domain User	

Domain Name	Host Name	User Name	Security Logs	System Logs
agentqa	10.0.111.115	administrator	true	true

Name	Output Events (SLC)	Output EPS (SLC)	Cached	Type	Location	Device Location	Comment	Parameters
WUC	0	0.0	0	ArcSight Logger SmartMessage (encrypted)			Host Name/IP: 192.168.35.167 Port: 443 Receiver Name: Lir Compression Mode: Disabled	
WUC	0	0.0	0	CEF Syslog (cleartext)			Ip/Host: 1.2.3.4 Port: 5478 Protocol: UDP	



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**.



Configuration parameters depend on the type of SmartConnector being configured.

- 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.



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 **Upload CSV 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	1.1.1.1.1	admin	password	TRUE	FALSE

- 5 Click **Done** when complete.

The updated table parameters display in the Table Parameters section of the Connector page.

Managing Connector Destinations

Connectors can forward events to more than one destination, such as ArcSight ESM 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
- Manage alternate configurations for a destination
- Send a command to a destination

Adding a Primary Destination to a Specific Connector

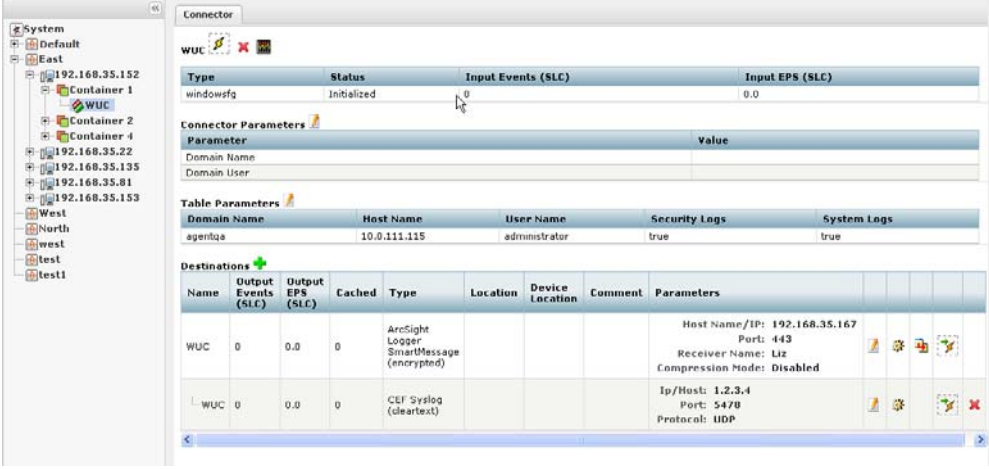
When you add a primary destination to a connector, you need to enter details for the destination, such as the destination hostname and port used.

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.



The screenshot shows the 'Connector' configuration page for a 'WUC' connector. The left sidebar shows a tree view with 'System', 'Default', 'East', and several containers. The main area displays the 'Destinations' section, which is currently expanded. It shows a table with columns: Name, Output Events (SLC), Output EPS (SLC), Cached, Type, Location, Device Location, Comment, and Parameters. Two destinations are listed: 'WUC' and 'WUC'. The 'WUC' destination is configured with Host Name/IP: 192.168.35.167, Port: 443, Receiver Name: U2, and Compression Mode: Disabled. The 'WUC' destination is configured with Ip/Host: 1.2.3.4, Port: 5470, and Protocol: UDP.



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.

- 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.

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 Update Connector Destination wizard.

- 5 Choose between 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.

- 6 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**:

- a Select the primary destination that applies to your failover.
- b Click the check box in the table header to modify all of the displayed connectors.
- c Click **Next** to update the configuration.

- 7 Click **Done** when complete.

Removing Destinations

You can remove a destination from a connector at any time.


To remove a destination from 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 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.

Editing Destination Parameters

The following procedures describe how to edit destination parameters.




You cannot change the SmartConnector type. However, you can remove the unwanted SmartConnector configuration and create a new one.

To edit destination parameters for 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 In the Destinations section, click () next to the destination you want to edit to display the Edit Destination Parameters page.

- 4 Make your changes and click **Next**.
- 5 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 [Appendix H, Destination Runtime Parameters, on page 429](#). 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 344](#).

- 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).
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 Connector Parameter Update 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

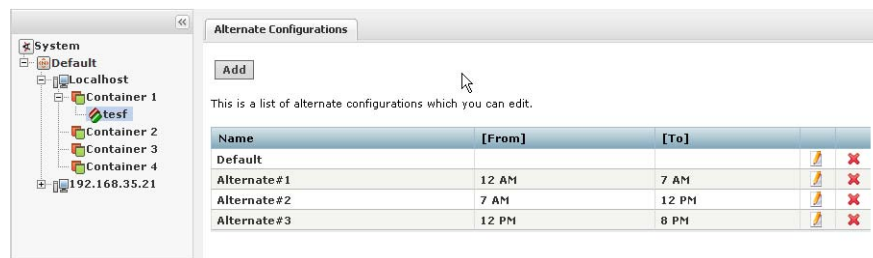
An alternate configuration is a set of runtime parameters that is used instead of the default configuration during a specified portion of every day. For example, you might want to specify different batching schemes (by severity or size) for different times of a day. You can define more than one alternate 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).

User Interface Options	Path
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](#).



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 345](#).

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 342](#).


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).
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).

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 Action column for the connector.

If you are on a specific Connector page,  is on top of the page.

- 4 Select the command from the Command Type drop-down list.
- 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).
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 SmartConnectors 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**.

Configuration Suggestions for SmartConnector Types

The following table provides configuration suggestions for different types of SmartConnectors.

SmartConnector 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 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>
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 SmartConnectors that are user-defined. FlexConnectors can be hosted on the system if they are compatible with a Linux platform. 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.

For more information, consult the *FlexConnector Developer's Guide*, available from ArcSight Customer Support.

Configuring the Check Point OPSEC NG Connector

The Check Point FW-1/VPN-1 OPSEC NG SmartConnector can operate in clear channel or sslca mode.



- This procedure is supported only for ArcSight SmartConnector 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:

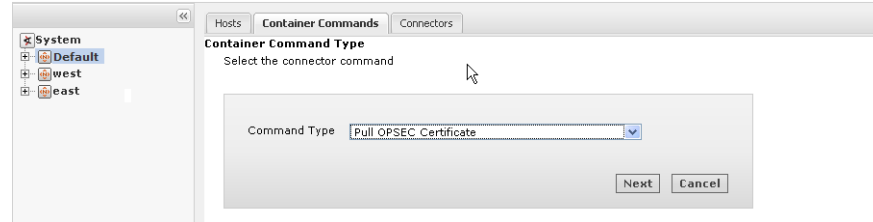
- 1 Create an OPSEC Application Object using the Check Point SmartDashboard. You need to provide these parameters when creating the application object.

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 327](#). 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 SmartConnector by following instructions described in [“Adding a Connector” on page 331](#). You need to provide the following information.

Parameters	Values to input
Type	Check Point FW-1/VPN-1 OPSEC NG
Connection Type	SSLCA
Connector Table Parameters	<p>Server IP: The IP address of the Check Point server.</p> <p>Server Port: The port on the server that listens for SSLCA connections. Use the default value 18184.</p> <p>OPSEC SIC Name: The name you noted in Step 1.</p> <p>OPSEC SSLCA File: The name you noted after pulling the certificate in Step 2.</p> <p>OPSEC Entity SIC Name: The name you noted in Step 1.</p>

- 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** check box. Click **Next**.

- 6 Continue to configure the rest of the connector. Go to [Step 6](#) in [“Adding a Connector” on page 331](#).

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.

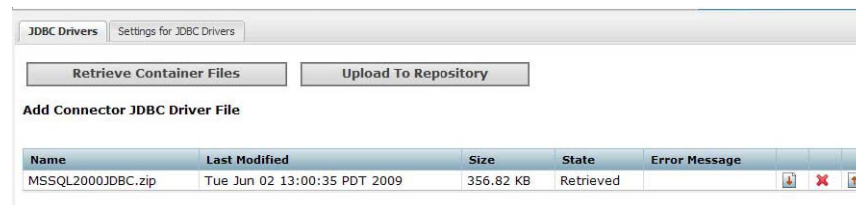
To install a JDBC Driver:

- 1 Download the MS SQL Server 2005 JDBC Driver 1.2 to a computer that can access Connector Appliance. You can download the driver from Microsoft at:

<http://www.microsoft.com/downloads/details.aspx?FamilyId=C47053EB-3B64-4794-950D-81E1EC91C1BA&displaylang=en>

- 2 Run the setup program to install the driver.
- 3 Follow the instructions in [“Uploading Files to a Repository” on page 367](#) to add the `sqljdbc.jar` file.

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 369](#).

After the driver file has been uploaded to a container, follow the instructions in [“Adding a Connector” on page 331](#) to add a connector that requires a JDBC driver.

Troubleshooting Connector Communication Issues

If your connectors are unable to communicate with an ArcSight Manager and you are:

- **Using a demo certificate on the Manager**

Enable the demo certificate on the connector. See [“Enabling or Disabling a Demo Certificate on a Container” on page 321](#) for detailed instructions.

- **Using a “self signed” certificate**

Add a CA certificate on the connector. See [“Managing Certificates on a Container” on page 321](#) for detailed instructions.

After you enable or disable FIPS mode on a container, check that the appropriate CA certificates are present in the trust store so that the connectors can validate their configured destinations successfully. For information on viewing and adding certificates, see [“Managing Certificates on a Container” on page 321](#).



If you see an error message indicating that the ESM Manager certificate is not trusted, connectors in FIPS mode are trying to communicate with an ESM Manager that is in non-FIPS mode. Enable FIPS mode on the ESM Manager.

- **Unable to resolve a hostname**

Update the Hosts file to include the required hostname. See [“Hosts” on page 253](#) for detailed instructions.

Managing Repositories in Connector Appliance

The information in this chapter is applicable only to Logger **appliance platforms with integrated Connector Appliance**.

This chapter discusses the following topics.

- [“Overview” on page 356](#)
- [“Logs Repository” on page 358](#)
- [“CA Certs Repository” on page 359](#)
- [“AUP Repository” on page 361](#)
- [“Content AUP Repository” on page 363](#)
- [“Emergency Restore” on page 364](#)
- [“User-Defined Repositories” on page 365](#)
- [“Pre-Defined Repositories” on page 370](#)

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 SmartConnector configuration and management.

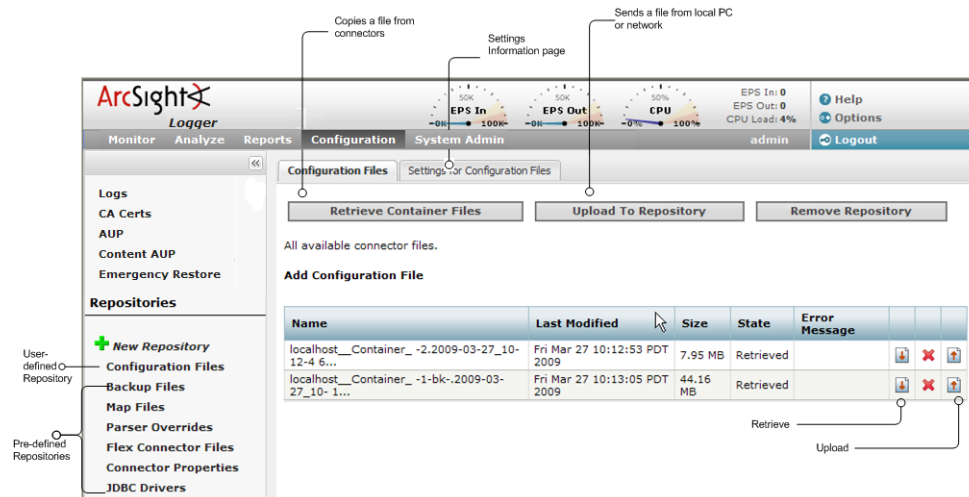



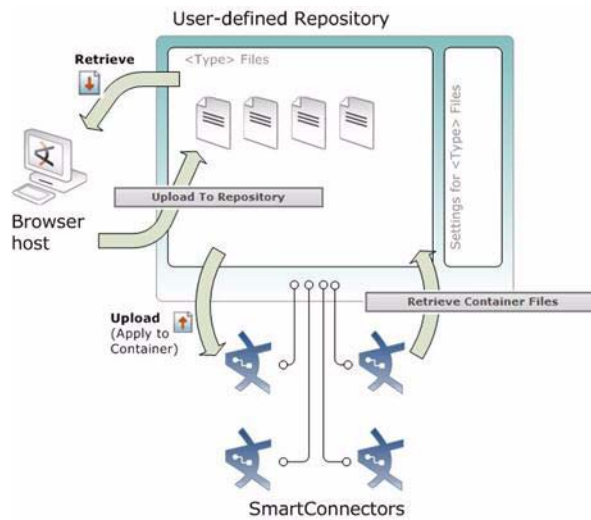
Figure 9-1 Repository Functions

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 SmartConnectors 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 SmartConnectors.



You can perform these operations using repositories:

- Manage logs in the Logs repository
- Manage CA certificates in the CA Certs repository
- Upgrade a SmartConnector using an upgrade file available in the Upgrade repository
- Apply a Content ArcSight Update Pack (AUP) on one or more SmartConnector
- Restore a container when it is damaged and irrecoverable
- Maintain centralized repositories of files for SmartConnector configuration and management

Logs Repository

When you want to view SmartConnector logs, you need to first **Load** the logs of the container that contains the SmartConnector to the Logs repository, then **Retrieve** the logs to view them.



If a container contains more than one SmartConnector, logs for all SmartConnectors are retrieved.

For information on loading, retrieving, and deleting the logs, see [“Viewing Container Logs” on page 329](#).

Uploading a File to the Logs Repository

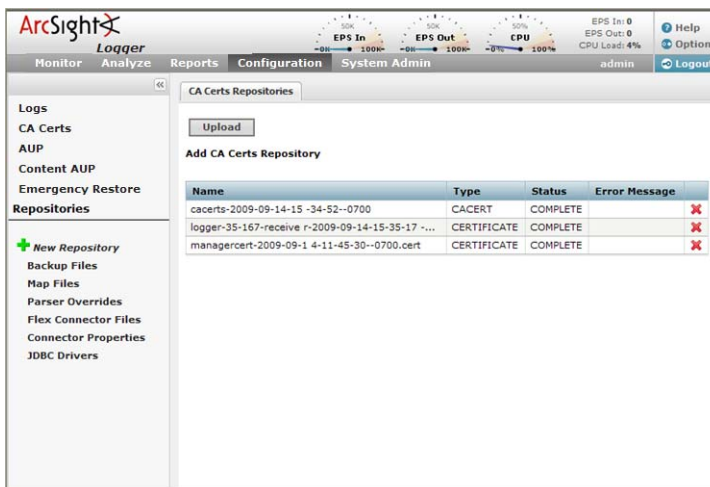
Uploading a file into the Log repository is useful for sharing annotated log or other files with other users. The file needs to be in `.zip` format.

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

SmartConnectors 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 [“Managing Certificates on a Container”](#) on page 321.



Note

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

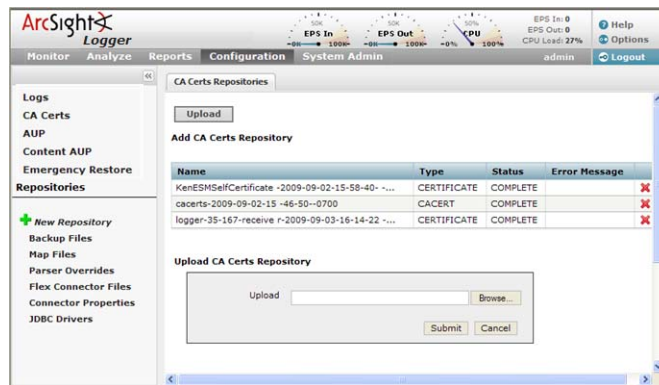
You can upload a CA Certs file or a single certificate to the CA Certs repository.



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.



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, see [“Managing Certificates on a Container” on page 321](#).

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 (✖).

AUP Repository

The Upgrade AUP repository enables you to maintain a number of SmartConnector 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 SmartConnectors 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 SmartConnectors and to upgrading a remotely-managed Connector Appliance.

To upgrade a SmartConnector 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 (see [“Upgrading a Container to a Specific Connector Version” on page 328](#)) or to a remote Connector Appliance (see [“Upgrading a Host Remotely” on page 315](#)).

Uploading an AUP Upgrade File to the Repository


To upload AUP upgrade files to the repository:

- 1 Download the upgrade AUP file for the SmartConnector or the remote Connector Appliance from the ArcSight Customer Support site at <https://support.arcsight.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 **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 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 328](#).
 - ◆ For a remotely-managed Connector Appliance upgrade, see [“Upgrading a Host Remotely” on page 315](#).

Removing a Connector Upgrade from the Repository

You can remove a SmartConnector upgrade file from the repository when you no longer need it. When you remove a SmartConnector upgrade file from the repository, it is deleted from the system.

To remove a SmartConnector upgrade from the repository:

- 1 Click **Configuration** > **Repositories** from the top-level menu bar.
- 2 Click **AUP** 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 SmartConnector 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 Customer Support site (<https://software.arcsight.com>). The AUP files are located under Content Subscription Downloads.

The ArcSight Content AUP feature enables you to apply an AUP file to applicable SmartConnector destinations that you are managing. Only the event categorization information can be applied to the SmartConnectors using this feature.

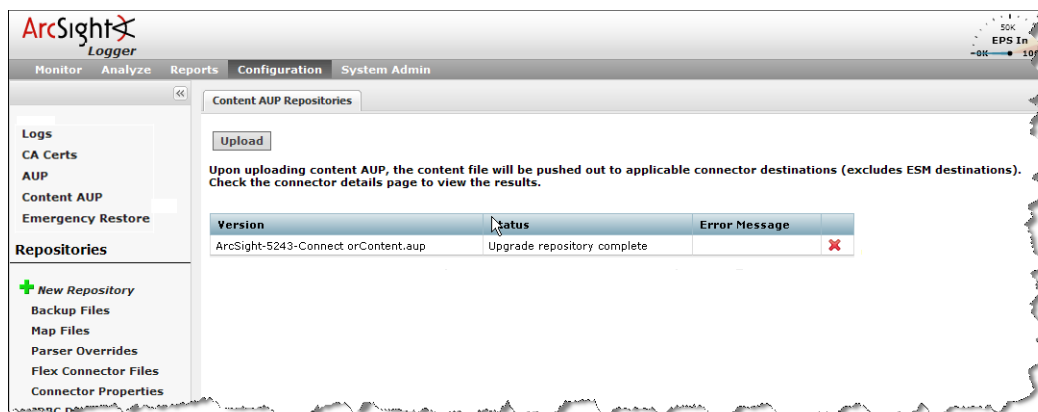
You can maintain a number of Content AUP files in the Content AUP repository. When an AUP file with a version number higher than the ones already in the repository is loaded, it is automatically pushed out to the SmartConnector destinations being managed. However, these SmartConnectors or connector destinations are skipped:

- SmartConnectors that are unavailable at the time of the AUP file push
- SmartConnectors whose current version does not fall in the range of versions that the Content AUP supports
- The ESM destination on a SmartConnector
- All destinations of a SmartConnector that have an ESM destination with the AUP Master flag set to Yes

Also, when a new SmartConnector 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 SmartConnectors



To apply a new Content AUP:

- 1 Download the new Content AUP version from ArcSight Customer Support site at <https://support.arcsight.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 SmartConnectors, or **Cancel** to quit.


You can verify the current Content AUP version on a SmartConnector by performing either of these steps:

- Run the `GetStatus` command on the SmartConnector 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 SmartConnector destination, see [“Sending a Command to a Destination” on page 346](#).
- Hover your mouse over a SmartConnector 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 SmartConnectors.

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.

Emergency Restore

The Container Restore wizard guides you through the process of restoring a modified container. This feature is supported only for SmartConnectors 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 SmartConnector is restored to the AUP version that you select.

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.

User-Defined Repositories

A *user-defined repository* is a user-named collection of settings that control upload and download of particular files from SmartConnectors to the repository. Each repository uses a specified path, relative to `$ARCSIGHT_HOME/user/agent`, for files to be uploaded or for locations to download files. ArcSight SmartConnectors use a standard directory structure, so map files, for example, are always found in `$ARCSIGHT_HOME/user/agent`, (that is, the root directory, `$ARCSIGHT_HOME`, of the SmartConnector 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 370](#)).

Files viewed in the user-defined repository can be bulk processed with specified SmartConnectors 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 typos. 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.

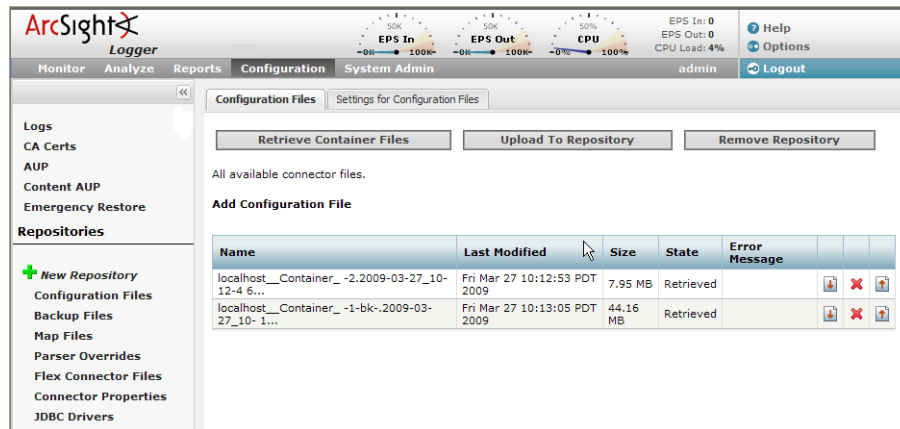
Parameter	Description
Filename Prefix	An identifying word that is included in the names of retrieved files. For example, map files are identified by Map in the file name: <code>localhost_Container_-1.Map-2009-04-06_12-22-25-607.zip</code>
Relative path (Download)	The path for download, relative to <code>\$ARCSIGHT_HOME</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.
Delete Groups	Whether to delete folders recursively in <code>\$ARCSIGHT_HOME/user/agent/map</code> directory.
Relative path (Upload)	The path for upload, relative to <code>\$ARCSIGHT_HOME/current/user/agent/flexagent/<connectorname></code>
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 SmartConnectors 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 left panel, 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.

Map Files **Settings for Map Files**

View Map Files Setting

Name: **map**

Display name: **Map Files**

Item display name: **Map File**

Recursive: ☐ No

Sort priority: **5**

Restart connector process: ☐ No

Filename prefix: **Map**

Download

Relative path: **map**

Include regular expression: **map\[0-9]+\\.properties\$**

Exclude regular expression:

Upload

Delete before upload: ☒ Yes

Delete groups: ☐ No

Relative path:

Delete relative path: **map**

Delete include regular expression: **map\[0-9]+\\.properties\$**

Delete exclude regular expression:

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.




SmartConnectors require correct properties and proper files. Applying incorrect files, including empty files or files with binary content, can prevent a SmartConnector from functioning correctly.



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 SmartConnectors, 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:** SmartConnector cloning (see [“Cloning Container Configuration” on page 371](#)).
- **Map Files:** enrich event data
- **Parser Overrides:** customize the parser
- **Flex Connector Files:** user-designed SmartConnector deployment
- **Connector Properties:** [agent.properties](#); subset of cloning
- **JDBC Drivers:** for database SmartConnectors

The following table lists the settings for each pre-defined repository.

Settings	Backup Files	Map Files	Parser Overrides	Flex Connector Files	Connector Properties
Name	backup	map	parser-overrides	flex-connectors	connector-properties
Display Name	Backup Files	Map Files	Parser Overrides	Flex Connector Files	Connector Properties
Item Display Name	Backup File	Map File	Parser Override	Flex Connector File	Connector Property File
Recursive	Checked	Un-checked	Checked	Checked	Un-checked
Sort Priority	0	5	10	15	20
Restart Connector Process	Checked	Checked	Checked	Checked	Checked
Filename Prefix	Connector-Backup	Map	Parsers	Flex-Connector	Connector-Properties
Download Relative Path		map	fcp	flexagent	
Download Include regular expression		map\[0-9]+\ .properties\$.*	.*	agent\..*
Download Exclude regular expression	(agentdata/ cwsapi_filesset_).*\$				
Delete before upload	Checked	Checked	Checked	Checked	No
Delete groups	Checked	Un-checked	Checked	Checked	Un-checked
Upload Relative Path					

Settings	Backup Files	Map Files	Parser Overrides	Flex Connector Files	Connector Properties
Delete Relative Path		map	fcf	flexagent	
Delete Include regular expression		map\[0-9]+\. properties\$.*	.*	agent\..*
Delete Exclude regular expression	(agentdata/ cwsapi_ fileset_).*\$				

Table 9-1 Pre-Defined Repository Settings

Cloning Container Configuration

Using the **Backup Files** repository, you can quickly copy a container to other containers. As a result, all SmartConnectors 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 are appended to the existing contents of the destination container.



Do not clone older, software-based SmartConnectors (such as build 4.0.8.4964) to containers with newer SmartConnector builds (such as 4.0.8.4976 or later).

Cloning a SmartConnector using the Backup repository only works if the SmartConnector 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 New Repository 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 369](#) 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 369](#) to upload the backup file to one or more containers.

The destination containers are unavailable while the backup file is applied and the SmartConnectors are restarted. After a minute or so, you can check the Connectors tab to see if the operation was successful.

Appendix A

Common Event Format

Common Event Format (CEF) is an industry standard for the interoperability of event- or log-generating devices. The myriad of formats used for event reporting, especially in the security world, greatly complicates integration. Each vendor has its own format for reporting event information, but these formats often lack key information necessary to integrate the events from their devices.

The CEF standard aims to improve the interoperability of infrastructure devices by better aligning the logging output from participating technology vendors. Vendors implementing the CEF standard log events in a format that is both useful, and more importantly, parse-able by ArcSight or any vendor following the standard. Further, this standard assures that an event and its semantics contain all necessary information.

Common Exchange Format

This specification defines a simple event format that can be readily adopted by vendors of both security and non-security devices. This format is intended to contain the most relevant information and make it easy for event consumers to parse and use events.

To simplify integration, we use syslog as a transport mechanism. This applies a common prefix to each message, containing the date and hostname:

```
Jan 18 11:07:53 zurich message
```

If an event producer is unable to write syslog messages, it is still possible to write the events to a file. In this case, omit the syslog header and start the message with the format defined below.

It is important to note that this part of the message need not be explicitly generated by the event producer. The remainder of the message is formatted using a common prefix composed of fields delimited by a bar ("|") character. The prefix is mandatory and all specified fields need to be present. Additional fields are specified in the Extension. The format is:

```
CEF:Version|Device Vendor|Device Product|Device Version|Device  
Event Class ID|Name|Severity|Extension
```

The *Extension* part of the message is a placeholder for additional fields. Those fields are documented in the Event Dictionary below and are logged as key-value pairs.

Here are definitions for the prefix fields:

Version is an integer and identifies the version of the CEF format. Event consumers use this information to determine what the following fields represent. Currently only version 0

(zero) is established in the above format. Experience may show that other fields need to be added to the "prefix" and therefore require a version number change. Adding new formats is handled through the standards body.

Device Vendor, **Device Product** and **Device Version** are strings that uniquely identify the type of sending device. No two products may use the same device-vendor and deviceproduct pair. There is no central authority managing these pairs. Event producers have to ensure that they assign unique name pairs.

Device Event Class ID is a unique identifier per event-type. This can be a string or an integer. Device Event Class ID identifies the type of event reported. In the intrusion detection system (IDS) world, each signature or rule that detects certain activity has a unique identifier assigned. This is a requirement for other types of devices as well, and helps correlation engines deal with the events.

Name is a string representing a human-readable and understandable description of the event. The event name should not contain information that is specifically mentioned in other fields. For example: "Port scan from 10.0.0.1 targeting 20.1.1.1" is not a good event name. It should be: "Port scan." The other information is redundant and can be picked up from the other fields.

Severity is an integer and reflects the importance of the event. Only numbers from 0 to 10 are allowed, where 10 indicates the most important event.

Extension is a collection of key-value pairs. The keys are part of a predefined set. The standard allows for including additional keys as outlined later. An event can contain any number of key-value pairs in any order, separated by spaces (" "). If a field contains a space, such as a file name, this is okay and can be logged in exactly that manner. For example: fileName=c:\Program Files\ArcSight is a valid token.

Here is a sample message to illustrate appearance:

```
Sep 19 08:26:10 zurich CEF:0|security|threatmanager|1.0|100|worm
successfully stopped|10|src=10.0.0.1 dst=2.1.2.2 spt=1232
```

Here are further details about character encoding:

The entire message has to be **UTF-8** encoded.

If a pipe (|) is used in the prefix, it has to be escaped with a backslash (\). But note that pipes in the extension do not need escaping. Here is an example message:

```
Sep 19 08:26:10 zurich
CEF:0|security|threatmanager|1.0|100|detected a \ in
message|10|src=10.0.0.1 act=blocked a | dst=1.1.1.1
```

If a backslash (\) is used in the prefix, it has to be escaped with another backslash (\). Again, note that backslashes in the extension do not need escaping. Here is an example:

```
Sep 19 08:26:10 zurich
CEF:0|security|threatmanager|1.0|100|detected a \\ in
packet|10|src=10.0.0.1 action=blocked a \ dst=1.1.1.1
```

If an equal sign (=) is used in the extensions, it has to be escaped with a backslash (\). Equal signs in the prefix need no escaping. For example:

```
Sep 19 08:26:10 zurich
CEF:0|security|threatmanager|1.0|100|detected a = in
message|10|src=10.0.0.1 action=blocked a \= dst=1.1.1.1
```

Multi-line fields can be sent by Common Event Format (CEF) by encoding the newline character as \n or \r. Note that multiple lines are only allowed in the value part of the extensions. See this example:

```
Sep 19 08:26:10 zurich
CEF:0|security|threatmanager|1.0|100|Detected a threat. No action
needed.|10|src=10.0.0.1 message=Detected a threat.\nNo action
needed.
```

Common Extension Dictionary

The following table contains predefined keys that establish usages for both event producers and consumers. The standard allows for defining additional keys, with the understanding that those fields may not be interpreted by other event consumers.

The table below contains key names as well as the full name for each key. The key name is the one that is required in events.

Key Name	Full Name	Data Type	Meaning
act	deviceAction	String	Action mentioned in the event.
app	applicationProtocol	String	Application level protocol. Example values include: HTTP, HTTPS, SSHv2, Telnet, POP, IMAP, IMAPS.
in	bytesIn	Integer	Number of bytes transferred inbound. Inbound relative to the source to destination relationship, meaning that data was flowing from source to destination.
out	bytesOut	Integer	Number of bytes transferred outbound. Outbound relative to the source to destination relationship, meaning that data was flowing from destination to source.
dst	destinationAddress	IPv4 Address	Identifies destination that the event refers to in an IP network. The format is an IPv4 address, such as "192.168.10.1".
dhost	destinationHostName	FQDN	Identifies the destination that the event refers to in an IP network. The format is a fully qualified domain name (FQDN) associated with the destination node, such as "zurich.domain.com".
dmac	destinationMacAddress	String	Six colon-separated hexadecimal numbers, such as "00:0D:60:AF:1B:61"
dntdom	destinationNtDomain	String	The Windows domain name of the destination address.

Key Name	Full Name	Data Type	Meaning
dpt	destinationPort	Integer	The valid port numbers are between 0 and 65535.
dproc	destination ProcessName	String	The name of the process which is the event's destination, such as "telnetd" or "sshd".
duid	destinationUserId	String	Identifies the destination user by ID. For example, in Unix, the root user is generally associated with ID 0.
dpriv	destination UserPrivileges	String	The allowed values are: "Administrator", "User", and "Guest". This identifies the destination user's privileges. In Unix, for example, activity executed on the root user would be identified with destinationUserPrivileges of "Administrator".
duser	destinationUserName	String	Identifies the destination user by name. This is the user associated with the event's destination. E-mail addresses are also mapped into the UserName fields. The recipient is a candidate to put into destinationUserName.
end	endTime	TimeStamp	The time at which the activity related to the event ended. The format is MMM dd yyyy HH:mm:ss or milliseconds since epoch (January 1, 1970). An example would be reporting the end of a session.
fname	fileName	String	Name of the file.
fsize	fileSize	Integer	Size of the file.
msg	message	String	An arbitrary message giving more details about the event. Multi-line entries can be produced by using '\n' as the newline separator.
rt	receiptTime	TimeStamp	The time at which the event related to the activity was received. The format is MMM dd yyyy HH:mm:ss or milliseconds since epoch (January 1, 1970).
request	requestURL	String	In the case of an HTTP request, this field contains the URL accessed. The URL should contain the protocol as well, such as "http://www.security.com".
src	sourceAddress	IPv4 Address	Identifies the source that the event refers to in an IP network. The format is an IPv4 address, such as "192.168.10.1".

Key Name	Full Name	Data Type	Meaning
shost	sourceHostName	FQDN	Identifies the source that the event refers to in an IP network. The format is a fully qualified domain name (FQDN) associated with the source node, such as "zurich.domain.com".
smac	sourceMacAddress	String	Six colon-separated hexadecimal numbers, such as "00:0D:60:AF:1B:61"
sntdom	sourceNtDomain	String	The Windows domain name of the source address.
spt	sourcePort	Integer	The valid port numbers are between 0 and 65535.
suid	sourceUserId	String	Identifies the source user by ID. This is the user associated with the source of the event. For example, in Unix, the root user is generally associated with ID 0.
spriv	sourceUserPrivileges	String	<p>The allowed values are: "Administrator", "User", and "Guest". This identifies the source user's privileges. In Unix, for example, activity executed on the root user would be identified with sourceUserPrivileges of "Administrator".</p> <p>This is an idealized and simplified view of privileges and can be extended in the future.</p>
suser	sourceUserName	String	Identifies the source user by name. E-mail addresses are also mapped into the UserName fields. The sender is a candidate to put into sourceUserName.
start	startTime	TimeStamp	The time when the activity the event referred to started. The format is MMM dd yyyy HH:mm:ss or milliseconds since epoch (January 1, 1970).
proto	transportProtocol	String	Identifies the Layer-4 protocol used. The possible values are protocol names such as TCP or UDP.

Appendix B

Regular Expressions

Regular String Search expressions (Perl Regex) are used to compose Logger filters. The following describes the syntax of regular expressions in Perl.

Regex Overview

A regular expression is a string of characters which tells the searcher which string (or strings) you are looking for. The following explains the format of regular expressions in detail. If you are familiar with Perl, you already know the syntax. If you are familiar with Unix, you should know that there are subtle differences between Perl's regular expressions and Unix' regular expressions.

The following is a list of the characters used in a regular expression syntax and their meaning.

Predefined Character Classes:

. Any character (may or may not match line terminators)

\d A digit: [0-9]

\D A non-digit: [^0-9]

\s A whitespace character: [\t\n\x0B\f\r]

\S A non-whitespace character: [^\s]

\w A word character: [a-zA-Z_0-9]

\W A non-word character: [^\w]

Standard quantifiers include '?' (zero or one), '+' (one or more), '*' (zero or more), {n} (exactly n), and {n,m} (at least n, but no more than m).

Boundary Matchers:

^ The beginning of a line

\$ The end of a line

\b A word boundary

\B A non-word boundary

\A The beginning of an event

\G The end of the previous match

\Z The end of an event

Simple Regular Expressions

In its simplest form, a regular expression is just a word or phrase to search for. For example,

`gauss`

would match any event with the string "gauss" in it, or which mentioned the word "gauss." Thus, events with "gauss", "gaussian" or "degauss" would all be matched, as would an event containing the phrases "de-gauss the monitor" or "gaussian elimination." Here are some more examples:

`carbon`

Finds any event with the string "carbon" in it, or which mentions carbon (or carbonization or hydrocarbons or carbon-based life forms) in the event.

`hydro`

Finds any event with the string "hydro" in it. Events with "hydro", "hydrogen" or "hydrodynamics" are found, as well as events containing the words "hydroplane" or "hydroelectric".

`oxy`

Finds any event with the string "oxy" in it. This could be used to find event on oxygen, boxy houses or oxymorons.

`top ten`

Note that spaces may be part of the regular expression. The above expression could be used to find top ten lists. (Note that they would also find articles on how to stop tension.)

Metacharacters

Some characters have a special meaning to the searcher. These characters are called **metacharacters**. Although they may seem confusing at first, they add a great deal of flexibility and convenience to the searcher.

The **period** (.) is a commonly used metacharacter. It matches exactly one character, regardless of what the character is. For example, the regular expression:

`2,.-Dimethylbutane`

will match "2,2-Dimethylbutane" and "2,3-Dimethylbutane". Note that the period matches **exactly one** character-- it will not match a string of characters, nor will it match the null string. Thus, "2,200-Dimethylbutane" and "2,-Dimethylbutane" will **not** be matched by the above regular expression.

But what if you wanted to search for a string containing a period? For example, suppose we wished to search for references to pi. The following regular expression would **not** work:

`3.14 (THIS IS WRONG!)`

This would indeed match "3.14", but it would also match "3514", "3f14", or even "3+14". In short, any string of the form "3x14", where x is any character, would be matched by the regular expression above.

To get around this, we introduce a second metacharacter, the **backslash** (\). The backslash can be used to indicate that the character immediately to its right is to be taken literally. Thus, to search for the string "3.14", we would use:

```
3\.14 (This will work.)
```

This is called "quoting". We would say that the period in the regular expression above has been quoted. In general, whenever the backslash is placed before a metacharacter, the searcher treats the metacharacter literally rather than invoking its special meaning.

(Unfortunately, the backslash is used for other things besides quoting metacharacters. Many "normal" characters take on special meanings when preceded by a backslash. The rule of thumb is, quoting a metacharacter turns it into a normal character, and quoting a normal character **may** turn it into a metacharacter.)

Let's look at some more common metacharacters. We consider first the **question mark** (?). The question mark indicates that the character immediately preceding it either zero times or one time. Thus

```
m?ethane
```

would match either "ethane" or "methane". Similarly,

```
comm?a
```

would match either "coma" or "comma".

Another metacharacter is the **star** (*). This indicates that the character immediately to its left may be repeated any number of times, including zero. Thus

```
ab*c
```

would match "ac", "abc", "abbc", "abbbc", "abbbbbbbbc", and any string that starts with an "a", is followed by a sequence of "b"s, and ends with a "c".

The **plus** (+) metacharacter indicates that the character immediately preceding it may be repeated one or more times. It is just like the star metacharacter, except it doesn't match the null string. Thus

```
ab+c
```

would **not** match "ac", but it **would** match "abc", "abbc", "abbbc", "abbbbbbbbc" and so on.

Metacharacters may be combined. A common combination includes the period and star metacharacters, with the star immediately following the period. This is used to match an arbitrary string of any length, including the null string. For example:

```
cyclo.*ane
```

would match "cyclodecane", "cyclohexane" and even "cyclones drive me insane." Any string that starts with "cyclo", is followed by an arbitrary string, and ends with "ane" will be matched. Note that the null string will be matched by the period-star pair; thus, "cycloane" would be matched by the above expression.

If you wanted to search for articles on cyclodecane and cyclohexane, but didn't want to match articles about how cyclones drive one insane, you could string together three periods, as follows:

`cyclo...ane`

This would match "cyclodecane" and "cyclohexane", but would not match "cyclones drive me insane." Only strings eleven characters long which start with "cyclo" and end with "ane" will be matched. (Note that "cyclopentane" would not be matched, however, since cyclopentane has twelve characters, not eleven.)

Here are some more examples. These involve the backslash. Note that the placement of backslash is important.

`a\.*z`

Matches any string starting with "a", followed by a series of periods (including the "series" of length zero), and terminated by "z". Thus, "az", "a.z", "a..z", "a...z" and so forth are all matched.

`a.*z`

(Note that the backslash and period are reversed in this regular expression.)

Matches any string starting with an "a", followed by one arbitrary character, and terminated with "*z". Thus, "ag*z", "a5*z" and "a@*z" are all matched. Only strings of length four, where the first character is "a", the third "*", and the fourth "z", are matched.

`a\++z`

Matches any string starting with "a", followed by a series of plus signs, and terminated by "z". There must be at least one plus sign between the "a" and the "z". Thus, "az" is **not** matched, but "a+z", "a++z", "a+++z", etc. will be matched.

`a\+\+z`

Matches only the string "a++z".

`a+\+z`

Matches any string starting with a series of "a"'s, followed by a single plus sign and ending with a "z". There must be at least one "a" at the start of the string. Thus "a+z", "aa+z", "aaa+z" and so on will match, but "+z" will not.

`a.?e`

Matches "ace", "ale", "axe" and any other three-character string beginning with "a" and ending with "e"; will also match "ae".

`a\.?e`

Matches "ae" and "a.e". No other string is matched.

`a.\?e`

Matches any four-character string starting with "a" and ending with "?e". Thus, "ad?e", "a1?e" and "a%?e" will all be matched.

`a\.\?e`

Matches only "a.?e" and nothing else.

Earlier it was mentioned that the backslash can turn ordinary characters into metacharacters, as well as the other way around. One such use of this is the **digit**

metacharacter, which is invoked by following a backslash with a lower-case "d", like this: "\d". The "d" **must be lower case**, for reasons explained later. The digit metacharacter matches exactly one digit; that is, exactly one occurrence of "0", "1", "2", "3", "4", "5", "6", "7", "8" or "9". For example, the regular expression:

```
2,\d-Dimethylbutane
```

would match "2,2-Dimethylbutane", "2,3-Dimethylbutane" and so forth. Similarly,

```
1\.\d\d\d\d\d\d
```

would match any six-digit floating-point number from 1.00000 to 1.99999 inclusive. We could combine the digit metacharacter with other metacharacters; for instance,

```
a\d+z
```

matches any string starting with "a", followed by a string of numbers, followed by a "z". (Note that the plus is used, and thus "az" is not matched.)

The letter "d" in the string "\d" must be lower-case. This is because there is another metacharacter, the **non-digit** metacharacter, which uses the uppercase "D". The non-digit metacharacter looks like "\D" and matches any character **except** a digit. Thus,

```
a\Dz
```

would match "abz", "aTz" or "a%z", but would **not** match "a2z", "a5z" or "a9z". Similarly,

```
\D+
```

Matches any non-null string which contains **no** numeric characters.

Notice that in changing the "d" from lower-case to upper-case, we have reversed the meaning of the digit metacharacter. This holds true for most other metacharacters of the format backslash-letter.

There are three other metacharacters in the backslash-letter format. The first is the **word** metacharacter, which matches exactly one letter, one number, or the underscore character (_). It is written as "\w". Its opposite, "\W", matches any one character **except** a letter, a number or the underscore. Thus,

```
a\wz
```

would match "abz", "aTz", "a5z", "a_z", or any three-character string starting with "a", ending with "z", and whose second character was either a letter (upper- or lower-case), a number, or the underscore. Similarly,

```
a\Wz
```

would not match "abz", "aTz", "a5z", or "a_z". It **would** match "a%z", "a{z", "a?z" or any three-character string starting with "a" and ending with "z" and whose second character was not a letter, number, or underscore. (This means the second character must either be a symbol or a whitespace character.)

The **whitespace** metacharacter matches exactly one character of whitespace. (Whitespace is defined as spaces, tabs, newlines, or any character which would not use ink if printed on a printer.) The whitespace metacharacter looks like this: "\s". Its opposite, which matches any character that is **not** whitespace, looks like this: "\S". Thus,

```
a\s z
```

would match any three-character string starting with "a" and ending with "z" and whose second character was a space, tab, or newline. Likewise,

```
a\Sz
```

would match any three-character string starting with "a" and ending with "z" whose second character was **not** a space, tab or newline. (Thus, the second character could be a letter, number or symbol.)

The **word boundary** metacharacter matches the boundaries of words; that is, it matches whitespace, punctuation and the very beginning and end of the text. It looks like "\b". It's opposite searches for a character that is **not** a word boundary. Thus:

```
\bcomput
```

will match "computer" or "computing", but not "supercomputer" since there is no spaces or punctuation between "super" and "computer". Similarly,

```
\Bcomput
```

will **not** match "computer" or "computing", unless it is part of a bigger word such as "supercomputer" or "recomputing".

Note that the underscore (_) is considered a "word" character. Thus,

```
super\bcomputer
```

will **not** match "super_computer".

There is one other metacharacter starting with a backslash, the **octal** metacharacter. The octal metacharacter looks like this: "\nnn", where "n" is a number from zero to seven. This is used for specifying control characters that have no typed equivalent. For example,

```
\007
```

would find all events with an embedded ASCII "bell" character. (The bell is specified by an ASCII value of 7.) You will rarely need to use the octal metacharacter.

There are three other metacharacters that may be of use. The first is the **braces** metacharacter. This metacharacter follows a normal character and contains two number separated by a comma (,) and surrounded by braces ({}). It is like the star metacharacter, except the length of the string it matches must be within the minimum and maximum length specified by the two numbers in braces. Thus,

```
ab{3,5}c
```

will match "abbbc", "abbbbc" or "abbbbbc". No other string is matched. Likewise,

```
.{3,5}pentane
```

will match "cyclopentane", "isopentane" or "neopentane", but not "n-pentane", since "n-" is only two characters long.

The alternative metacharacter is represented by a vertical bar (|). It indicates an either/or behavior by separating two or more possible choices. For example:

```
isopentane|cyclopentane
```

will match any event containing the strings "isopentane" or "cyclopentane" or both. However, it will not match "pentane" or "n-pentane" or "neopentane." The last

metacharacter is the **brackets** metacharacter. The bracket metacharacter matches one occurrence of any character inside the brackets ([]). For example,

```
\s[cmt]an\s
```

will match "can", "man" and "tan", but not "ban", "fan" or "pan". Similarly,

```
2,[23]-dimethylbutane
```

will match "2,2-dimethylbutane" or "2,3-dimethylbutane", but not "2,4-dimethylbutane", "2,23-dimethylbutane" or "2,-dimethylbutane". Ranges of characters can be used by using the dash (-) within the brackets. For example,

```
a[a-d]z
```

will match "aaz", "abz", "acz" or "adz", and nothing else. Likewise,

```
textfile0[3-5]
```

will match "textfile03", "textfile04", or "textfile05" and nothing else.

If you wish to include a dash within brackets as one of the characters to match, instead of to denote a range, put the dash immediately before the right bracket. Thus:

```
a[1234-]z
```

and

```
a[1-4-]z
```

both do the same thing. They both match "a1z", "a2z", "a3z", "a4z" or "a-z", and nothing else.

The bracket metacharacter can also be inverted by placing a caret (^) immediately after the left bracket. Thus,

```
textfile0[^02468]
```

matches any ten-character string starting with "textfile0" and ending with anything except an even number. Inversion and ranges can be combined, so that

```
\W[^f-h]ood\W
```

matches any four letter word ending in "ood" **except** for "food", "good" or "hood". (Thus "mood" and "wood" would both be matched.)

Note that within brackets, ordinary quoting rules do not apply and other metacharacters are not available. The only characters that can be quoted in brackets are "[", "]", and "\". Thus,

```
[\\[\]]abc
```

matches any four letter string ending with "abc" and starting with "[", "]", or "\".

Forbidden Characters

Because of the way the searcher works, the following metacharacters should not be used, even though they are valid Perl metacharacters. They are:

- \$ (allowed within brackets)

- `\n`
- `\r`
- `\t`
- `\f`
- `\b`
- `()` (allowed within brackets. Note that if you wish to search for parentheses within text outside of brackets, you should quote the parentheses.)
- `\1, \2 ... \9`
- `\B`
- `:`
- `!`

Things To Remember

Here are some other things you should know about regular expressions.

Because regular expressions can be complex, it can be more work mastering a search than just sifting through a long list of matches (unless you're already familiar with regular expressions).

The search is case insensitive; thus

`mopac`

and

`Mopac`

and

`MOPAC`

all search for the same set of strings. Each will match "mopac", "MOPAC", "Mopac", "mopaC", "MoPaC", "mOpAc" and so forth. Thus you need not worry about capitalization. (Note, however, that metacharacters must still have the proper case. This is especially important for metacharacters whose case determines whether their meaning is reversed or not.)

Outside of the brackets metacharacter, you must quote parentheses, brackets and braces to get the searcher to take them literally.

Appendix C

Using the Rex Operator

This appendix describes the `rex` search operator in detail.

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.

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 "(?<src_ip>[^\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
```

Syntax of the rex Operator

```
| rex "text1(?field_1>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.

field_1—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 upto *text2* that matches the specified *regex* (regular expression) and assign TO *field_1*.

Notes:

- *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 is found in the event. In this case, a space.

- 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 "(?<src_ip>[^\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" | rex field=src_ip
"(?<net_id>\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 "(?<field_1>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 “src_ip” is the name assigned to the capture.

```
failed | rex "(?<src_ip>[^\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
```

Once named, use “src_ip” for further processing as follows:

```
failed | rex "(?<src_ip>[^\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
| top src_ip
```

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 70](#). 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 "\[client(?<clientip_1>[^\]]*)" "
```

Samples of rex Expressions

This sections 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 Regexp 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
CEF:0|ArcSight|Logger|4.5.0.4836.0|eps:100|Logger Internal Event|1| cat=/Monitor/Receiver/A11/EP5 cs2=SinceLastMonitorEvent cnt=1 dvc=192.168.36.3 cs1=
2010/07/01 13:46:00 PDT      unknown      Local      ArcSight      Logger      4.5.0.4836.0
CEF:0|ArcSight|Logger|4.5.0.4836.0|eps:100|Logger Internal Event|1| cat=/Monitor/Forwarder/A11/EP5 cs2=SinceLastMonitorEvent cnt=1 dvc=192.168.36.3 cs1=
```

- 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>[^\s]*)" | where message is not null |
rex "dvc=(?<msgip>[^\s]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" |
where msgip is not null
```

Note: 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 03:04:24 192.168.20.114 192.168.20.112 192.168.20.112 C007:4D28:EvilPackets:Line 16:"New Group","My 80280150","11/10/2005 11:02:05.000"
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"
```

- Extract the first two IP addresses from an event and assign them to two different fields, IP1 and IP2.

```
| rex "(?<IP1>[^\s]\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>[^\s]\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.

Note: 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
"((?<IP1>[^\s]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})(?<IP2>[^\s]\d{1,3}\.
\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}))"
```

Note: 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>[^\s]\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>[^\s]\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
```

Note: 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 "(?<src_ip>[^\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}])" | rex field=src_ip
"(?<net_id>\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 - _count
```

Notes pertinent to this example:

- ◆ 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
RAW	Feb 25 14:03:24 beach login(pam_unix)[123]: session closed for user root
2	2010/04/06 22:11:46 CDT 10.0.10.222 [Raw_Syslog] Local
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
3	2010/04/06 22:11:46 CDT 10.0.10.222 [Raw_Syslog] Local
RAW	Feb 25 14:03:24 beach sshd(pam_unix)[123]: session closed for user p4admin
4	2010/04/06 22:11:46 CDT 10.0.10.222 [Raw_Syslog] Local
RAW	Feb 25 14:03:24 beach sshd(pam_unix)[123]: session opened for user 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 D

Restoring Factory Settings

ArcSight Logger can be restored to its original factory settings using built-in Acronis True Image software.



Restoring Logger to factory settings will irrevocably delete all event data and configuration settings.

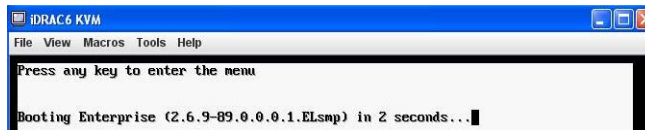


The screens shown here are examples only. Your Logger partitions might vary, and the overall capacity might be different.

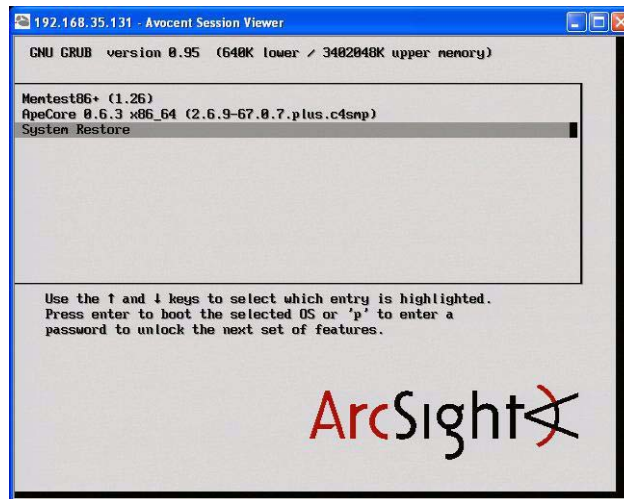
To restore factory settings

To restore Logger to its original factory settings, perform these steps:

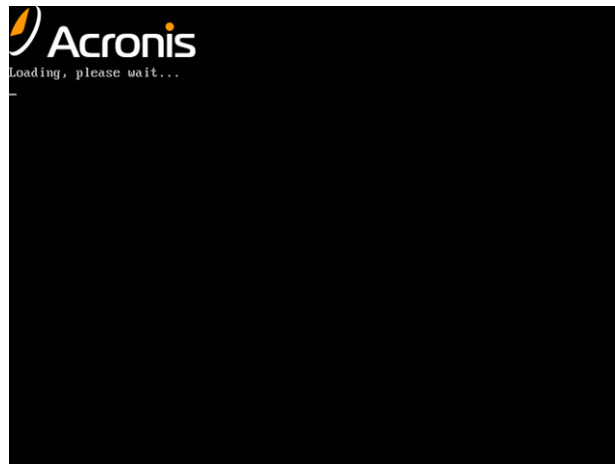
- 1 Attach a keyboard, monitor, and mouse directly to the ArcSight Logger appliance or, if your Logger is configured for DRAC, use that functionality to access the Logger Console. See [“Configure DRAC for Remotely Accessing a Logger Appliance” on page 19](#) for information about DRAC.
- 2 Reboot Logger from the web interface by clicking the **System Admin** tab, the **System Reboot** command in the sub-menu, and the **Start Reboot Now** button.
- 3 Once the following screen is displayed, press any key on your keyboard. This screen is displayed for a very short time, therefore, make sure you press a key on your keyboard quickly; otherwise, the Logger will continue to boot normally.



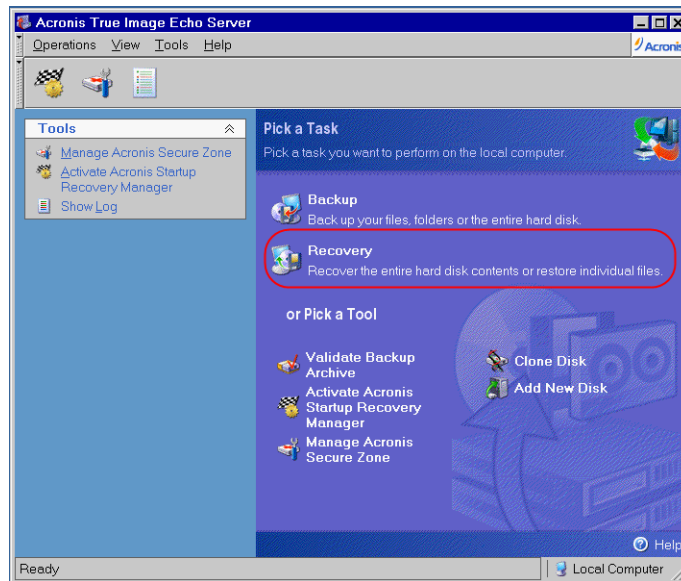
- 4 A screen similar to the following is displayed on the attached monitor. Use the mouse or arrow keys to select System Restore and press the **Enter** key on your keyboard.



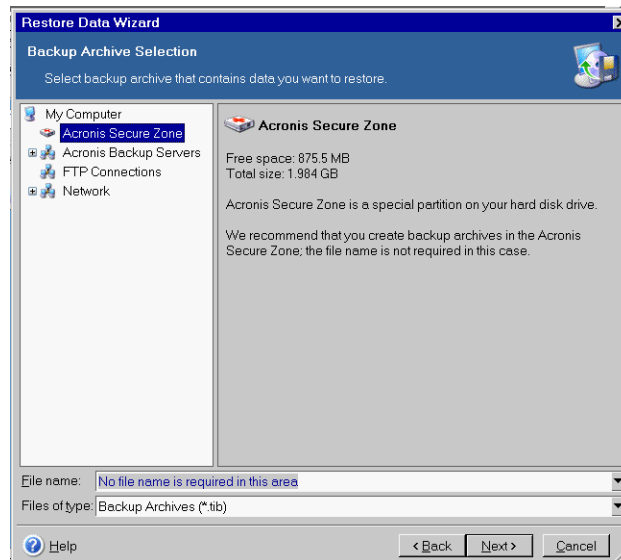
- 5 Select **Acronis True Image Server** to continue.



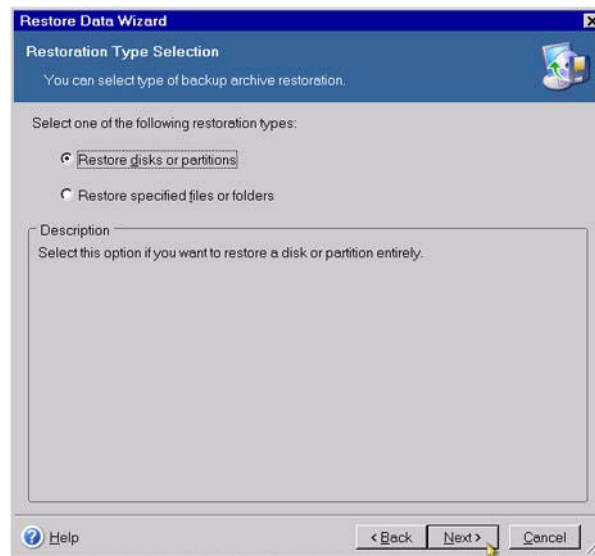
- 6 On the Pick a Task list, as shown in the following figure, choose **Recovery**. On the next page (Welcome to the Restore Data Wizard), click **Next** to continue.



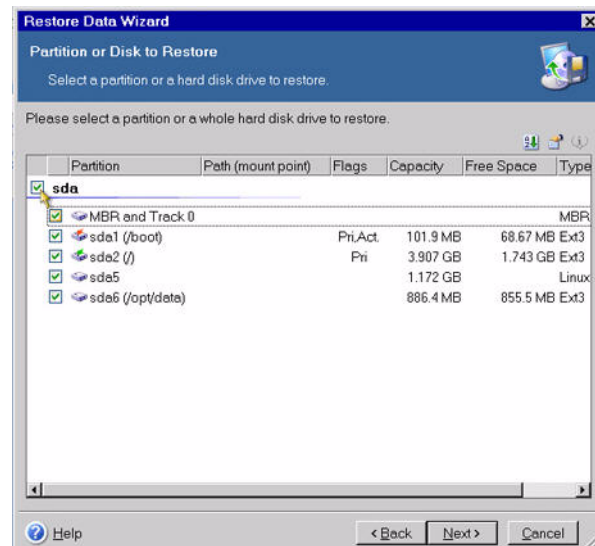
- 7 Select the Acronis Secure Zone, as shown in the following figure, and click **Next**. You will have a chance to review the choices you make on this page and the wizard pages that follow before initiating the restore process.



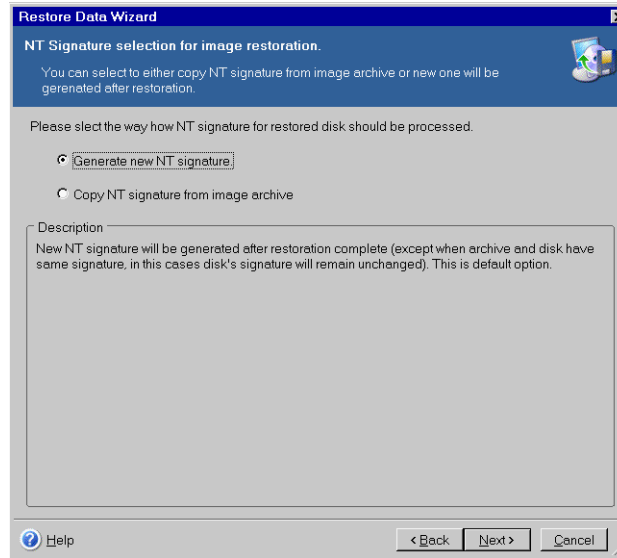
- 8 Select **Restore disks or partitions** and click **Next**. Only choose other options if specifically directed to do so by ArcSight Customer Support.



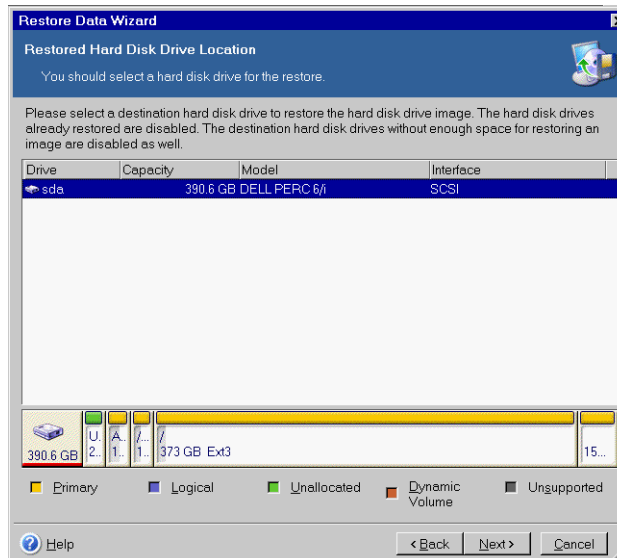
- 9 Select the entire drive, labeled 'sda' in the following figure. Click **Next** to continue.



- 10 Select the way in which the NT signature for the restored disk should be processed and click **Next**.



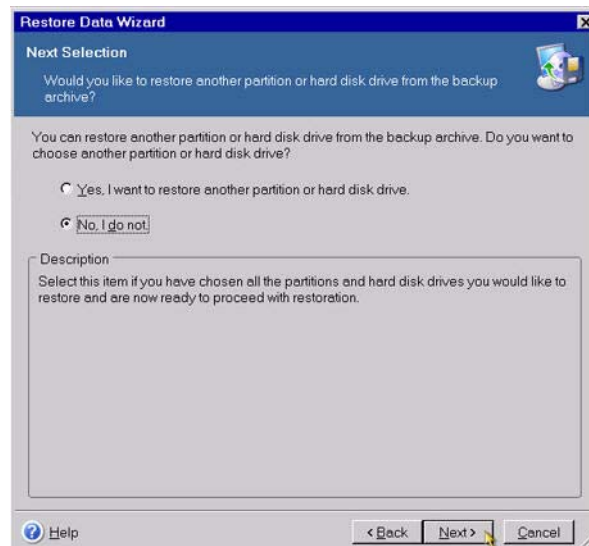
- 11 Choose the drive to restore ('sda') and click **Next**.



- 12** Select, "Yes, I want to delete all partitions on the destination hard disk drive before restoring", as shown in the following figure.

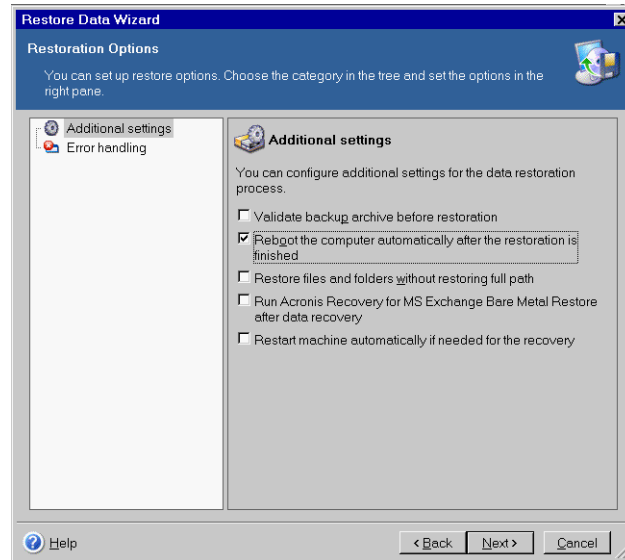


- 13** Because there are no other partitions or disks to restore, choose "No, I do not," on the Next Selection page of the wizard. Click **Next**.

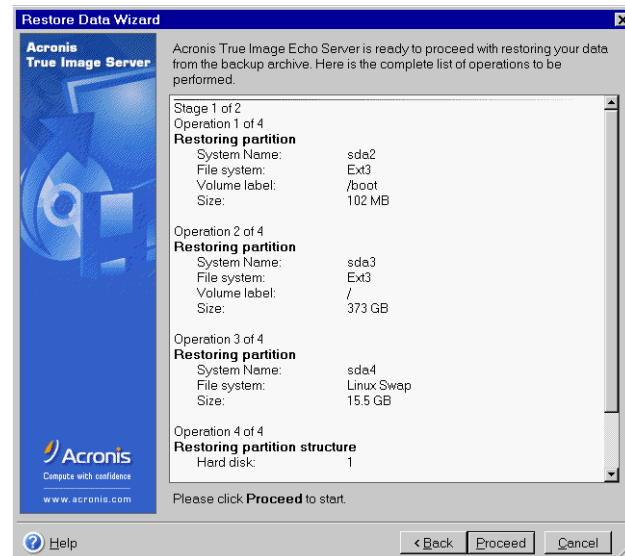


- 14** Validating the archive before restoring is optional. Check the box to validate the archive or leave it unchecked to skip this step. Check the box labeled "Reboot the

computer automatically after the restoration is finished” to automatically reboot. Click **Next**.



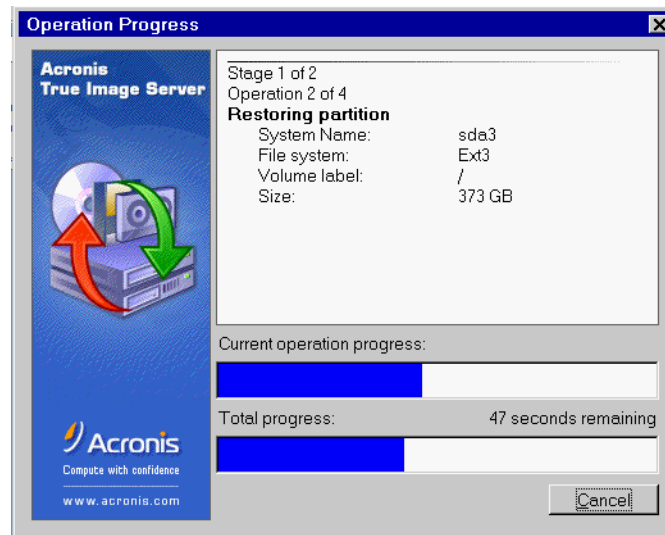
- 15 Review the checklist of operations to be performed, as shown in the following figure, and click **Proceed** to begin the restore process, or click **Back** to revisit previous wizard pages.



Caution

Do not interrupt or power-down Logger during the restore process. Interrupting the restore process may force the system into a state from which it cannot be recovered.

- 16** The progress bars (shown in the following figure) display the status of the current and total operations. When the restoration is complete, an alert is displayed that says "Data was successfully restored." Click **OK**.



If you specified automatic reboot in Step 13, Logger will reboot when the restore is complete. Otherwise, reboot manually.

Appendix E

Logger Audit Events

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 [“Logs - Audit Forwarding” on page 262](#)) to forward the events. For a detailed understanding of the format of CEF events, see [Appendix A, Common Event Format, on page 373](#).

The following events are logged and available for Audit Forwarding to ArcSight ESM.

[“Platform Events” on page 402](#)

[“Logger Application Events” on page 408](#)

Types of Audit Events

Two types of audit events are generated on Logger:

- Platform Events—related to the Logger hardware/system
- Logger Application Events—related to Logger functions and configuration changes on it

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.



Note

Platform audit events are not stored on the Logger appliance; therefore, you cannot search for them using the search facility available on the appliance. However, you can search on Alerts that are configured for these audit events.

Logger application events are searchable as usual.

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—(key name 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—User name
- duid—User ID
- src—IP address of client
- dst—IP address of appliance
- cat—Device Event Category
- cn1—Session number
- cn1label—Session

In addition, additional fields (if applicable) are listed in the following table.

Device Event Class ID	Severity	Message	Device Event Category (cat)	Additional Fields (listed only if applicable)
platform:200	7	Failed password change	/Platform/Authentication/Failure/Password	
platform:201	7	Failed login attempt	/Platform/Authentication/Failure/Login	
platform:202	5	Password changed	/Platform/Authentication/Password	
platform:205	5	Access enabled for support personnel	/Platform/Support/Enable	
platform:206	1	Access disabled for support personnel	/Platform/Support/Disable	
platform:210	3	Global login settings modified	/Platform/Configuration/Global/Login	
platform:211	3	Password policy modified	/Platform/Configuration/Global/Policy	
platform:212	5	Authentication settings modified	/Platform/Configuration/Global/RADIUS	cs1=server IP:port cs1label=RADIUS server
platform:213	7	Audit forwarding modified	/Platform/Configuration/Global/AuditEvents	
platform:220	5	Installed certificate	/Platform/Certificate/Install	

Device Event Class ID	Severity	Message	Device Event Category (cat)	Additional Fields (listed only if applicable)
platform: 221	7	Certificate mismatch failure	/Platform/Certificate/Mismatch	
platform: 222	1	Created certificate signing request	/Platform/Certificate/Request	
platform: 223	5	Certificate request expired	/Platform/Certificate/Expired	
platform: 225	7	Uploaded file damaged or corrupt	/Platform/Update/Failure/CorruptFile	fname=filename fsize=size
platform: 226	7	Uploaded package damaged or corrupt	/Platform/Update/Failure/CorruptPackage	cs1=corrupt checksum cs1label=Error cs2=time cs2label=Unpack time fname=filename fsize=size
platform: 227	5	Applied appliance update	/Platform/Update/Applied	cs1=flag cs1label=Reboot required cs2=time cs2label=Unpack time cs3=time cs3label=Install time fname=filename fsize=size
platform: 228	5	Failed to install package	/Platform/Update/Failure/Installation	cs2=time cs2label=Unpack time cs3=time cs3label=Install time fname=filename fsize=size
platform: 230	5	Successful login	/Platform/Authentication/Login	cs1label=Radius Server cs1value=server_ip: port
platform: 231	5	Successful login (RADIUS)	/Platform/Authentication/Login/RADIUS	cs1label=Radius Server cs1value=server_ip: port
platform: 232	7	Failed login attempt (BADUSER)	/Platform/Authentication/Failure/BADUSER	cs1label=Radius Server cs1value=server_ip: port
platform: 233	7	Failed login attempt (BADPASS)	/Platform/Authentication/Failure/BADPASS	cs1label=Radius Server cs1value=server_ip: port

Device Event Class ID	Severity	Message	Device Event Category (cat)	Additional Fields (listed only if applicable)
platform: 234	7	Failed login attempt (LOCKED)	/Platform/Authentication/Failure/LOCKED	cs1label=Radius Server cs1value=server_ip:port
platform: 235	7	Failed login attempt (INTERNAL)	/Platform/Authentication/Failure/INTERNAL	cs1label=Radius Server cs1value=server_ip:port
platform: 236	7	Failed login attempt (EBADAUTH)	/Platform/Authentication/Failure/EBADAUTH	cs1label=Radius Server cs1value=server_ip:port
platform: 237	7	Failed login attempt (ETIMEOUT)	/Platform/Authentication/Failure/ETIMEOUT	cs1label=Radius Server cs1value=server_ip:port
platform: 238	7	Failed login attempt (NOACCESS)	/Platform/Authentication/Failure/NOACCESS	cs1label=Radius Server cs1value=server_ip:port
platform: 239	1	User logout	/Platform/Authentication/Logout	
platform: 240	3	Added user group	/Platform/Groups/Add	fileID=ID fileType=type fname=filename
platform: 241	3	Updated user group	/Platform/Groups/Update	fileID=ID fileType=type fname=filename
platform: 242	3	Removed all members from group	/Platform/Groups/Membership/Remove	fileID=ID fileType=type fname=filename
platform: 243	3	Modified user group membership	/Platform/Groups/Membership/Update	fileID=ID fileType=type fname=filename
platform: 244	3	Deleted user group	/Platform/Groups/Remove	fileID=ID fileType=type fname=filename
platform: 245	3	Added user	/Platform/Users/Add	fileID=ID fname=filename
platform: 246	3	Updated user	/Platform/Users/Update	fileID=ID fname=filename
platform: 247	3	Deleted user	/Platform/Users/Delete	fileID=ID fname=filename

Device Event Class ID	Severity	Message	Device Event Category (cat)	Additional Fields (listed only if applicable)
platform: 250	5	Added remote mount point	/Platform/Storage/RFS/Add	cs1=IP_address cs1label=Server cs2=remote_directory_Path cs2label=Remote directory cs3=mount_point cs3label=Mount point cs4=type cs4label=Mount type cs5=permission cs5label=Username
platform: 251	5	Edited remote mount point	/Platform/Storage/RFS/Edit	cs1=IP_address cs1label=Server cs2=remote_directory_path cs2label=Remote directory cs3=mount_point cs3label=Mount point cs4=type cs4label=Mount type
platform: 252	7	Failed to create remote mount point	/Platform/Storage/RFS/Failure	cs1=IP_address cs1label=Server cs2=remote_directory_path cs2label=Remote directory cs3=mount_point cs3label=Mount point cs4=type cs4label=Mount type cs5=permission cs5label=Username
platform: 253	5	Removed remote mount point	/Platform/Storage/RFS/Remove	cs1=IP_address cs1label=Server cs2=fileservers cs2label=Remote directory cs3=trump cs3label=Mount point cs4=type cs4label=Mount type
platform: 254	5	Destroyed SAN Logical Unit	/Platform/Storage/SAN/Destroy	cs1=WWW_device cs1label=Device cs2=WWN_number cs2label=WWN cs3=WWW_label cs3label=Label

Device Event Class ID	Severity	Message	Device Event Category (cat)	Additional Fields (listed only if applicable)
platform: 255	5	Attached SAN Logical Unit	/Platform/Storage/SAN /Attach	cs1=WWW_device cs1label=Device cs2=WWN_number cs2label=WWN cs3=WWW_label cs3label=Label
platform: 256	7	Detached SAN Logical Unit	/Platform/Storage/SAN /Detach	cs1=WWW_device cs1label=Device cs2=WWN_number cs2label=WWN cs3=WWW_label cs3label=Label
platform: 257	5	Removed SAN Logical Unit	/Platform/Storage/SAN /Remove	cs1=WWW_device cs1label=Device cs2=WWN_number cs2label=WWN cs3=WWW_label cs3label=Label
platform: 259	5	Reattached SAN Logical Units	/Platform/Storage/SAN /Reattach	cs1=WWW_device cs1label=Device cs2=WWN_number cs2label=WWN cs3=WWW_label cs3label=Label
platform: 260	5	Static route modified	/Platform/Configuration /Network/Route /Update	cn2=route_ID cn2label=Route ID cs1=int_name cs1label=Interface cs2=dest_IP cs2label=Destination cs3=subnet_mask cs3label=Subnet cs4=gateway_IP cs4label=Gateway
platform: 261	5	Static route deleted	/Platform/Configuration /Network/Routes/ Remove	cn2=route_ID cn2label=Route ID cs1=int_name cs1label=Interface cs2=dest_IP cs2label=Destination cs3=subnet_mask cs3label=Subnet cs4=gateway_IP cs4label=Gateway
platform: 262	5	Appliance time modified	/Platform/Configuration /Time	

Device Event Class ID	Severity	Message	Device Event Category (cat)	Additional Fields (listed only if applicable)
platform: 263	5	Network settings modified	/Platform/Configuration/Network	cs1=NIC0_IP_mask cs1label=NIC0 cs2=NIC12_IP_mask cs2label=NIC1 cs4=default_gateway cs4label=Default gateway cs5=flag cs5label=Multi-homing
platform: 264	5	NTP server settings modified	/Platform/Configuration/Network/NTP	
platform: 265	5	DNS settings modified	/Platform/Configuration/Network/DNS	
platform: 266	5	Hosts file modified	/Platform/Configuration/Network/Hosts	
platform: 267	5	SMTP settings modified	/Platform/Configuration/Network/SMTP	cs1=SMTP_IP cs1label=SMTP Server cs2=outgoing_address cs2label=Outgoing Address
platform: 268	5	Static route added	/Platform/Configuration/Network/Route/Add	cn2=ID cn2label=Route ID cs1=interface_name cs1label=Interface cs2label=Destination cs3=subnet_mask cs3label=Subnet cs4=default_gateway cs4label=Gateway
platform: 270	9	Stopped process '<process>'	/Platform/Process/Control/Stop	dproc=nullmailer
platform: 271	7	Restarted process '<process>'	/Platform/Process/Control/Restart	dproc=nullmailer
platform: 272	5	Started process '<process>'	/Platform/Process/Control/Start	dproc=nullmailer
platform: 280	7	Appliance reboot initiated	/Appliance/State/Reboot/Initiate	
platform: 281	3	Appliance reboot canceled	/Appliance/State/Reboot/Cancel	
platform: 282	9	Appliance poweroff initiated	/Appliance/State/Shutdown	

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	Message	Device Event Category (cat)	Additional Fields
Filters			
logger: 500	Filter [filterName] has been added	/Logger/Resource/Filter/Configuration/Add	fname=filterName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Filter fileId=filterId
logger: 501	Filter [filterName] has been deleted	/Logger/Resource/Filter/Configuration/Delete	fname=filterName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Filter fileId=filterId
logger: 502	Filter [filterName] has been updated	/Logger/Resource/Filter/Configuration/Update	fname=filterName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Filter fileId=filterId
Devices			
logger: 510	Device [deviceName] has been added	/Logger/Resource/Device/Configuration/Add	fname=deviceName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Device fileId=deviceId
logger: 511	Device [deviceName] has been deleted	/Logger/Resource/Device/Configuration/Delete	fname=deviceName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Device fileId=deviceId
logger: 512	Device [deviceName] has been updated	/Logger/Resource/Device/Configuration/Update	fname=deviceName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Device fileId=deviceId

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
Groups			
logger: 513	Group [groupName] has been added	/Logger/Resource/Group/Configuration/Add	fname=groupName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Group fileId=groupId
logger: 514	Group [groupName] has been deleted	/Logger/Resource/Group/Configuration/Delete	fname=groupName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Group fileId=groupId
logger: 515	Group [groupName] has been updated	/Logger/Resource/Group/Configuration/Update	fname=groupName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Group fileId=groupId
Archives			
logger: 520	Archive [archiveName] has been added	/Logger/Resource/Archive/Configuration/Add	fname=archiveName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger: 521	Archive [archiveName] has been deleted	/Logger/Resource/Archive/Configuration/Delete	fname=archiveName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger: 523	Archive [archiveName] has been loaded	/Logger/Resource/Archive/Configuration/Load	fname=archiveName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger: 524	Archive [archiveName] has been unloaded	/Logger/Resource/Archive/Configuration/Unload	fname=archiveName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
logger: 525	Archive [archiveName] has been archived	/Logger/Resource/Archive/Configuration/Archive	fname=archiveName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger: 526	Event archive settings added	/Logger/Resource/Archive/Add	fname=archiveName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger: 527	Daily archive task settings updated	/Logger/Resource/Archive/Update	fname=archiveName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
Storage Groups			
logger: 530	Storage group [storageGroupName] has been added	/Logger/Resource/StorageGroup/Configuration/Add	fname=storageGroupName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Group fileId=storageGroupId
logger: 532	Storage group [storageGroupName] has been updated	/Logger/Resource/StorageGroup/Configuration/Update	fname=storageGroupName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Group fileId=storageGroupId
Storage Rule			
logger: 533	Storage rule [name] has been added	/Logger/Resource/StorageRule/Configuration/Add	fname=storageRuleName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Rule
logger: 535	Storage rule [name] has been updated	/Logger/Resource/StorageRule/Configuration/Update	fname=storageRuleName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Rule

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
Storage Volume			
logger: 536	Storage volume [name] has been added	/Logger/Resource/StorageVolume/Configuration/Add	fname=storageVolumeName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Volume fileId=storageVolumeId
Saved Search			
logger: 540	Saved search [name] has been added	/Logger/Resource/SavedSearch/Configuration/Add	fname=savedSearchName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Saved Search fileId=savedSearchId
logger: 541	Saved search [name] has been deleted	/Logger/Resource/SavedSearch/Configuration/Delete	fname=savedSearchName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Saved Search fileId=savedSearchId
logger: 542	Saved search [name] has been updated	/Logger/Resource/SavedSearch/Configuration/Update	fname=savedSearchName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Saved Search fileId=savedSearchId
Peer Loggers			
logger: 550	Peer Logger [name] has been added	/Logger/Resource/PeerLogger/Configuration/Add	fname=peerName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger fileId=peerLoggerId
logger: 551	Peer Logger [name] has been deleted	/Logger/Resource/PeerLogger/Configuration/Delete	fname=peerName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger fileId=peerLoggerId
logger: 570	Peer Logger authorization [name] has been added	/Logger/Resource/PeerLogger/Authorizations/Configuration/Add	fname=peerName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger Authorization

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
logger:571	Peer Logger authorization [name] has been deleted	/Logger/Resource/PeerLogger/Authorizations/Configuration/Delete	fname=peerName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger Authorization fileId=peerLoggerId
Event Input/Output			
logger:600	Receiver [name] has been added	/Logger/Component/Receiver/Configuration/Add	fname=receiverName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort
logger:601	Receiver [name] has been deleted	/Logger/Component/Receiver/Configuration/Delete	fname=receiverName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort
logger:602	Receiver [name] has been updated	/Logger/Component/Receiver/Configuration/Update	fname=receiverName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort
logger:603	Receiver [name] has been enabled	/Logger/Component/Receiver/Configuration/Enable	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	Message	Device Event Category (cat)	Additional Fields
logger:604	Receiver [name] has been disabled	/Logger/Component/Receiver/Configuration/Disable	fname=receiverName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort
logger:605	Forwarder [name] has been added	/Logger/Component/Forwarder/Configuration/Add	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	Forwarder [name] has been deleted	/Logger/Component/Forwarder/Configuration/Delete	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	Forwarder [name] has been updated	/Logger/Component/Forwarder/Configuration/Update	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	Message	Device Event Category (cat)	Additional Fields
logger:608	Forwarder [name] has been enabled	/Logger/Component/Forwarder/Configuration/Enable	fname=forwarderName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName e cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger:609	Forwarder [name] has been disabled	/Logger/Component/Forwarder/Configuration/Disable	fname=forwarderName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName e cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger:663	Forwarder [name] has been paused	/Logger/Component/Forwarder/Configuration/Pause	fname=forwarderName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName e cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger:664	Forwarder [name] has been resumed	/Logger/Component/Forwarder/Configuration/Resume	fname=forwarderName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName e cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
logger: 640	ESM destination [name] has been added	/Logger/Component/EsmDestination/Configuration/Add	fname=esmDestinationName e_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=Connector Name cs1=connectorName cs2Label=Connector Location cs2=connectorLocation cs3Label=Logger Location cs3=loggerLocation
logger: 641	ESM destination [name] has been deleted	/Logger/Component/EsmDestination/Configuration/Delete	fname=esmDestinationName e_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=Connector Name cs1=connectorName cs2Label=Connector Location cs2=connectorLocation cs3Label=Logger Location cs3=loggerLocation
logger: 643	Certificate [name] has been added	/Logger/Component/Certificate/Configuration/Add	fname=alias duser=username duid=userId cs4=sessionId cs4Label=Session ID fileType=Certificate
logger: 650	Certificate [name] has been deleted	/Logger/Component/Certificate/Configuration/Delete	fname=alias duser=username duid=userId cs4=sessionId cs4Label=Session ID fileType=Certificate
logger: 651	Certificate [name] has been updated	/Logger/Component/Certificate/Configuration/Update	fname=alias duser=username duid=userId cs4=sessionId cs4Label=Session ID fileType=Certificate

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
logger: 644	SNMP destination [name] has been added	/Logger/Component/SnmpDestination/Configuration/Add	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=Connector Name cs1=connectorName cs2Label=Connector Location cs2=connectorLocation cs3Label=Logger Location cs3=loggerLocation
logger: 645	SNMP destination [name] has been deleted	/Logger/Component/SnmpDestination/Configuration/Delete	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=Connector Name cs1=connectorName cs2Label=Connector Location cs2=connectorLocation cs3Label=Logger Location cs3=loggerLocation
logger: 647	Syslog destination [name] has been added	/Logger/Resource/SyslogDestination/Configuration/Add	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	Message	Device Event Category (cat)	Additional Fields
logger: 648	Syslog destination [name] has been deleted	/Logger/Component/SyslogDestination/Configuration/Delete	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	Syslog destination [name] has been updated	/Logger/Component/SyslogDestination/Configuration/Update	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
Alerts			
logger: 610	Alert [name] has been added	/Logger/Component/Alert/Configuration/Add	fname=alertName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmpIpAddr dvchost=syslogOrSnmpHost Name cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
logger: 611	Alert [name] has been deleted	/Logger/Component/Alert/Configuration/Delete	fname=alertName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmpIpAddr dvchost=syslogOrSnmpHost Name cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
logger:612	Alert [name] has been updated	/Logger/Component/Alert/Configuration/Update	fname=alertName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmpIpAddr dvchost=syslogOrSnmpHost Name cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
logger:613	Alert [name] has been enabled	/Logger/Component/Alert/Configuration/Enable	fname=alertName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmpIpAddr dvchost=syslogOrSnmpHost Name cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
logger:614	Alert [name] has been disabled	/Logger/Component/Alert/Configuration/Disable	fname=alertName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmpIpAddr dvchost=syslogOrSnmpHost Name cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
logger: 615	Alert [name] has been sent	/Logger/Component/Alert/Configuration/Sent	fname=alertName duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmpIpAddr dvchost=syslogOrSnmpOrEsmHostName cn1Label=SyslogoOrSnmpOrEsmDestination Port cn1=syslogOrSnmpOrEsmPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
Configuration Backup			
logger: 660	Configuration backup has been updated	/Logger/Component/ConfigBackup/Configuration/Update	fname=Configuration Backup duser=username duid=userId cs4=sessionId cs4Label=Session ID fileType=Configuration Backup
logger: 661	Configuration backup has been enabled	/Logger/Component/ConfigBackup/Configuration/Enable	fname=Configuration Backup duser=username duid=userId cs4=sessionId cs4Label=Session ID fileType=Configuration Backup
logger: 662	Configuration backup has been disabled	/Logger/Component/ConfigBackup/Configuration/Disable	fname=Configuration Backup duser=username duid=userId cs4=sessionId cs4Label=Session ID fileType=Configuration Backup
Search			
logger: 680	Search indices have been added	/Logger/Search/Index/Update	
	OR		
	Search index has been added		
logger: 690	Search options have been updated	/Logger/Search/Options/Update	
Maintenance Mode			
logger: 700	Maintenance mode entered	/Logger/Server/MaintenanceMode/Enter	fname=Maintenance Mode duser=username duid=userId cs4=sessionId cs4Label=Session ID fileType=Maintenance Mode

Appendix F

Event Field Name Mappings

The nomenclature for field names depends on the function area of the Logger. The following table provides the mapping between those names.

Database Name	Search Results	Actual Events (CEF Key)	Reports
arc_deviceVendor	deviceVendor	Device Vendor	Device Vendor
arc_deviceProduct	deviceProduct	Device Product	Device Product
arc_deviceVersion	deviceVersion	Device Version	Device Version
arc_deviceEventClassId	deviceEventClassId	Signature ID	Signature Id
arc_name	name	Name	Name
arc_agentSeverity	agentSeverity	Severity	Severity
arc_agentAddress	agentAddress	agt	Agent Address
arc_agentHostName	agentHostName	ahost	Agent Host Name
arc_agentNtDomain	agentNtDomain	agentNtDomain	Agent NT Domain
arc_agentType	agentType	at	Agent Type
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
arc_categoryObject	categoryObject	categoryObject	Category Object
arc_categoryOutcome	categoryOutcome	categoryOutcome	Category Outcome
arc_categorySignificance	categorySignificance	categorySignificance	Category Significance
arc_categoryTechnique	categoryTechnique	categoryTechnique	Category Technique

Database Name	Search Results	Actual Events (CEF Key)	Reports
arc_destinationAddress	destinationAddress	dst	Destination Address
arc_destinationDnsDomain	destinationDnsDomain	destinationDnsDomain	Destination DNS Domain
arc_destinationHostName	destinationHostName	dhost	Destination Host Name
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_destinationUserId	destinationUserId	duid	Destination User ID
arc_destinationUserName	destinationUserName	duser	Destination User Name
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
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

Database Name	Search Results	Actual Events (CEF Key)	Reports
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_deviceExternalId	deviceExternalId	deviceExternalId	Device External Id
arc_deviceHostName	deviceHostName	dvchost	Device Host Name
arc_deviceInboundInterface	deviceInboundInterface	deviceInboundInterface	Device Inbound Interface
arc_deviceOutboundInterface	deviceOutboundInterface	deviceOutboundInterface	Device Outbound Interface
arc_deviceReceiptTime	deviceReceiptTime	rt	Device Receipt Time
arc_deviceSeverity	deviceSeverity	deviceSeverity	Device Severity
arc_deviceZoneURI	deviceZoneURI	deviceZoneURI	Device Zone URI
arc_endTime	endTime	deviceEndTime	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_flexNumber1	flexNumber1	flexNumber1	Flex Number1

Database Name	Search Results	Actual Events (CEF Key)	Reports
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_priority	priority	priority	Priority
arc_sourceAddress	sourceAddress	src	Source Address
arc_sourceHostName	sourceHostName	shost	Source Host Name
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_sourceUserId	sourceUserId	suid	Source User Id
arc_sourceUserName	sourceUserName	suser	Source User Name
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
arc_agentZone	agentZone	agentZone	Agent Zone
arc_agentZoneName	agentZoneName	agentZoneName	Agent Zone Name
arc_agentZoneResource	agentZoneResource	agentZoneResource	Agent Zone Resource
arc_destinationZone	destinationZone	destinationZone	Destination Zone
arc_destinationZoneName	destinationZoneName	destinationZoneName	Destination Zone Name
arc_destinationZoneResource	destinationZoneResource	destinationZoneResource	Destination Zone Resource

Database Name	Search Results	Actual Events (CEF Key)	Reports
arc_deviceZone	deviceZone	deviceZone	Device Zone
arc_deviceZoneName	deviceZoneName	deviceZoneName	Device Zone Name
arc_deviceZoneResource	deviceZoneResource	deviceZoneResource	Device Zone Resource
arc_sourceZone	sourceZone	sourceZone	Source Zone
arc_sourceZoneName	sourceZoneName	sourceZoneName	Source Zone Name
arc_sourceZoneResource	sourceZoneResource	sourceZoneResource	Source Zone Resource

Connector Appliance Documentation

This information is applicable only to Logger **appliance platforms with integrated Connector Appliance**.

Connector Appliance documentation is available as follows:

- The [Chapter 8, Managing Connectors on Connector Appliance, on page 303](#) and [Chapter 9, Managing Repositories in Connector Appliance, on page 355](#) chapters in this guide.
- Through the Help icon (🔍) on any user interface page, when you are in the Connector Appliance context. When you click this icon, a PDF of the Connector Appliance Administrator's Guide is displayed. **All information in this guide except system administration is applicable to your product.**
- Through the ArcSight Customer Support site at <https://support.arcsight.com>.

Destination Runtime Parameters

The information in this chapter is applicable only to Logger **appliance platforms with integrated Connector Appliance**.

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 341](#).

Name Fields	Value Fields
Batching	SmartConnectors can batch events to increase performance and optimize network bandwidth. When activated, SmartConnectors create blocks of events and send them when they either (1) reach a certain size or (2) the time window expires, whichever occurs first. You can also prioritize batches by severity, forcing the SmartConnector to send the highest-severity event batches first and the lowest-severity event batches later.
Enable Batching (per event)	Create batches of events of this specified size (5, 10, 20, 50, 100 , 200, 300 events).
Enable Batching (in seconds)	The SmartConnector sends the events if this time window expires (1, 5 , 10, 15, 30, 60).
Batch By	This is Time Based if the SmartConnector should send batches as they arrive (the default) or Severity Based if the SmartConnector should send batches based on severity (batches of Highest Severity events sent first).
Time Correction	The values you set for these fields establish forward and backward time limits, that if exceeded, cause the SmartConnector 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 SmartConnector can adjust the time reported by the device Detect Time , using this setting. This is useful when a remote device's clock isn't synchronized with the ArcSight ESM 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 SmartConnector can also adjust the time reported by the SmartConnector itself, using this setting. This is for informational purposes only and allows you to modify the local time on the SmartConnector. This should be a temporary setting. The recommended way to synchronize clocks between Manager and SmartConnectors is the NTP protocol. 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 SmartConnector is doing so. If this is not true, or the device isn't reporting correctly, you can switch this option from Disabled to GMT or to a particular world time zone. That zone is then applied to the time reported. 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 SmartConnector checks its future and past thresholds at intervals specified by this number of seconds. Default is 1 minute (60 seconds).

Cache

	Changing these settings will not affect the events cached, it will only affect new events sent to the cache.
Cache Size	SmartConnectors use a compressed disk cache to hold large volumes of events when the ArcSight ESM Manager is down or when the SmartConnector receives bursts of events. This parameter specifies the disk space to use. The default is 1 GB which, depending on the connector, can hold about 15 million events, but it also can go down to 5 MB . When this disk space is full, the SmartConnector drops the oldest events to free up disk cache space. (5 MB, 50 MB, 100 MB, 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 that the heartbeat is also used to communicate with the SmartConnector; therefore, if its frequency is set to 10 minutes , then it could take as much as 10 minutes to send any configuration information or commands back to the SmartConnector.
Enable Name Resolution	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 Mbits/sec.)
Transport Mode	You can configure the SmartConnector to cache to disk all the processed events it receives. This is equivalent to pausing the SmartConnector. However, you can use this setting to delay event-sending during particular time periods. For example, you could use this setting to cache events during the day and send them at night. You can also set the connector to cache all events, except for those marked with a very-high severity, during business hours, and send the rest at night. (Normal Cache Cache (but send Very High severity events)).
Address-based Zone Population Defaults Enabled	This field applies to v3.0 ArcSight ESM Managers. This field is not relevant in ESM v3.5 because the system has integral zone mapping. Default: Yes .
Address-based Zone Population	This field applies to v3.0 ArcSight ESM Managers. This field is not relevant in ESM v3.5 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 ESM 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 SmartConnector's source address. This field is present for ESM v3.0 compatibility. It is not relevant in ESM v3.5 because of integral zone mapping.

Source Translated Zone URI	When populated, this field shows the URI of the zone associated with the SmartConnector's translated source address. The translation is presumed to be NAT (network address translation). This field is present for ESM v3.0 compatibility. It is not relevant in ESM v3.5 because of integral zone mapping.
Destination Zone URI	When populated, this field shows the URI of the zone associated with the SmartConnector's destination address. This field is present for ESM v3.0 compatibility. It is not relevant in ESM v3.5 because of integral zone mapping.
Destination Translated Zone URI	When populated, this field shows the URI of the zone associated with the SmartConnector's translated destination address. The translation is presumed to be NAT (network address translation). This field is present for ESM v3.0 compatibility. It is not relevant in ESM v3.5 because of integral zone mapping.
Connector Zone URI	When populated, this field shows the URI of the zone associated with the SmartConnector's address. This field is present for ESM v3.0 compatibility. It is not relevant in ESM v3.5 because of integral zone mapping.
Connector Translated Zone URI	When populated, this field shows the URI of the zone associated with the SmartConnector's translated address. The translation is presumed to be NAT (network address translation). This field is present for ESM v3.0 compatibility. It is not relevant in ESM v3.5 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 v3.0 compatibility. It is not relevant in ESM v3.5 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 v3.0 compatibility. It is not relevant in ESM v3.5 because of integral zone mapping.

Field Based Aggregation 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.

SmartConnector aggregation significantly reduces the amount of data received, and should be applied only when you use less than the total amount of information the event offers. For example, you could enable field-based aggregation to aggregate "accepts" and "rejects" in a firewall, but you should use it only if you are interested in the count of these events, instead of all the information provided by the firewall.

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 (i.e., 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	
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).	
SmartConnector aggregation significantly reduces the amount of data received, and should be applied only when you use less than the total amount of information the event offers.	
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 (i.e., 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

If your configuration, reporting, and analytic usage permits, you can greatly accelerate the transfer of a sensor's event information through SmartConnectors by choosing one of two "turbo" (narrower data bandwidth) modes. The default transfer mode is called **Complete**, which passes all the data arriving from the device, including any additional data (custom, or vendor-specific).

Complete mode does indeed use all the database performance advances of ArcSight ESM v3.x.

The first level of Turbo acceleration is called **Faster** and drops just additional data, while retaining all other information. The **Fastest** mode eliminates all but a core set of event attributes, in order to achieve the best throughput. Consider the possible effects such a restricted data set might have from a given device (e.g., on reports, rules, threat resolution) before selecting it.

The specific event attributes that apply to these modes in your enterprise are defined in the self-documented [\\$ARCSIGHT_HOME/config/connector/agent.properties](#) file for the ArcSight ESM Manager. Because these properties might have been adjusted for your needs, you should refer to this file for definitive lists. Only scanner SmartConnectors need to run in **Complete** mode, to capture the additional data.



SmartConnector Turbo Modes are superseded by the Turbo Mode in use by the ArcSight ESM Managers processing their events. For example, a Manager set to **Faster** will not pass all the data possible for a SmartConnector that is set for the default of **Complete**.

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 SmartConnector's burden on the CPU by reducing its processing rate. This can also be a means of dealing with the effects of event bursts.</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 SmartConnector in use, as described in the SmartConnector 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**.

Logger Search From An ESM Console

If you have ArcSight Logger and ArcSight ESM deployed in your network infrastructure, you can perform a Logger search operation directly from your ESM Console.

Understanding the Integrated Search Functionality

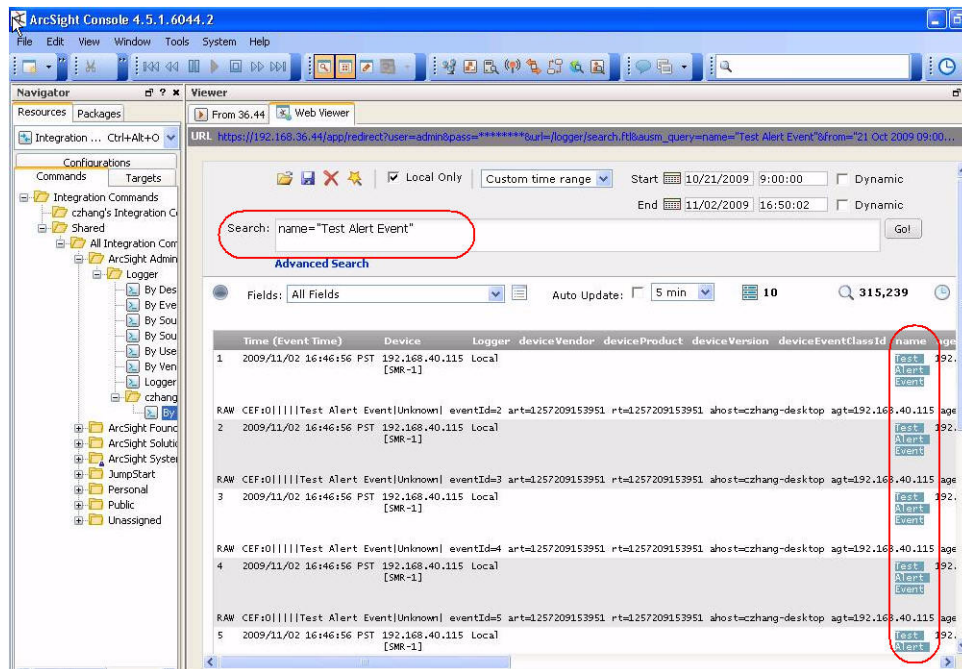
There are two ways to perform a search operation on Logger from an ESM Console:

- 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 ESM Console active channel; you are not prompted for any search options.

To run a Logger search, you right click on an event in an active channel of the ESM Console to display a menu. You select the search method—Logger Search or Logger Quick Search—from the menu.

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 ESM Console, as shown in the following figure:



Prerequisites

The Logger on which search will be performed must be running Logger v4.0. The ESM Manager from which search will be initiated must be running ESM v4.5 SP1 Patch2.

Setup and Configuration

ESM

Follow these instructions to set up and configure ArcSight ESM Manager to run integrated search operations:

- 1 Ensure that the ESM Manager is running v4.5 SP1 Patch2.
- 2 Follow instructions in the *ArcSight ESM v4.5 SP1 Patch2 Release Notes* to set up ESM Console for integrated searches on Logger.

The ESM release notes are available from the ArcSight Customer Support web site at <https://support.arcsight.com>.

Logger

Make sure:

- 1 Your Logger is running v4.0.
- 2 A Logger user name that you specified when creating an integration parameter on ESM Console (Step 2 of ESM in "Setup and Configuration" on page 438) 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 ESM 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 ESM 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 ESM Console

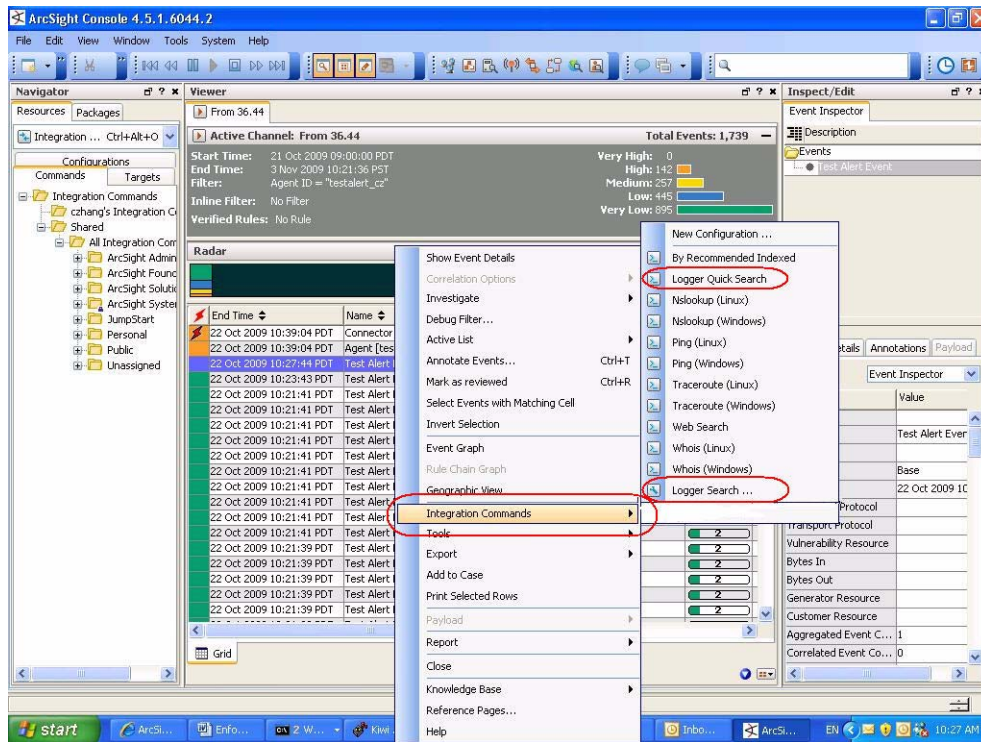
- A field-based search query is used to perform search on the Logger.
- A search operation only from an active channel of an ESM Console is supported; search operation from other ESM resources is not supported.
- Multiple search options (see [“Supported Search Options” on page 439](#)) cannot be specified 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 ESM 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 ESM Console.

Searching on Logger From ESM Console

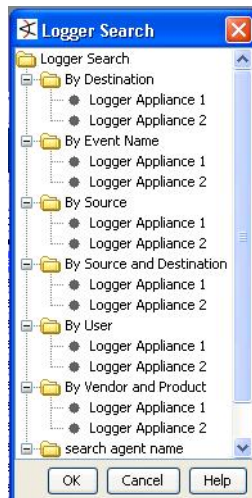
- To run a **Quick Search** on Logger (as described in [“Understanding the Integrated Search Functionality” on page 437](#)):
 - a Right click on the event field in an active channel of the ESM Console.

- 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 ESM 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.

The search results are displayed in the ESM Console Web Viewer.

Index

Symbols

.aup file for content update 363

A

accounts, user. See User.

Acronis True Image Server 393

agents. See SmartConnectors

Alerts

about 207

adding 210

disable 211, 218

enable 211, 218

remove 212

apache status 260, 292

archive, scheduled 186

ArcSight ESM 2, 7, 31, 33, 199, 207

ArcSight Logger Streaming SmartConnector 200

ArcSight Manager 8, 31, 33, 199, 205

ArcSight SmartConnectors 7, 31

Audit forwarding 401

Audit Log 262, 293, 401

AUP upgrade process 361

authentication, RADIUS 282

automatic timeout 280

B

backup, configuration 235

batching 429

browser requirements 35

bulk copy (see cloning) 371

C

CA certificate

applying on container 322

demo 321

invalid errors 327

managing 321

removing from container 324

viewing list 325

CAC support 273

CACERTS for ESM Destination 205

canonical equality check 37

cases 349

case-sensitive search 37

Categories tab 435

CEF 86, 373

CEF event filters 86

certificate revocation list 274

Certificate Signing Request 271

Certificate, installation 272

changing container credentials 319

character encoding 196

CIFS, configuring 331

CLI 17

command table 18

cloning SmartConnectors 371

Comma Separated Values file, uploading 308

command line interface (CLI) 17

Common Event Format (CEF) 86, 373

Common Extension Dictionary 375

Configuration Backup 235

configuration monitoring, reports for 97

Configuration tab 181, 251

Connector Appliance

remote upgrade 361

Connector Forwarder 200

connectors supported 331

connectors. See SmartConnectors

constraints, search 44

containers

adding 317

changing credentials 319

definition 316

deleting 318

editing 317

running commands 327

updating properties 318

upgrading 328

viewing all 316

viewing logs 329

content AUP 363

copying (see cloning) 371

CSR

generating a certificate signing request 272

CSV file information 308

current time, changing 256

customers 431

D

dashboard

reports 100

reports, preference for display 112

date/time format 199

default gateway 254

Default Storage Group 2, 28

demo certificate 321

deploying

report package 177

Device 181, 182

delete 182

- edit 182
- pre-defining 182
- Device Group 183
 - creating 183
 - deleting 184
 - editing 184
- device group
 - maximum number 183
- devices
 - maximum number 182
- DNS Settings 253
 - changing 253, 254
- dynamic search 55

E

- e-mailing
 - reports 122
- encoding 196
- Error Log 262, 293
- ESM (ArcSight Enterprise Security Manager) 2, 7, 31, 33, 199
- ESM Destination 203
 - creating 204
 - deleting 205
 - updating CACERTS 205
- ESM SmartConnector status 260, 292
- etc/hosts.txt 253
- Event Archive 184, 185
 - adding 185
 - deleting 185
 - loading 185
 - settings 186
 - unloading 185
- event archive, scheduled 186
- Event Input/Output 193
- event storage, remote 191
- events
 - search 43
- export
 - search results 80

F

- factory settings, restoring 393
- field query
 - indexing fields 81
- field set, search 44
- fields, indexing 81
- File events to ESM 207
- Filter 222
 - copying 223
 - creating 222
 - deleting 223
 - editing 223
- Filter, Report Category 180
- filter, search 44
- filtering information on UI page 306
- filters, system 86
- find, events 43
- FIPS 140-2
 - enabling on Connector Appliance 275
 - enabling on container 320
- Firefox (web browser) 35
- Forwarder 199

- creating 200
- deleting 203
- editing 202
- Forwarder status 260, 292
- forwarding file events to ESM 207
- function tabs 36

G

- gateway, default 254
- gauge range 37
- gauges 36
- gid 265
- Global Settings 279

H

- health, system 88
- Help
 - how to use Console Online Help xv
- help 36
- hosts
 - adding 311
 - definition 310
 - deleting 314
 - editing 315
 - moving to different location 314
 - scanned 311
 - scanning 313
 - software-type 310
 - upgrading remotely 315
 - viewing all 310
- Hosts file 253

I

- i18n options 37
- indexing fields 81
- initialization 16
- insp status 260, 292
- interface homing 255
- Internal Storage Group 2, 187
- Internet Explorer (web browser) 35
- intrusion monitoring, reports for 96
- invalid certificate errors 327
- IP addresses
 - assigning 16
 - changing 254

L

- localhost 253
- locations
 - adding 308
 - definition 307
 - deleting 309
 - editing 309
 - viewing all 307
- Logfu utility 329
- Logger
 - rebooting 252
- login 35
- Login Settings 279
 - changing 280
- logout 36, 37

logout, automatic 37
 Logs 262, 293
 logs, internal 246
 retrieving 246

M

maintenance mode 238
 Manager 8, 31, 33, 199, 205
 Monitor tab 37
 multi-homing 255

N

navigation 36
 Network Settings 254
 changing 254
 network speed 254
 NFS, configuring 331
 NTP Server 257
 NTP setting 257, 258, 291

O

Online Help
 see *Help*
 online help 36
 options 36, 37

P

package contents 15
 parameter value groups, in reports 170
 parameters
 in report queries 163
 quick run report 117
 run reports 119
 Password policy
 changing 281
 Password, changing 290, 291, 300, 301
 PCI Storage Group 2, 187
 Peer Logger 231
 adding 233
 deleting 235
 peer Logger, searching 73
 postgresql status 260, 292
 predefined filters 86
 Process Status 260, 292

Q

queries
 in reports 140
 query
 events 43
 query controls 36

R

RADIUS authentication 282
 RAID controller status 260, 270, 292
 range, gauge 37
 rebooting Logger 252
 Receiver 193
 creating 194
 deleting 195

 editing 195
 types 193, 195
 Receiver status 260, 292
 refreshing UI screen 306
 regular expressions (regex)
 predefined 86
 tutorial 380
 Remote Authentication Dial-In User Service (RADIUS)
 282
 remote event storage 191
 remote file system mount
 adding 266, 268
 deleting 265, 267
 editing 265, 267, 268
 remote upgrade 361
 Report Category Filter 180
 reports 226
 access rights 140
 administration 178
 categories 94
 configuration monitoring 97
 creating new 125
 dashboard 100
 delivery options 122
 designing 124
 editing 138
 e-mailing 122
 exporting 123
 file formats 120
 foundation 95
 groups 94
 intrusion monitoring 96
 navigating to 93
 parameter value groups 170
 PCI solution add-on 97
 publishing 121
 query parameters 163
 quick run parameters 117
 remove scheduled 174
 run parameters 119
 running 113
 SANS Top 5 95
 saving 123
 scheduling 173
 solution add-ons 97
 template styles 172
 user-created 98
 viewing published 124
 viewing, editing schedules 174
 repositories, user-defined 365
 reset to factory settings 393
 restore to factory settings 393
 restoring a SAN 269
 Retrieve Logs 246

S

safety precautions 16
 SAN Storage 267
 SAN, restore 269
 SANS Top 5, reports for 95
 saved
 filter 84
 search 84
 Saved Search 225

- adding 225
 - deleting 226
 - editing 226
- Saved Search Files 229
- Saved Search Job 226
 - adding 226
 - deleting 229
 - editing 229
- scan a host 311, 313
- scheduled event archive 186
- Scheduled Task 221
 - currently running 222
 - finished 222
- scheduling
 - export of search results 73
 - reports 173
- SCP file receiver 195
- search
 - constraints 44
 - events 43
 - exporting results 80
 - field set 44
 - filter 44, 84
 - peer Loggers 73
 - query, defining a 44
 - results, scheduling export of 73
 - saved 84
 - system filters 86
 - time range 44
- Search Group Filter
 - associating with user group 224
 - report category filter 180
- Search Group Filters 223
- Search Results tab 74
- servers status 260, 292
- severity level 429, 431, 435
- SFTP file receiver 195
- SmartConnectors 7, 31, 430, 431
 - batching 429
 - defined 331
 - scanner 434
 - zones 431, 432
- SmartMessage 193
- software-type Host 310
- solutions
 - reports 97
- speed, network 254
- SSL
 - Certificate Signing Request 271
- SSL Settings 270
- Static Route 259
 - adding 259
- statistics 36
- status
 - 3Ware RAID Controller 260, 270, 292
 - process 260, 292
- Storage 187
- Storage Area Network 267
- Storage Group 187
 - adding 188
 - Default 28
 - editing 188
- Storage Group, default 2, 28
- Storage Group, internal 2, 187

- Storage Group, PCI 2, 187
- Storage Rule 28, 189
 - adding 190
 - deleting 191
 - editing 191
- Storage Settings 191
- Storage Volume 191
 - settings 191
- streaming SmartConnector 200
- subnet mask 254
- supported connectors 331
- System Admin tab 251
- system definition 305
- system filters 86
- system health, monitoring 88
- System Information 263
- System Reboot 252
- System Update 259

T

- template styles
 - for reports 172
- time configuration 257, 258, 291
- time range, dynamic 55
- time range, search 44
- Time Settings 256
 - changing 257, 258, 291
- time, changing 256
- timeout, automatic 280
- timezone 257

U

- uid 265
- Unicode options 37
- update, content 363
- updating container properties 318
- upgrade
 - Connector Appliance 361
 - host 361
 - remote 361
- US-ASCII encoding 196
- User 284, 289, 294, 299
 - changing password 291, 301
 - creating 289, 299
 - deleting 290, 300
 - editing 290, 300
- User Group 284, 294
 - associating with Search Group Filter 224
 - creating 287, 298
 - deleting 288, 298
 - editing 288, 298
- user interface 36
 - filtering information to display 306
 - refresh 306
 - Search Results tab 74
- User password, changing 291, 301
- user-defined repositories 365
- UTF-8 encoding 196

V

- version, component 259

W

Web

- viewing Online Help in Web browser xv
- web browser requirements 35

- web status 260, 292

- What's New 9

- widgets

- in report dashboards 109

