

Administrator's Guide

ArcSight Logger™ v5.1

May 14, 2011



Administrator's Guide ArcSight Logger™ v5.1

May 14, 2011

Copyright © 2006-2011 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
05/14/11	v5.1	v5.1 release.
11/08/10	v5.0 Patch 2	A special Patch 2 update for this guide to update remote access to Logger information.
09/16/10	v5.0 GA	v5.0 GA.
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.

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	http://www.arcsight.com/supportportal
Protect 724 Community	https://protect724.arcsight.com

Logger Administrator's Guide

Contents

Chapter 1: Overview	15
Introduction	15
Logger Features	17
Storage Configuration	17
Receiver Configuration	17
Analyzing Events	18
Grouping Events	18
Exporting	19
Forwarder Configuration	19
User Management	20
Other Setup and Maintenance	20
Deployment Scenarios	20
What's New in Logger v5.1	22
Chapter 2: Installation and Initialization	23
Section 1: Installing the Logger Appliance	23
Installing the Logger Appliance	23
Setting an IP Address on the Appliance	24
Using a Browser to Set an IP Address	24
Using the Command Line Interface to Configure IP Address	25
Initializing the Logger Appliance	27
Planning	27
Initialization Sequence for Logger appliances	28
1 License	28
2 SAN	29
3 Storage Volume	29
3 Storage Groups	30
4 Time Settings	30

5 Index Fields and Full-text Indexing	31
6 Locale	31
7 Reboot	32
8 Receivers	32
9 Devices	32
10 Device Groups	32
11 Storage Rules	32
Configure Logger for Remote Access	33
Section 2: Installing the Software Version of Logger	33
Supported Platforms and Browsers	33
Downloading the Software	34
How Licensing Works on the Software Version of Logger	35
Deployment Planning for the Software Version of Logger	36
Storage Strategy	36
Retention Policy	36
Installing and Configuring the Software Version of Logger	37
Prerequisites for Installation	37
Installation Modes	38
Installation Steps	38
Using the GUI Mode to Install Software Logger	38
Using the Console Mode to Install Software Logger	46
Using the Silent Mode to Install Software Logger	47
Connecting to the Software Logger User Interface	49
Starting and Stopping the Software Logger	50
Uninstalling the Software Logger	52
Applying a License on the Software Version of Logger	52
Enabling or Disabling Logger as a System Service	53
Best Practices for the Minimal Install	54
Installing SmartConnectors to Send Events to Logger	54
SmartMessage	55
Forwarding Logger Events to an ESM Manager	56
Configuring SmartConnectors to Send Events to Both Logger and an ESM Manager	56
Configuring SmartConnectors for Failover Destinations	56
Sending Events from ArcSight ESM to Logger	57
Chapter 3: Using the User Interface	61
Connecting to the Logger User Interface	61
Browser Requirements	62
Navigating the User Interface	63
Help	63
Options	65
Logout	65
Monitor	66

Platform	68
Network	68
Logger	69
Receivers	69
Forwarders	69
Storage	69
Chapter 4: Searching and Analyzing Events	71
The Need to Search Events	71
The Process of Searching Events	71
Elements of a Search Query	72
Query Expression	72
Syntax Reference for Query Expression	98
Using the Search Builder Tool	102
Accessing Search Builder	103
Nested Conditions	105
Alternate Views for Query Building in Search Builder	105
Search Analyzer	106
Regex Helper Tool	108
Searching for Events on Logger	110
Advanced Search Options	111
Searching Peer Loggers (Distributed Search)	111
Tuning Search Performance	112
Understanding the Search Results Display	112
Multi-line Data Display	115
Auto Updating Search Results	115
Exporting Search Results	116
Scheduling an Export Operation	119
Indexing	119
How indexing works	119
Full-text Indexing (Keyword Indexing)	119
Field-based Indexing	120
Saving Queries (Saved Filters and Searches)	123
Saving a Query	123
Using a Saved Filter or a Saved Search	124
System Filters/Predefined Filters	124
Using a System Filter	127
Monitoring System Health	127
System Health Events	128
Alerts	130
Viewing Alerts	130
Receiving Alerts for Events	130
Base Event Fields	131

Go, Export, and Auto Update Options	131
Chapter 5: Reporting	133
Navigating to Reports	133
Report Groups	134
Foundation Reports	135
SANS Top 5 Reports	135
Network Monitoring Reports	136
Intrusion Monitoring Reports	136
Configuration Monitoring Reports	137
Solution Reports	137
Device Monitoring Reports	137
Anti-Virus Reports	138
Cross Device Reports	138
Database Reports	138
Firewall Reports	138
Identity Management Reports	138
IDS-IPS Reports	138
Network Reports	138
Operating System Reports	138
VPN Reports	138
User Reports	138
Reports Home Page	139
Using the Dashboard	140
Viewing the Dashboard	140
Designing Dashboards	141
What items can a dashboard include?	142
Quick Start - Creating a New Dashboard	142
Add an Empty Dashboard	143
Creating Widgets	145
Placing Dashboard Items on the Layout	145
Placing a Report on a Dashboard	146
Placing a Use Case on a Dashboard	149
Placing an External Link on a Dashboard	150
Swapping Items on Widgets	152
Setting Dashboard Preferences	152
Working with Available Dashboards	153
Selecting a Dashboard View	153
Modifying or Removing Existing Dashboards	154
Running, Viewing, and Publishing Reports	154
Best Practices	154
Finding Reports	155
Task Options on Available Reports	155

Running and Viewing Reports	156
About the Pagination of Reports	157
Quick Run / Run In Background Report Parameters	158
Run Report Parameters	160
Report File Formats	161
Publishing Reports	161
Report Delivery Options	162
Refreshing a Report	162
E-mailing a Report	162
Exporting and Saving a Report	163
Viewing the Output of a Published Report	164
Designing Reports	164
Opening the Report Designer	165
Creating New Reports	165
Quick Start: Base a New Report on an Existing One	165
Designing New Reports	167
Select Filter Criteria	170
Select Grouping	172
Select Totals	174
Sort Order	174
Highlighting	175
Create Matrix	176
Create Chart	177
Editing a Report	179
Adhoc Report Designer	179
Setting Access Rights on Reports	180
Setting up Queries	181
How Search and Report Queries Differ	182
Overview of Query Design Elements	182
Creating a Copy of an Existing Query	183
Designing a New SQL Query	183
Modifying a Query Object	197
Deleting a Query Object	198
Defining SQL in the Editor	198
Working with Parameters	205
Creating New Parameters	206
Modifying a Parameter	211
Deleting a Parameter	212
Configuring Parameter Value Groups	212
Applying Report Template Styles	214
Defining a New Template	215
Scheduling Reports	215
Viewing and Editing Scheduled Reports	216

Scheduling a Report	216
Deploying a Report Package	219
Report Server Administration	220
Using Report Category Filters	222
Backup and Restore of Report Content	222
Chapter 6: Configuration	223
Devices	223
Devices	223
Device Groups	225
Event Archives	226
Guidelines for Archiving Events	228
Archiving Events	229
Scheduled Event Archive	230
Archive Storage Settings	231
Loading and Unloading Archives	232
Storage	233
Storage Groups	233
Storage Rules	235
Storage Volume	236
Event Input/Output	238
Receivers	239
Forwarders	246
ESM Destinations	251
Forwarding Log File Events to ESM	254
Alerts	255
Configuring and Managing Real Time Alerts	258
Creating and Managing Saved Search Alerts	260
SNMP Destinations	266
Syslog Destinations	267
ESM Destinations	268
Export	269
Scheduled Tasks	269
Scheduled Tasks	269
Currently Running Tasks	270
Finished Tasks	270
Filters	270
Filters	270
Search Group Filters	272
Export	273
Saved Searches	273
Saved Searches	273
Scheduled Saved Search	275

Saved Search Files	278
Search Optimization	278
Add Search Indexes	278
Tuning Advanced Search Options	279
Deleting Custom Field Sets	280
Peer Loggers	280
Guidelines	281
Configuration Backup and Restore	284
Running a Configuration Backup (Ad-hoc or Scheduled)	285
Restoring from a Configuration Backup	286
Editing Configuration Backup Settings	286
System Maintenance	287
Database Defragmentation	288
Storage Volume Size Increase	293
License Information	295
Retrieve Logs	296
Exporting and Importing Content	297
Guidelines for Exporting and Importing	298
Exporting Content	298
Importing Content	299
Chapter 7: System Admin	301
Section 1: Logger Appliance System Administration	301
System Locale	302
Reboot	302
DNS Settings	303
Hosts	303
Network	304
Time/NTP	306
Static Routes	308
SMTP Settings	308
License & Update	309
Process Status	310
SSH Access to Logger	311
Enabling or Disabling SSH Access	311
Connecting to Logger using SSH	312
Logs - Audit Logs	313
Logs - Audit Forwarding	313
Storage	314
CIFS Settings	314
Network File System (NFS) Settings	316
SAN	318
Multipath	322

RAID Controller	323
Security	323
SSL Server Certificate	323
SSL Client Authentication (CAC Authentication)	326
FIPS 140-2	328
Users/Groups	333
Authentication Settings	333
User Groups	337
Users	343
Users/Groups - Change Password	345
Section 2: Software Version Logger Administration	345
System - System Locale	345
System - SMTP Settings	346
System - Process Status	346
Logs - Audit Logs	348
Logs - Audit Forwarding	348
Users/Groups - User Management	349
User Groups	349
Managing a User Group	352
Users	354
Users/Groups - Change Password	355
Using a CA-signed Certificate on Software Version of Logger	356
Applying a License on the Software Version of Logger	357
Chapter 8: Managing Connectors	359
Connector Overview	360
Navigating the Manage Connectors Tab	361
Locations	363
Viewing All Locations	363
Viewing Hosts, Containers, and Connectors in a Location	363
Adding a Location	364
Exporting and Importing Remote Management Configuration	364
Adding Locations and Hosts from a File	365
Editing a Location	366
Deleting a Location	366
Adding Hosts to a Location	366
Hosts	367
Viewing All Hosts	367
Viewing Containers and Connectors in a Host	367
Adding a Host	368
Scanning a Host	370
Deleting a Host	372
Moving a Host to a Different Location	372

Editing a Host	372
Upgrading a Host Remotely	372
Adding a Container to a Host	373
Containers	374
Viewing All Containers	374
Viewing Connectors in a Container	375
Adding a Container	375
Adding a Connector to a Container	375
Editing a Container	375
Deleting a Container	376
Updating Container Properties	376
Changing Container Credentials	377
Enabling and Disabling FIPS on a Container	378
Managing Certificates on a Container	379
Enabling or Disabling a Demo Certificate on a Container	379
Adding CA Certificates on a Container	380
Adding a CA Certs File on a Container	381
Removing CA Certificates from a Container	382
Viewing Certificates on a Container	383
Resolving Invalid Certificate Errors	385
Running a Command on a Container	385
Upgrading a Container to a Specific Connector Version	386
Viewing Container Logs	387
Deleting Container Logs	387
Running Logfu on a Container	388
Running Diagnostics on a Container	389
Connectors	390
Viewing all Connectors	390
Adding a Connector	390
Editing Connector Parameters	394
Updating Simple Parameters for a Specific Connector	394
Updating Table Parameters for a Specific Connector	396
Updating Simple and Table Parameters for Multiple Connectors	397
Managing Destinations	398
Adding a Primary Destination to a Specific Connector	398
Adding a Failover Destination to a Specific Connector	401
Adding a Primary or Failover Destination to Multiple Connectors	402
Removing Destinations	403
Re-Registering Destinations	404
Editing Destination Parameters	405
Editing Destination Runtime Parameters	407
Managing Alternate Configurations	409
Sending a Command to a Destination	411

Removing a Connector	412
Sending a Command to a Connector	413
Running Logfu on a Connector	414
Changing the Network Interface Address for Events	414
Developing FlexConnectors	415
Editing FlexConnectors	418
Sharing Connectors (ArcExchange)	419
Packaging and Uploading Connectors	419
Downloading Connectors	422
Configuration Suggestions for Connector Types	424
Deploying FlexConnectors	425
Configuring the Check Point OPSEC NG Connector	425
Adding the MS SQL Server JDBC Driver	428
Chapter 9: Managing Repositories	429
Overview	430
Logs Repository	432
Uploading a File to the Logs Repository	432
CA Certs Repository	433
Uploading CA Certificates to the Repository	434
Removing CA Certificates from the Repository	435
Upgrade AUP Repository	436
About the AUP Upgrade Process	436
Uploading an AUP Upgrade File to the Repository	436
Removing a Connector Upgrade from the Repository	437
Content AUP Repository	437
Applying a New Content AUP	438
Applying an Older Content AUP	439
Remote Management AUP Repository	439
Downloading Remote Management AUP Files	440
Uploading Remote Management AUP Files	440
Deleting Remote Management AUP Files	440
Emergency Restore	441
User-Defined Repositories	442
Creating a User-Defined Repository	442
Retrieving Container Files	444
Uploading Files to a Repository	444
Deleting a Repository	445
Updating Repository Settings	445
Managing Files in a Repository	446
Retrieving a File from the Repository	446
Uploading a File from the Repository	446
Pre-Defined Repositories	447

Settings for Backup Files	447
Settings for Map Files	448
Settings for Parser Overrides	449
Settings for FlexConnector Files	450
Settings for Connector Properties	451
Settings for JDBC Drivers	452
Cloning Container Configuration	453
Adding Parser Overrides	454
Appendix A: Common Event Format	455
Common Exchange Format	455
Common Extension Dictionary	457
Appendix B: Regular Expressions	461
Regex Overview	461
Simple Regular Expressions	462
Metacharacters	462
Forbidden Characters	467
Things To Remember	468
Appendix C: Using the Rex Operator	469
Syntax of the rex Operator	469
Understanding the rex Operator Syntax	469
Ways to Create a rex Expression	470
Creating a rex Expression Manually	470
Samples of rex Expressions	471
Appendix D: Restoring Factory Settings	475
Appendix E: Logger Audit Events	483
Types of Audit Events	483
Information in an Audit Event	483
Platform Events	484
Logger Application Events	490
Appendix F: Examples of System Health Events	503
Appendix G: Connector Appliance Documentation	509
Appendix H: Destination Runtime Parameters	511
Appendix I: Event Field Name Mappings	519
Appendix J: Logger Search From An ESM Console	525
Understanding the Integrated Search Functionality	525
Prerequisites	526

Setup and Configuration	527
ESM	527
Logger	527
Supported Search Options	528
Guidelines	528
Searching on Logger From ESM Console	528
Index	531

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 15](#)

["Logger Features" on page 17](#)

["Deployment Scenarios" on page 20](#)

["What's New in Logger v5.1" on page 22](#)

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. The software-based solution is similar in feature and functionality to the appliance-based solution, however, the software solution enables you to install ArcSight Logger on a supported platform of your choice.

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

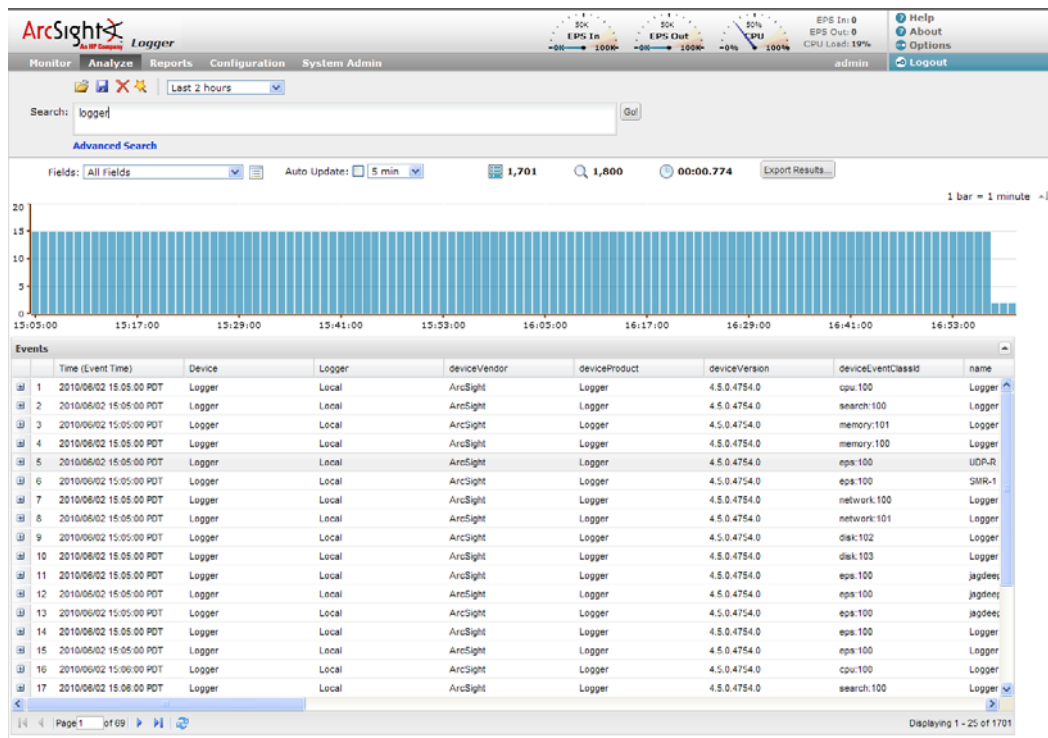


Figure 1-1 Logger web interface, Analyze tab

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 network, with search queries distributed across all 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 455](#) 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 280](#)
- ["Common Event Format" on page 455](#)

Logger Features

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

Storage Configuration

The **Logger appliance** includes onboard storage for events. Some Logger models include RAID 1 or RAID 5 storage systems. (See the *Appliance Specifications* document for more details. This document is available from the ArcSight Download Center at <https://arcsight.subscribenet.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 33](#) to store events. SAN is not supported for storage on software Loggers. The disk space needs to be on the partition where the `<install_dir>` directory exists. Specifically, most of this space should be available for the `<install_dir>/data/logger` directory. 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 27](#)
- [“Initializing the Logger Appliance” on page 27](#)
- [“Storage” on page 233](#)

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 software 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 239](#)
- [“Installing SmartConnectors to Send Events to Logger” on page 54](#)
- [“Sending Events from ArcSight ESM to Logger” on page 57](#)

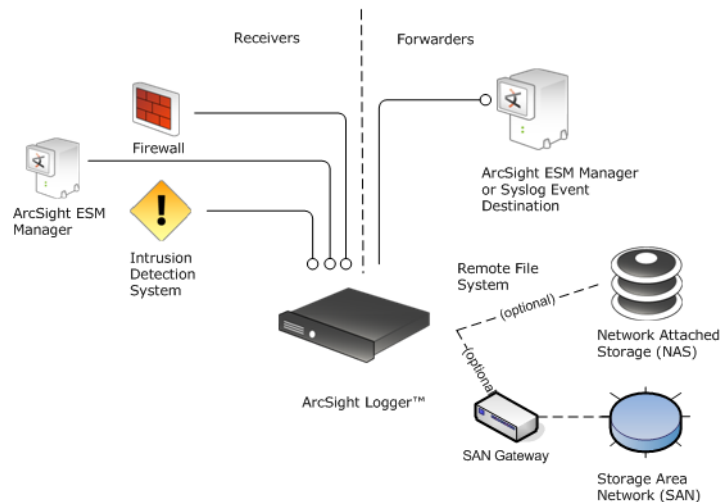


Figure 1-2 Logger appliance has multiple network interface cards (NICs) so that it can receive events on one subnet and forward events on a different subnet.

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 Loggers are configured. However, you can configure it to distribute a query across 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 110](#)
- [“Saving Queries \(Saved Filters and Searches\)” on page 123](#)
- [“Filters” on page 270](#)
- [“Saved Searches” on page 273](#)
- [“Peer Loggers” on page 280](#)

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 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 223](#)
- [“Storage Rules” on page 235](#)
- [“Searching Peer Loggers \(Distributed Search\)” on page 111](#)

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 `<install_dir>/data/logger` directory) or to the browser from which you connect to the Logger. The `<install_dir>/data/logger` directory can be mounted to an NFS or CIFS, 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 116](#)
- [“Impact of Daylight Savings Time Change on Logger Operations” on page 307](#)
- [“Scheduled Saved Search” on page 275](#)

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 246](#)
- [“ESM Destinations” on page 251](#)

User Management

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

- [“Users/Groups” on page 333](#)
- [“Users/Groups - Change Password” on page 345](#)
- [“Search Group Filters” on page 272](#)

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. For the **software version of Logger**, typically, Logger service and related processes need to be restarted. Follow instructions in [“Starting and Stopping the Software Logger” on page 50](#) to start, stop, or restart Logger service on a software Logger.

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

- [“Configuration Backup and Restore” on page 284](#)
- [“Retrieve Logs” on page 296](#)
- [“Storage” on page 314](#)
- [“System Locale” on page 302](#)
- [“License & Update” on page 309](#)
- [“Network” on page 304](#)

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.



Note

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.

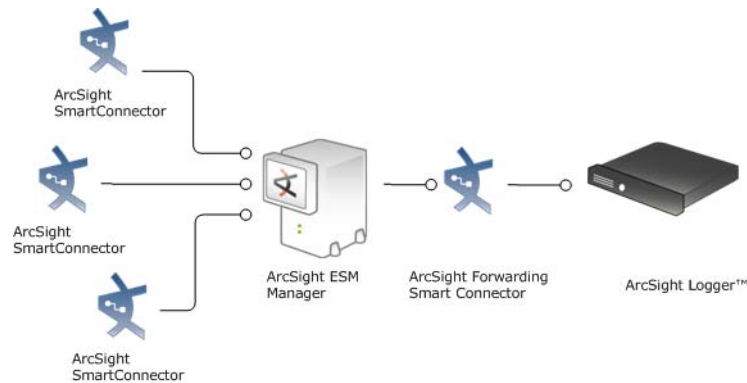


Figure 1-4 Logger can store events sent by ArcSight ESM.




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

What's New in Logger v5.1

Please see the release notes for this release that are available at the ArcSight Customer Support web site at <http://www.arcsight.com/supportportal>.

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 23](#)
- [“Section 2: Installing the Software Version of Logger” on page 33](#)

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

- [“Installing SmartConnectors to Send Events to Logger” on page 54](#)
- [“Sending Events from ArcSight ESM to Logger” on page 57](#)

Section 1: Installing the Logger Appliance

Installing the Logger appliance includes these steps:

- 1 Installing the appliance, as described in [“Installing the Logger Appliance” on page 23](#).
- 2 Setting an IP address, as described in [“Setting an IP Address on the Appliance” on page 24](#)
- 3 Initializing the appliance, as described in [“Initializing the Logger Appliance” on page 27](#).
- 4 Configuring your Logger so it can be accessed remotely as described in [“Configure Logger for Remote Access” on page 33](#).

Installing the Logger Appliance

The ArcSight Logger appliance package includes detailed rack installation instructions, which you can use to rack mount your appliance. Refer to the enclosed *Getting Started with ArcSight Logger Appliance* to configure an IP address on your Logger appliance, install a license file, and connect to it the first time using the IP address you configured.

Safety Precautions

Read carefully through the instructions, cautions, and warnings that are included with the appliance shipment.

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



The following instructions are also included in the *Getting Started with ArcSight Logger Appliance* document that is enclosed with Logger shipment.

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.

Using a Browser to Set an IP Address

- 1 Open a modern, Flash-enabled browser. 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. 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.

Using the Command Line Interface to Configure IP Address

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`.

Other CLI Commands

The following commands are available at the CLI prompt:

Category	Command	Description
System Commands	exit	Logout
	halt	Stop and power down the Logger appliance
	reboot	Reboot the Logger appliance
Admin	show admin	Show the default administrator user's name
Config	show config	Show host name, IP address, DNS, and default gateway for this Logger
Date	show date	Show the currently configured date on the Logger
Default Gateway	show defaultgw [nic]	Display the default gateway for all or the specified network interface
	set defaultgw <IP> [nic]	Set the default gateway for one or all network interfaces
DNS	show dns	Show the currently configured DNS servers on the Logger
	set dns <dn1> [, <dn2>] [, <dn3>] ns1 [, ns2]	Set DNS name server(s). dn=search domain name, ns=nameserver
Hostname	show hostname	Show the currently configured hostname on the Logger
	set hostname <host>	Set Logger's host name
IP	show ip [nic]	Show the IP addresses of all or the specified network interface
	set ip <nic> <IP> [/prefix] [netmask]	Set Logger's IP address for a specific network interface
Password	set password	Set the password the current user's account
SSL Certificate	show sslcert	Show the currently loaded SSL certificate on Logger
	reset sslcert	Install and restart the HTTPS server with the default, self-signed certificate that Logger shipped with
	restart sslcert	Restart the HTTPS server
	diag sslcert	Display the SSL session information
Status	show status	Show the Logger configuration

Initializing the Logger Appliance

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
- 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.

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

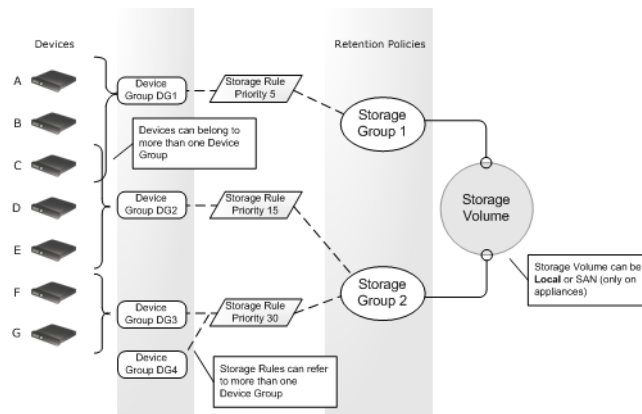
You can also configure the Logger to read event data or log files from a CIFS host.

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.



The previous figure illustrates the relationship between ArcSight components and retention policies. Devices, on the left, are grouped by Device Groups. Storage Groups implement different retention policies on the Storage Volume. Storage Rules, in the middle, create a mapping between Device Groups and Storage Groups. In the example shown, Device C is a member of both Device Group 1 and Device Group 2. Storage Rules are defined that send Device Group 1 events to Storage Group 1 and Device Group 2 events to Storage Group 2.

There is no ambiguity, however, because each Storage Rule has a unique priority value, and the lower value has the higher priority. In the example, events from Device C are stored in Storage Group 1 because that Storage Rule has a priority of 5, which is lower than the other matching Storage Rule, which has a priority of 15.



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

Initialization Sequence for Logger appliances

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 475](#)). Be sure you know how you want Logger storage set up before performing the first steps of the initialization sequence (up to rebooting).

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

- 1 License
- 2 SAN (on selected Logger appliance models)
- 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



When you connect to the Logger appliance for the first time, an End User License Agreement screen prompts you to read through the licensing information and accept the agreement. Once you accept, you can proceed further.

1 License

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 “ArcSight Appliance License & System Update” page and the “Change Password” page (accessible from the

System Admin menu option) are available on it. You cannot use any of the Logger system administration or 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 each appliance has a license file with its unique serial number installed. 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.

To apply a license file on a Logger:

- a** Download the license file from the ArcSight software download site at <https://arcsight.subscribenet.com> to a computer from which you can connect to Logger.
- 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 > **License & Update**.
- d** Browse to the *license* file you downloaded earlier and click **Upload Update**.

Wait until a 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.)

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

When you multipath a LUN, you create two different network paths to it from Logger. Doing so reduces the possibility of a single point of failure causing the LUN to become unavailable. **If you want to configure multipathing on Logger, you must configure it before attaching the LUN.**

See [“SAN” on page 318](#) for detailed information about connecting LUN and multipathing.

3 Storage Volume

Establish the Logger's Storage Volume. See [“Storage Volume” on page 236](#). Choose **Local** to use Logger's built-in storage. 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 can configure up to a 5 TB volume on Logger's that support SAN. To make use of this enhanced capacity, make sure you allocate a 5.4 TB LUN during initialization. The additional 0.4 TB need to be allocated to accommodate Logger system files.

You can choose to pre-allocate your Storage Volume to enhance performance. 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 293](#).
 - Even if your LUN is larger than 5.4 TB in size, you can only allocate a maximum of 5.4 TB and pre-allocate a maximum of 5TB.
 - The size of a LUN cannot be changed after it has been set during Logger initialization. For more information, see [“About Increasing Storage Volume Size on a SAN Logger” on page 294](#).
-

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.** If you are decreasing the size of the storage group and the new size is less than the currently used space on the storage group, you will need to delete data to achieve the new size. Logger UI guides you in this situation to delete sufficient data. See [“Storage Groups” on page 233](#) for the details of adding and resizing Storage Groups.

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

4 Time Settings

Configure the system time manually. Follow instructions in [“Time/NTP” on page 306](#).

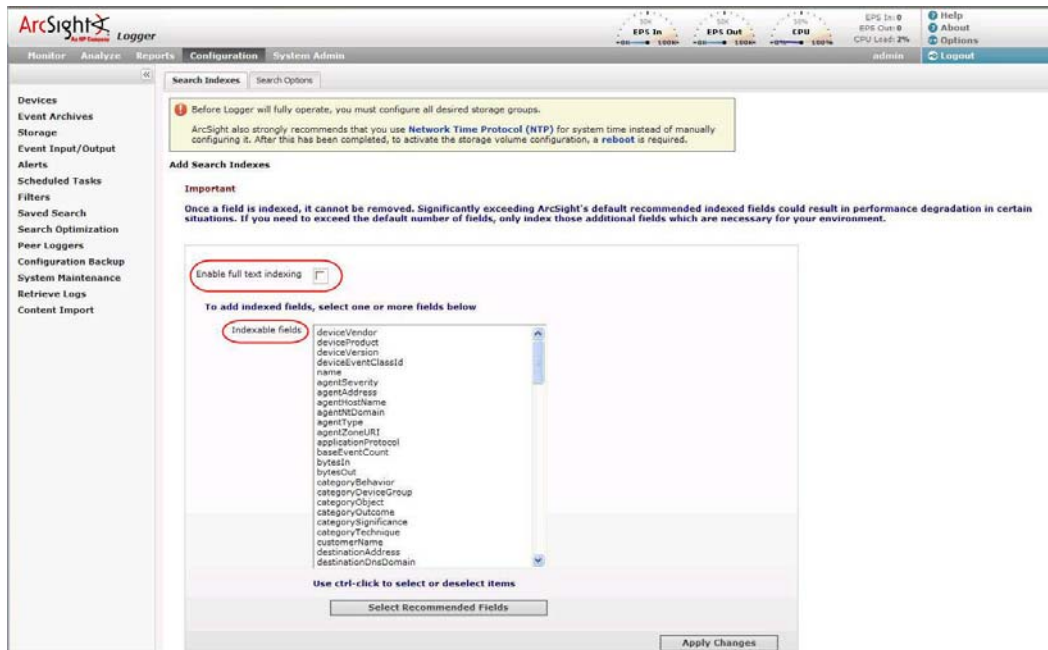
Optionally, configure NTP time settings. Configuring an NTP server will ensure precise time stamping of events, which is a key log management function. ArcSight strongly recommends that you use Network Time Protocol (NTP) for system time. See [“Time/NTP” on page 306](#) for more information.

5 Index Fields and Full-text Indexing

As shown in the following figure, 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 119](#) for more information.

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



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 Locale

The Locale setting ensures that the user interface displays information such as date, time, numbers, and messages in the format and language appropriate for the selected country. For more information about setting up Locale on your Logger appliance, see [“System Locale” on page 302](#).

7 Reboot

After the storage, indexing, and locale settings 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 302](#).



When Logger is rebooted, the storage, indexing, and locale 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 293](#). Additional fields can be added to the index, but not removed. Once the locale is set, it cannot be changed.

8 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 239](#).

9 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 223](#).

10 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 225](#).

11 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 235](#) 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.

Configure Logger for Remote Access

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

If your Logger appliance **is an** Lx4xx model, it is equipped with an HP ProLiant Integrated Lights-Out (iLO) Advanced remote management card. Follow the directions in the *HP ProLiant Integrated Lights-Out User Guide* to set up your appliance for remote access. The guide is available at <http://www.hp.com/go/iLO>.

If your Logger appliance **is not** an Lx4xx model, contact ArcSight Customer Support for assistance in setting up remote access for your appliance.

Section 2: Installing the Software Version of Logger

The information in this section pertains to installing the software version of 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 Download Center at <https://arcsight.subscribenet.com>.

A VM installation of the operating systems listed in the table below is supported. ArcSight strongly recommends allocating 4 GB RAM per VM instance. Additionally, the sum of memory configurations of the active VMs on a VM server must not exceed the total physical memory on the server.

Specification	Details
Certified Operating Systems	<ul style="list-style-type: none"> Red Hat Enterprise Linux (RHEL), version 5.4 and 5.5, 64-bit Oracle Enterprise Linux (OEL) 5.4, 64-bit CentOS, versions 5.4 and 5.5, 64-bit
Other Supported Operating Systems	<ul style="list-style-type: none"> Red Hat Enterprise Linux (RHEL), version 4.x, 64-bit Oracle Enterprise Linux (OEL) 5.5, 64-bit CentOS, version 4.x, 64-bit

Specification	Details
CPU, Memory, Disk Space	<p>For Small Deployments <i>(Only for ArcSight Logger—Downloadable Version)</i></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: 10 GB (minimum) <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 you will install the Logger software. Using NFS as primary storage for events on the software version of Logger is not recommended. The system on which you are installing the software version of Logger must not have more than two CPUs.
Browsers	<ul style="list-style-type: none"> Internet Explorer: Versions 7 and 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>
Other Applications	For optimal performance, make sure no other applications are running on the system on which you install the software version of Logger.

Downloading the Software

The web site from where you can download the software version of Logger depends on your method of purchase. Use the following table to determine the web site from where you can download the software.

Purchased through...	Download from...
ArcSight Sales group	https://arcsight.subscribenet.com
ArcSight corporate web site (www.arcsight.com)	http://www.arcsight.com/products/products-logger/

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. Once you have downloaded the software, you receive an email from ArcSight that contains the license file you will need to install the software. A license file is uniquely generated for each download; therefore, you cannot use the same license file to install multiple instances of software Logger.

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 limits for the following:

- **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 5.1.3.5870.0 or later 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.

			<div> <div>50K</div> <div>EPS In</div> <div>0% 100%</div> </div> <div> <div>50K</div> <div>EPS Out</div> <div>0% 100%</div> </div> <div> <div>50%</div> <div>CPU</div> <div>0% 100%</div> </div> <div> <div>EPS In: 0</div> <div>EPS Out: 0</div> <div>CPU Load: 31%</div> </div>
			dev_admin
License Information			
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 License on the Software Version of Logger”](#) on page 52 to apply the new license on your Logger.

Deployment Planning for the Software Version of Logger

Storage Strategy

Logger events can be stored locally.

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

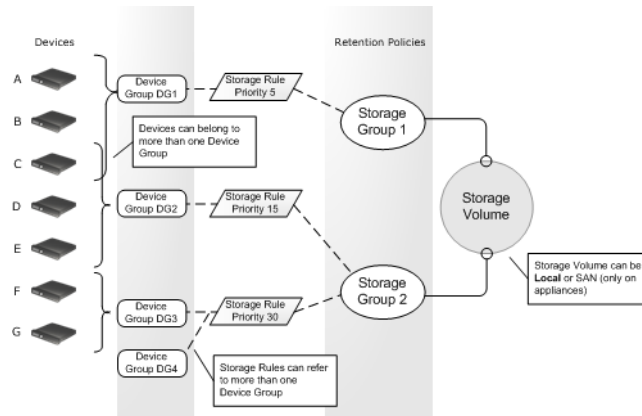
You can also configure the Logger to read event data or log files from a CIFS host.

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.



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



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

Note

Installing and Configuring the Software Version of Logger

Prerequisites for Installation

Make sure these prerequisites are met before you install the software version of Logger:

- You received a license file from ArcSight. The file is attached to an email that you received from ArcSight after downloading the Logger software. You will need this file for installation.
- You can be logged in as a root user or a non-root user on the system on which you are installing the software. When you install the software as a root user, you can select the port on which Logger listens for secure web connections. However, when you install it as a non-root user, Logger can only listen for connections on port 9000. You cannot configure the port to a different value. Additionally, you can configure Logger to start as a service when you install as a root user.

The ability to install as a root user is new in version 5.1; only non-root user installation was supported prior to this release. Therefore, if you are upgrading from a previous version of Logger to 5.1, you cannot change the previous install to a root-user

installation. You will need to use the previously configured port 9000 for accessing software Logger. A non-root user account is still required to complete the Logger installation as a root user. Therefore, make sure a non-root user account exists on the system on which you are installing Logger.

- The hostname of the machine on which you are installing Logger cannot be "localhost". If it is, change the hostname before proceeding with the installation.
- You must not have an instance of MySQL or PostgreSQL installed on the Linux machine on which you will install Logger. If instances of these exist on that machine, uninstall them before proceeding with the installation.
- Ensure that the umask setting in the `/etc/bashrc` file and your local `.bashrc` file has not been modified on the system which you are installing the Logger software. If this setting is modified, Logger might not function as expected.
- If you want to use the GUI mode of installation (described in ["Installation Modes" on page 38](#)) and will be installing Logger software over an SSH connection, make sure that you have enabled X window forwarding using the `-X` option so that you can view the screens of the installation wizard. If you will be using PuTTY, you will also need an X client on the machine from which you are connecting to the Linux machine.

Installation Modes

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

- GUI—In this mode, a wizard steps you through the installation and configuration of software Logger.
- Console—In this mode, a command-line process steps you through the installation and configuration of software Logger.
- Silent—In this mode, you provide the input required for installation and configuration through a file. Therefore, you do not need to interact with the installer to complete the installation and configuration. However, before you can use this mode, you must run the installation and configuration using one of the other modes to record the input in a file.

Installation Steps

This section describes all three modes of software Logger installation.

Using the GUI Mode to Install Software Logger



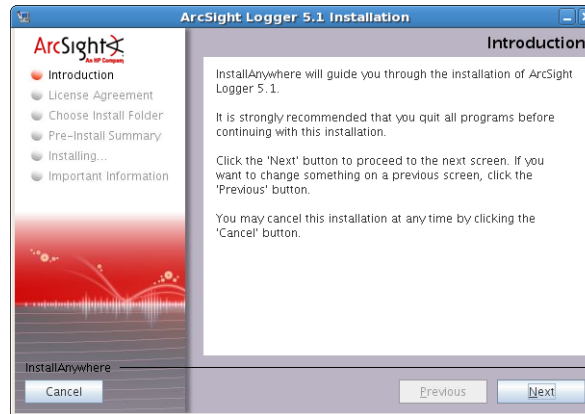
You can install software Logger as a root user or as a non-root user. See ["Prerequisites for Installation" on page 37](#) for details.

To install the software version of Logger using the GUI mode:

- 1 Make sure the machine on which you will be installing the software Logger complies with the requirements listed in ["Supported Platforms and Browsers" on page 33](#) and the prerequisites listed in ["Prerequisites for Installation" on page 37](#) are met.
- 2 Run these commands from the directory where you copied the Logger software:

```
chmod +x ArcSight-logger-5.1.0.XXXX.0.bin
./ArcSight-logger-5.1.0.XXXX.0.bin
```


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



- 3 On the next screen, read through the License Agreement details. Click **I accept the terms of the License Agreement**. Then, click **Next**.

Note: You need to scroll down to the end of the License Agreement details to enable the “I accept the terms of the License Agreement” button.

- 4 If a previous version of Logger is running on this machine, the following message is displayed. Click **Continue** to proceed with the installation of version 5.1 or **Quit** to exit the installation process.



- 5 Specify or browse to a folder where you want to install Logger, as shown in the following figure. By default, the home directory of the user who is logged in is specified.

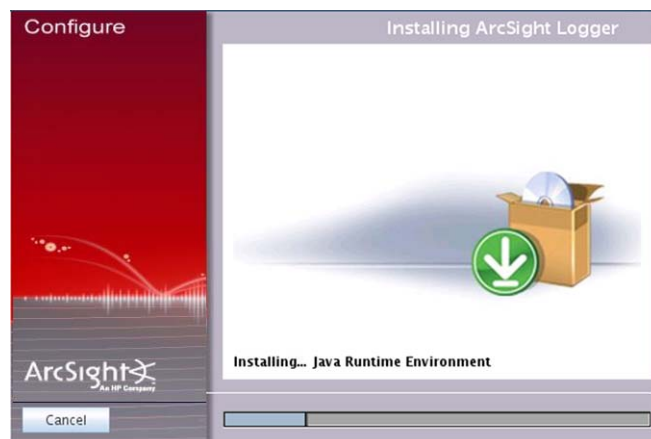


If sufficient space to install the software is not found at the location you specify for installation, the following error message is displayed. You need to either specify a

different location or make sufficient space in the location you specified before installation can proceed.

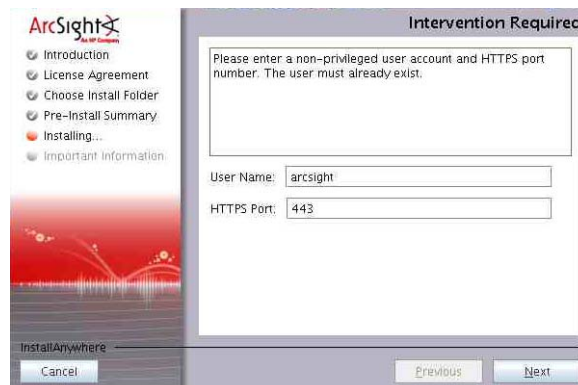


- 6 Review the summary of installation information on the Pre-Installation Summary screen and click **Install**.
- 7 The Logger software begins to install, as shown in the following figure.

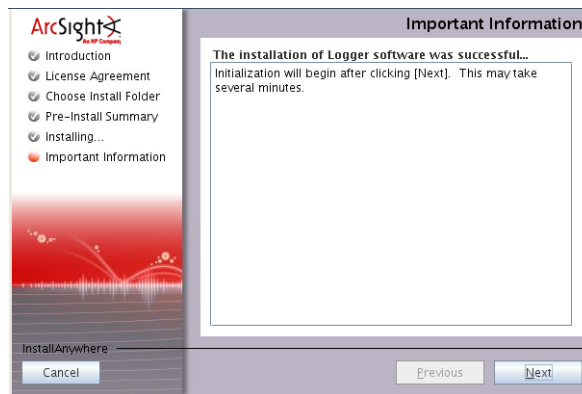


- 8 **If you are logged in as a root user** on the system on which you are installing Logger software, the following screen is displayed next. This screen enables you to specify a non-root user (that must exist on the system already) and configure a port on which Logger users will connect to it through the Logger UI. For example, you can enter 443, the standard HTTPS port, or any other that suits your needs. If any port except 443 is specified, your users will need to enter that port number in the URL they use to access the Logger UI.

Enter the user name of the non-root user and HTTPS port number, and click **Next**.



- 9 Once the software is installed, the following screen is displayed. Click **Next** to begin Logger initialization.




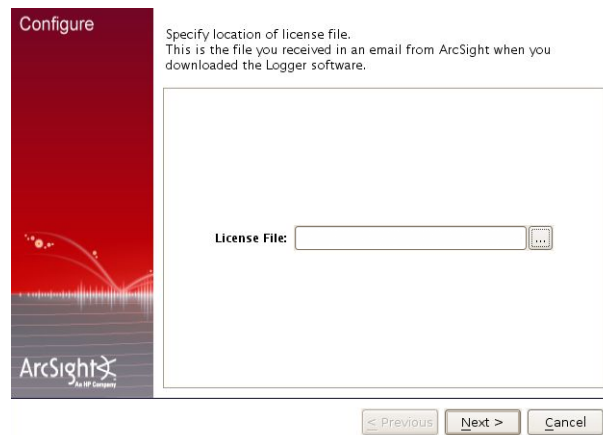
- 10 Once initialization is complete, the following screen is displayed. Click **Done** to launch the Logger Configuration wizard.

Note: The Configuration wizard launches automatically. However, if it does not, use this command to start it:

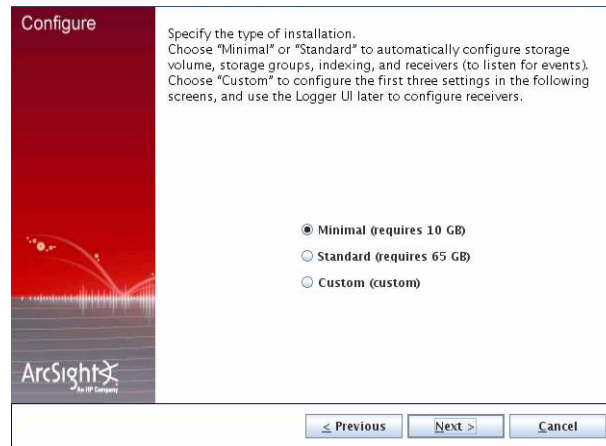
```
<install_dir>/current/arcsight/logger/bin/arcsight loggersetup
```



- 11 In the next screen, click the Browse () icon to locate the license file that you received from ArcSight. Click **Next**.



12 Select the type of installation.



You can select from these choices.

Setting	Minimal	Standard
Minimum Disk Space Required	10 GB	65 GB
Storage Volume	10 GB	65 GB
Storage Groups	2	6
Indexing	Full-text and field-based	Full-text and field-based
Receivers	2	2

- ◆ **Minimal or Standard**—If you choose either of these options, storage volume, storage groups, indexing, and receivers (to listen for events) are automatically configured for you during the installation process using the default values listed in the above table.

The following guidelines provide additional information for the above listed settings:

- **Storage Volume:** If needed, you can increase the storage volume size later if additional disk space is available.
- **Storage Groups:** Only two storage groups can be configured for the Minimal installation type. Once configured, the number of storage groups cannot be increased. Therefore, if you choose Minimal installation, ensure that two storage groups will be sufficient for your needs.
- **Indexing:** All recommended fields are included in the index and full-text indexing is enabled. A complete list of fields is available in the *Logger Administrator's Guide*.
- **Receivers:** A TCP and a UDP receiver is automatically configured and enabled. Make sure your events sources are configured to send events to these receivers on the ports that the configuration wizard displays after this step. You can add more receivers or modify the automatically configured ones later using the Logger user interface.
- ◆ **Custom**—If you choose this option, you need to configure indexing, storage groups, and the storage volume size, as described in the next step. For this type of installation, the receivers are configured using the Logger UI after the installation is complete.

13 If you selected “Custom”:

- a** Enter the maximum size of the Storage Volume. Click **Next**.

The maximum size you enter can only be equal to or less than the aggregated storage size allowed by your license.

Configure

Specify size of Storage Volume.

Storage Volume is where all event data is stored. You can increase the Storage Volume size after it has been created, but you cannot decrease it.

Maximum Size (GB) 50

< Previous Next > Cancel

- b** Add the storage groups to suit your needs. Click the Maximum Age (number of days to retain the event) and Maximum Size fields for each storage group you add and specify values. Click **Next**.

ArcSight recommends that you create the maximum allowed of four additional Storage Groups (in addition to the two that preexist—Default Storage Group and the Internal Storage Group) at this stage even if you do not need all of them because you cannot add storage groups later, although you can decrease or increase the size of a Storage Group at any time. Additionally, if you will not use all storage groups, keep the size of the spare groups to a minimum to optimize space for storage groups that you will use.

Configure

Specify Storage Groups and their retention policies.

- Once created, a storage group cannot be deleted, however its size can be changed.
- A Default Storage Group and an Internal Event Storage Group must always exist on Logger and are created automatically.

Storage Group Name	Maximum Age (Days)	Maximum Size (GB)
Default Storage Group	180	15
Internal Event Storage ...	365	5
Long retention 1	365	8
Long retention 2	365	8
Short retention 1	30	5
Short retention 2	30	5

Add Remove

< Previous Next > Cancel

- c** If you want to enable full-text indexing, click “Enable full text index”.

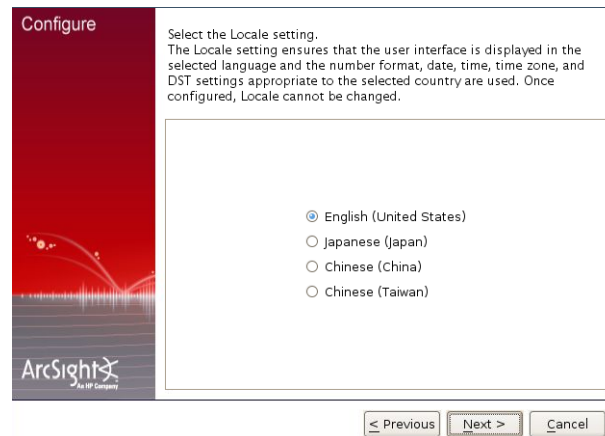
Tip: ArcSight strongly recommends that you enable indexing because indexing significantly improves search and reporting performance. When you add fields to the index, search queries yield significantly faster results.

Select the fields you want to index. You can drag and drop fields from the left column (Indexable Fields) to the right column (Selected Fields). You can also click

“Select Recommended Fields” to select the ArcSight recommended fields with one click and then drag them to the right column. Click **Next**.



- 14** Select the Locale for your Logger, as shown in the following figure. Click **Next**.



- 15** If you are logged in as a root user on the system on which you are installing Logger software, the following screen enables you to configure software Logger to run as a system service. By default, software Logger runs as a standalone application, which you need to launch manually after each system reboot.

When you install software Logger as a root user, a service called `arcsight_logger` can be configured created and enabled at run levels 2, 3, 4, and 5. Additionally, a few libraries are added using `ldconfig`. For a complete list of those libraries, see

`/etc/ld.so.conf.d/arcsight_logger.conf` and
`<install_dir>/current/arcsight/install/ldconfig.out`.




16 You have installed Logger. Click **Start Logger Now**.

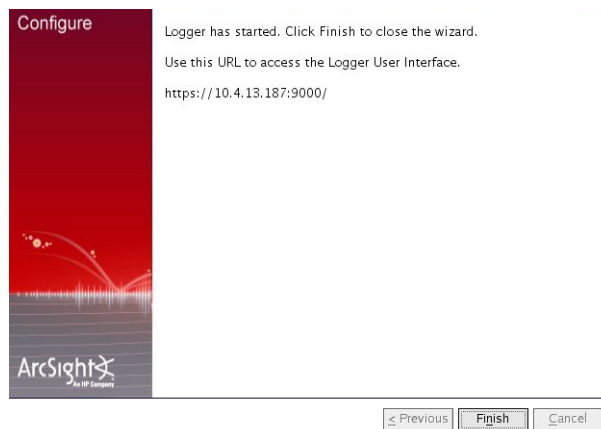
OR click **Start Logger later** and click **Finish**.

If you selected to start Logger later, you need to read the [“Starting and Stopping the Software Logger” on page 50](#) information to understand how to start Logger later.

17 If you selected “Start Logger Now”, the following screen is displayed.



You can click **Finish** to exit the wizard. Logger continues to start service and processes in the background. Once Logger service and processes have started up, the following screen is displayed.



Follow the instructions on the above screen or use instructions in [“Connecting to the Software Logger User Interface” on page 49](#) to connect to the Logger.

Using the Console Mode to Install Software Logger



You can install software Logger as a root user or as a non-root user. See [“Prerequisites for Installation” on page 37](#) for details.

To install the software version of Logger using the Console mode:

- 1 Make sure the machine on which you will be installing the software Logger complies with the requirements listed in [“Supported Platforms and Browsers” on page 33](#) and the prerequisites listed in [“Prerequisites for Installation” on page 37](#) are met.
- 2 Run these commands from the directory where you copied the Logger software:

```
chmod +x ArcSight-logger-5.1.0.XXXX.0.bin
./ArcSight-logger-5.1.0.XXXX.0.bin -i console
```

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

```
Introduction
-----
```

```
InstallAnywhere will guide you through the installation of
ArcSight Logger 5.1.
```

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

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

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

```
PRESS <ENTER> TO CONTINUE:
```

```
....
```

```
Once at the license information is displayed, press Enter until you see the following
information:
```

```
Select "I accept the terms of the License Agreement" below if
you recognize
that you have read the terms of this Agreement and attachments
and agree to be
bound by each of these terms.
```

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

The prompts that follow next are exactly similar to the ones described for the GUI mode install in [“Using the GUI Mode to Install Software Logger” on page 38](#). Follow the instructions provided for the GUI mode install to complete the installation.

Using the Silent Mode to Install Software Logger

Before you install software Logger in silent mode, you need to create two properties files that are required for the silent mode installation:

- A file to capture the installation properties
- A file to capture the configuration properties

Once you have generated the two files, you need to merge them into one file and use the resulting file for silent mode installations.

About Licenses for Silent Mode Installations

As for any Logger installation, each silent mode installation requires a unique license file. You must obtain licenses from ArcSight Customer Support and place them on the machines on which you will be installing Logger in silent mode, or ensure that the location where the license is placed is accessible from those machines.

Generating the Silent Install Properties File

This procedure generates the two properties files first and later instructs you to combine them into one file. The resulting file is used for future silent installations.

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

If you want the silent mode installations to be done as root user, make sure you are logged in as root in this step. Otherwise, log in as a non-root user.

- 2 Run this command:

```
./ArcSight-logger-5.1.0.XXXX.0.bin -r <directory_location>
```

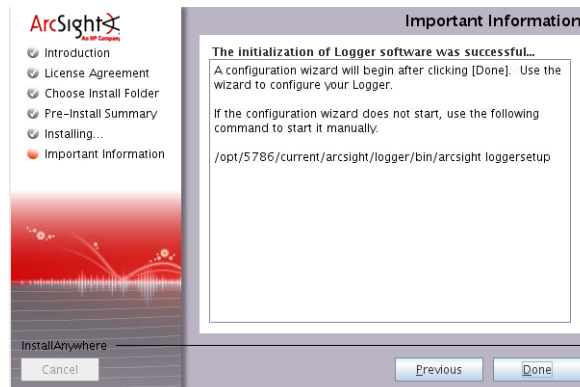
where `<directory_location>` is the location of the directory where the generated properties file will be placed.

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

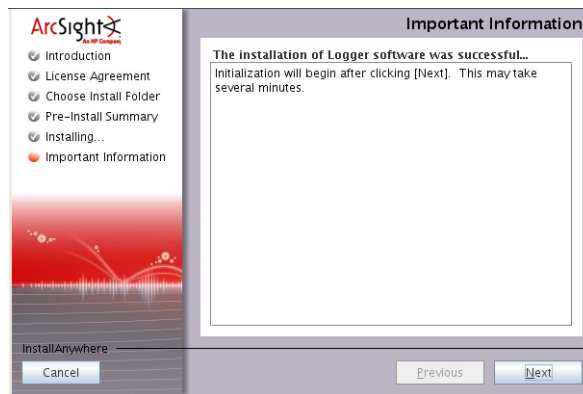
- 3 Install Logger in GUI mode, as described in [“Using the GUI Mode to Install Software Logger” on page 38](#).

At [Step 10 on page 41](#) of the installation procedure, do this:

- a Click **Previous** instead of clicking Done to proceed further, as shown in the following figure.



- b** Once you are at the previous step, click **Cancel** to stop the installation, as shown in the following figure.



- 4** Navigate to the directory location you specified for the `installer.properties` file earlier.

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

```
#####sample file#####
# Thu Mar 17 19:38:56 PDT 2011
# Replay feature output
# -----
# This file was built by the Replay feature of InstallAnywhere. It contains
variables that #were set by Panels, Consoles or Custom Code.
#Choose Install Folder
#-----
USER_INSTALL_DIR=/opt/softlogger
#Install
#-----
fileOverwrite_/opt/softlogger/UninstallerData/Uninstall_ArcSight_Logger_5.1.lax=Yes
#Intervention Required
#-----
USER_AND_PORT_1=arcsight
USER_AND_PORT_2=9000
#####
```

- 5** Start the configuration wizard with the option to record configuration properties:

```
<install_dir>/current/arcsight/logger/bin/arcsight loggersetup
-i recorderui
```

When prompted to enter a file name to capture the configuration properties, enter a meaningful name; for example, `silentproperties.properties`.

- 6** Step through the configuration wizard, as described starting at [Step 10 on page 41](#) of the [“Using the GUI Mode to Install Software Logger” on page 38](#).
- 7** Once the configuration properties file has been generated, append the contents of this file to the `installer.properties` file generated in the previous procedure, [“Generating the Silent Install Properties File” on page 47](#), to create a combined file.

For example, you can use the `cat` command to concatenate both files:

```
cat installer.properties silentproperties.properties >
combinedProperties.properties
```

- 8** Include the following property in the combined file:

```
ARCSIGHT_LOGGER_SETUP_PROPERTIES=/<directory>/<combined_propert
ies_file>
```

where `<directory>` is the path of the directory where the combined file is located.

`<combined_properties_file>` is the file name of the combined file you created earlier.

Use the combined file for future Logger silent mode installations, as described in [“Installing Software Logger in Silent Mode” on page 49](#).

Installing Software Logger in Silent Mode

To install the software version of Logger using the Silent mode:

- 1 Make sure the machine on which you will be installing the software Logger complies with the requirements listed in [“Supported Platforms and Browsers” on page 33](#) and the prerequisites listed in [“Prerequisites for Installation” on page 37](#) are met.
- 2 A non-root user account, “arcsight”, exists on the machine on which you are installing software Logger. This account is required even if you will be installing as a root user.
- 3 Copy the silent mode properties file you generated previously to the same location where you have copied the Logger software.
- 4 Edit the `licensePanel.path` property in the silent mode properties file to include the location of license file for this instance of installation. (A unique license file is required for each instance of installation.)

OR

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

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

```
chmod +x ArcSight-logger-5.1.0.XXXX.0.bin
./ArcSight-logger-5.1.0.XXXX.0.bin -i SILENT -f
<combined_properties_file>
```

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

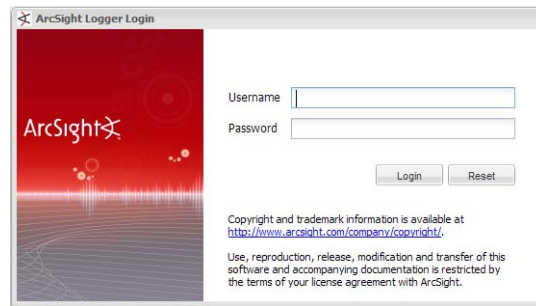
Connecting to the Software Logger User Interface

Because the software Logger user interface uses SSL, make sure you connect to it using this URL:

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

where `hostname` or `IP address` is of the system on which you installed Logger software.

Once you use the URL specified above, the following Login screen is displayed.



Use the following default credentials if you are connecting for the first time:

Username: `admin`
Password: `password`



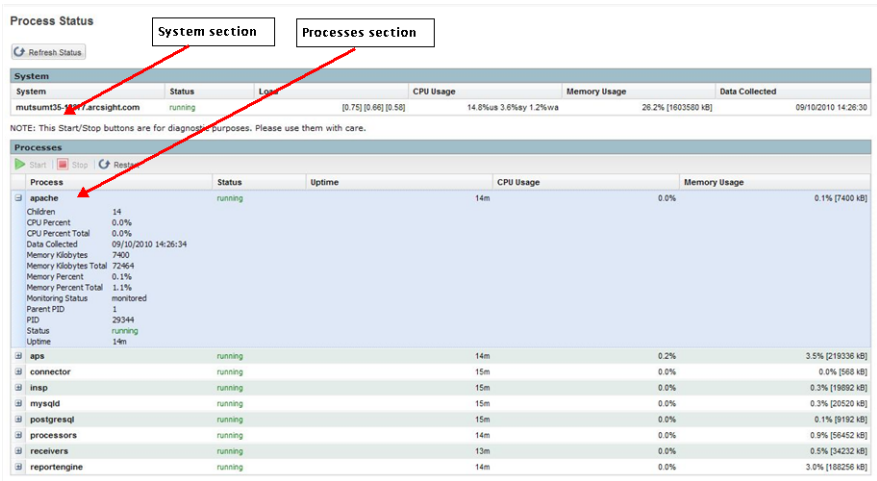
Change the credentials as soon as possible after connecting to your Logger for the first time.

Starting and Stopping the Software Logger

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

```
<install_dir>/current/arcsight/logger/bin/loggerd  
{start|stop|restart|status|quit}  
  
<install_dir>/current/arcsight/logger/bin/loggerd {start  
<process_name> | stop <process_name> | restart <process_name>}  
  
/etc/init.d/service arcsight_logger {start | stop | status}
```

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



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

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

Uninstalling the Software Logger



Note

If you will be uninstalling Logger software over an SSH connection, make sure that you have enabled X window forwarding using the -X option so that you can view the screens of the uninstall wizard.

If you will be using PuTTY, you will also need an X client on the machine from which you are connecting to the Linux machine.

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

```
./UninstallerData/Uninstall_Arcsight_Logger_5.1
```

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

Applying a License on the Software Version of Logger


To apply a license on the software version of Logger:

- 1 Save the license file you receive from ArcSight on a computer from which you access the Logger user interface through a browser.
- 2 Run this command:


```
<install_dir>/current/arcsight/logger/bin/arcsight loggersetup
```

The following Update License screen is displayed. If you are logged in as a root user, the following screen also provides an option to enable or disable Logger from starting as a system service. For more information about configuring Logger to start as a system service, see [“Enabling or Disabling Logger as a System Service”](#) on page 53.



- 3 Click the Browse () icon to locate the license file that you received from ArcSight. Click **OK**.

- 4 Click **Finish**. A message on the screen confirms that the license was applied, as shown in the following figure.



- 5 Restart the Logger service and related processes after applying the license. Use this command to restart the service and processes:

```
<install_dir>/current/arcsight/logger/bin/loggerd restart
```

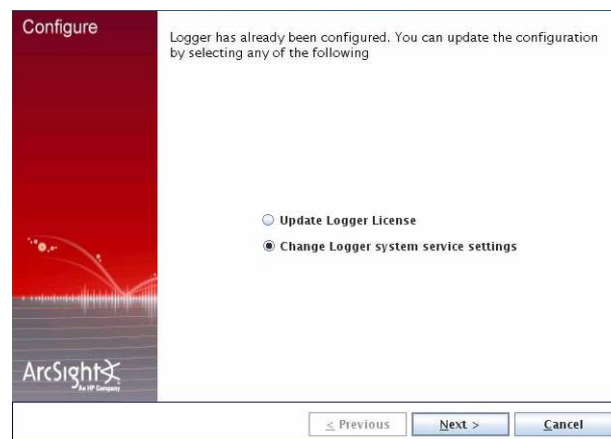
Enabling or Disabling Logger as a System Service

If you want to disable Logger from starting as a system service, or if you want to enable it to start as a system service after it has been installed, follow these steps.

- 1 Make sure you are logged in as a root user on the system on which Logger is installed.
- 2 Run this command:

```
<install_dir>/current/arcsight/logger/bin/arcsight loggersetup
```

The following screen is displayed.



- 3 Select **Change Logger system service settings** and click **Next**.
- 4 If Logger is currently installed as a service, the next screen provides you the option to disable it. Conversely, if Logger is currently installed as a standalone application, you can configure it to run as a service. Click **Finish** and reboot the Logger for changes to take effect.

Best Practices for the Minimal Install

In general, you should use the “Minimal” install for evaluation purposes or if you do not have a server with 65 GB of storage space available, which is the storage space required for a “Standard” install. To truly experience the power of Logger, ArcSight recommends installing software Logger using the “Standard” or “Custom” installation type. Doing so creates a storage system infrastructure on Logger that can handle large data storage needs, such as published reports, large amount of exported data, event archives, and so on.

However, if you must install Logger using the “Minimal” installation type, remember that Logger does not manage disk space availability for published reports, exported data, or event archives. You must ensure that there is sufficient disk space available for these functions. A Logger installed using the “Minimal” installation type can quickly run out of space if you publish a large number (or large) reports, export a large amount of data, or archive events for a long period.

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.

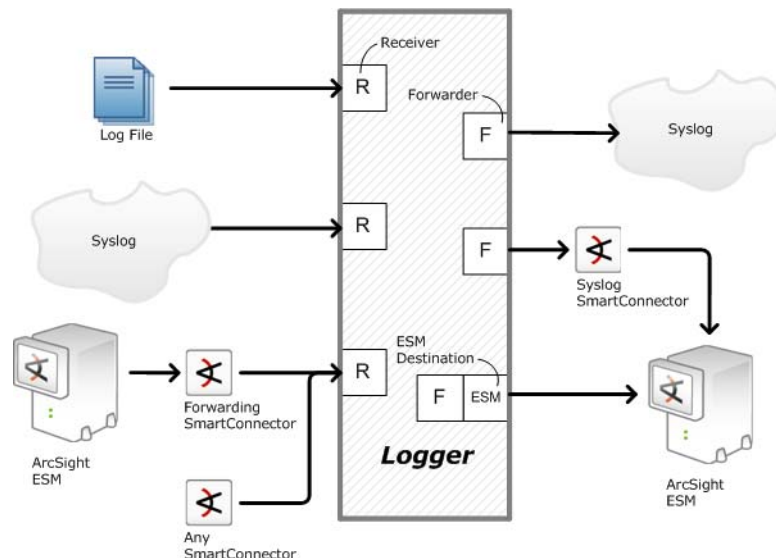


Figure 2-2 ArcSight SmartConnectors interact in a variety of ways with Logger Receivers (R) and Forwarders (F).

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

Downloading SmartConnectors

For Logger appliance, contact your ArcSight sales representative or ArcSight Customer Support for location to download SmartConnectors.

For software Logger, the SmartConnector software is available at the same location from where you downloaded the Logger software.

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.

To use SmartMessage to communicate between an ArcSight SmartConnector and a **Logger appliance**, configure the SmartConnector to use port 443. To communicate between an ArcSight SmartConnector and the **software version of Logger** using SmartMessage, configure the SmartConnector to use the port that you configured for the software Logger.

Set up the SmartMessage Receiver on Logger first (see [“Receivers” on page 239](#)) 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.

If you have “Logger 5.0 - Downloadable Version” installed, which you downloaded from <http://www.arcsight.com>, refer to the Syslog SmartConnector documentation available at the web site from where you downloaded the Logger software.



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. To communicate with a Logger appliance, configure port 443. To communicate with the software version of Logger, configure the port you have configured for it.

(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 455](#).

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 2-3](#)). 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 2-3 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
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 56.

For more information about the Common Event Format (CEF), see [“Common Event Format”](#) on page 455.

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 [“Connecting to the Logger User Interface” on page 61](#)

Performance monitoring: see [“Monitor” on page 66](#)

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

Connecting to the 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.

On the Logger appliance, enter the following URL from any of the supported browsers:

`https://Hostname or IP address of Logger appliance`

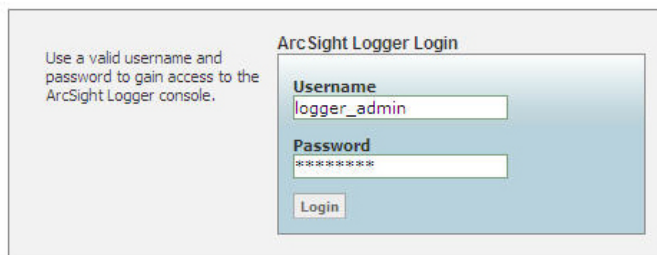


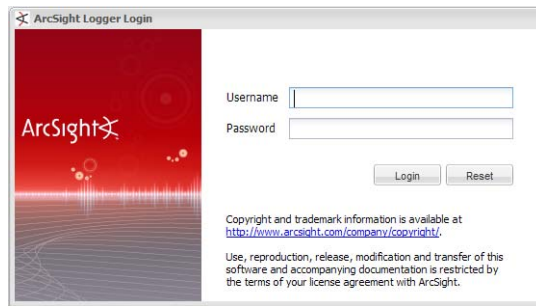
Figure 3-1 Users must login to authenticate themselves to Logger.

On the software version of Logger, make sure you connect using this URL:

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

where `hostname` or `IP address` is of the system on which you installed Logger software.

Once you use the URL specified above, the following Login screen is displayed.



Use the following default credentials if you are connecting for the first time:

Username: `admin`
Password: `password`



Change the credentials as soon as possible after connecting to your Logger for the first time.

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.

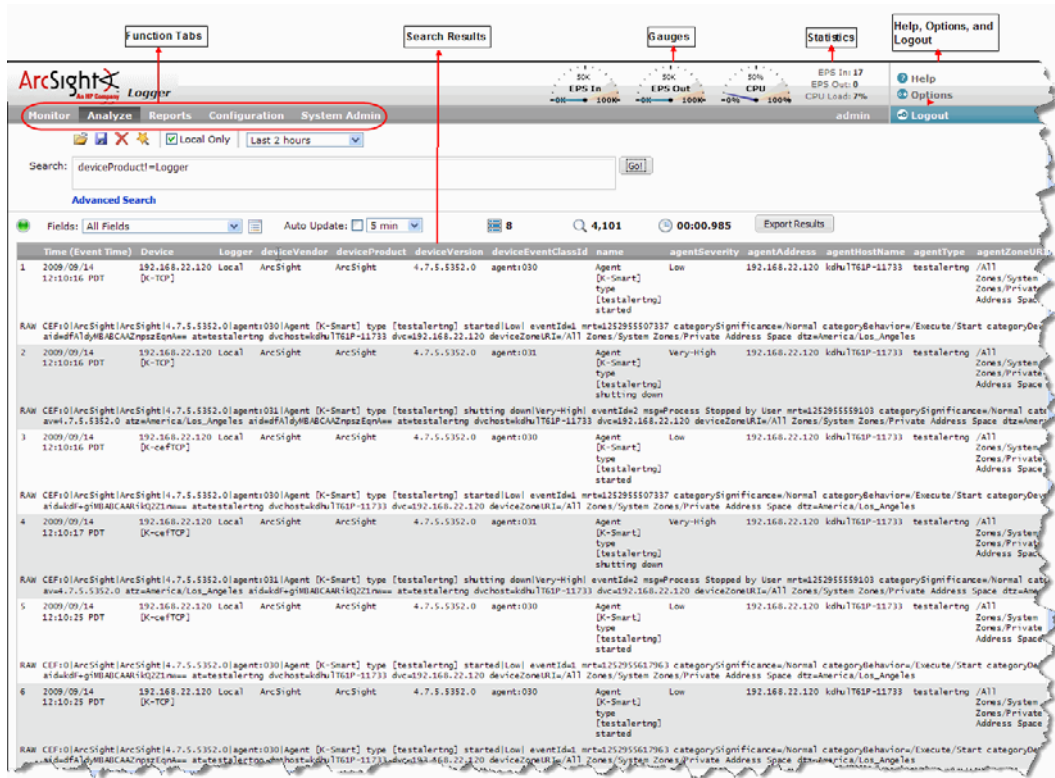


Figure 3-2 Overall layout of the Logger interface.

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.

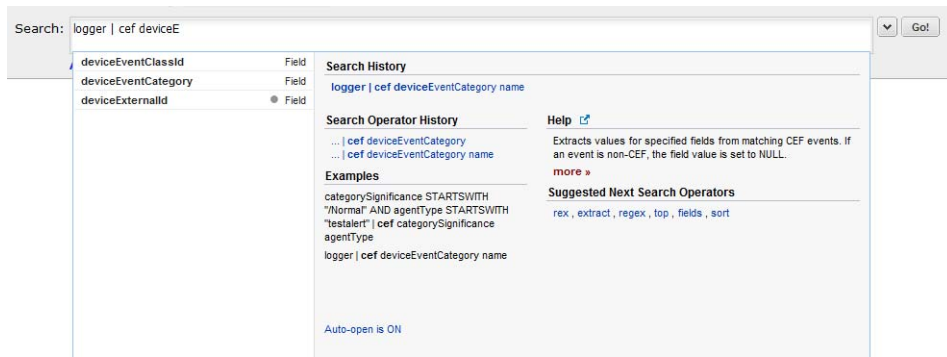
The menu list in the upper right includes links for Help, Options, and Logout.

Help

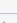
Clicking the Help link on any page displays online help for the current page.

In addition to context-sensitive Help for the current page, you can also access the PDF version of this guide from the Help link. To access the guide, click the "PDF version of Logger Documentation" in the left panel of the Help window.

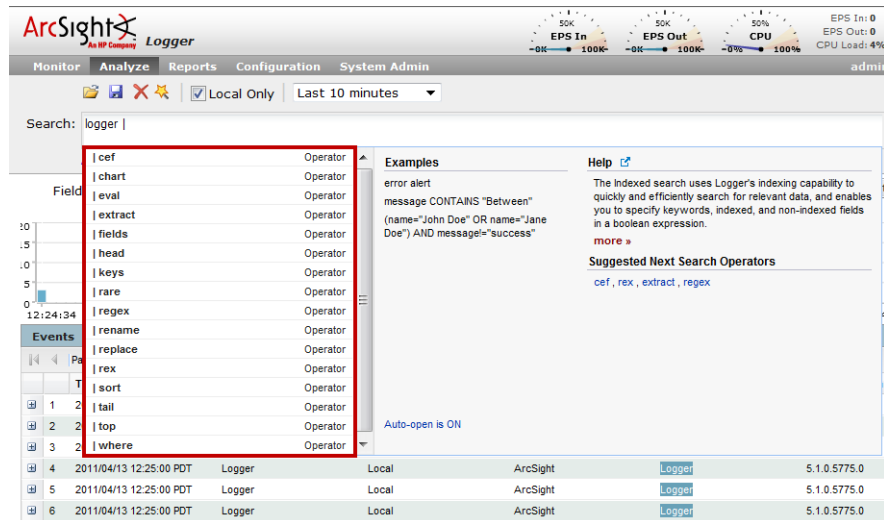
Search Helper



Search Helper is a search-specific utility that provides the following:

- **Search History**—Displays the recently run queries on Logger, thus enabling you to select and reuse previously run queries without typing them again.
- **Search Operator History**—Displays the fields used previously with the search operator that is currently typed in the Search text box. The Search Operator History only displays if you have previously used the operator you have currently typed to perform searches on this Logger.
- **Examples**—Lists examples relevant to the latest query operator you have typed in the Search text box.
- **Suggested Next Operators**—List of operators that generally follow the currently typed query. For example, if you type `logger |`, the operators that often follow are `cef`, `rex`, `extract`, or `regex`. You can select one of the listed operators to automatically append to the currently typed query in the Search text box. This list saves you from guessing the next possible operators and manually typing them in.
- **Help**—Provides context-sensitive help for the last-listed operator in the query that is currently typed in the Search text box. Additionally, if you click the  icon, Logger online Help is launched.
- **List of Fields and Operators**—Depending on the current query in the Search text box, a complete list of fields that possibly match the field name you are typing or a list of operators that are available on Logger is displayed. This list enables you to select a field (or operator) from the list instead of typing it in, thus enabling you to quickly build a query expression. The field list only contains fields in the Logger schema,

metadata terms (`_storageGroup`, `_deviceGroup`, `_Logger`), and the regular expression term (`|REGEX=`). (See [“Indexing” on page 119](#) for a complete list of fields.)



Search Helper is available by default and automatically displays relevant information based on the query currently entered in the Search text box. If you do not want the Search Helper to display this information automatically, click the “Auto-open is ON” link (in the Search Helper window). The link toggles to “Auto-open is OFF”. To access Search Helper on demand (once it has been turned off), click the down-arrow button to the right of the Search text box.

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.




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 333](#).



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.

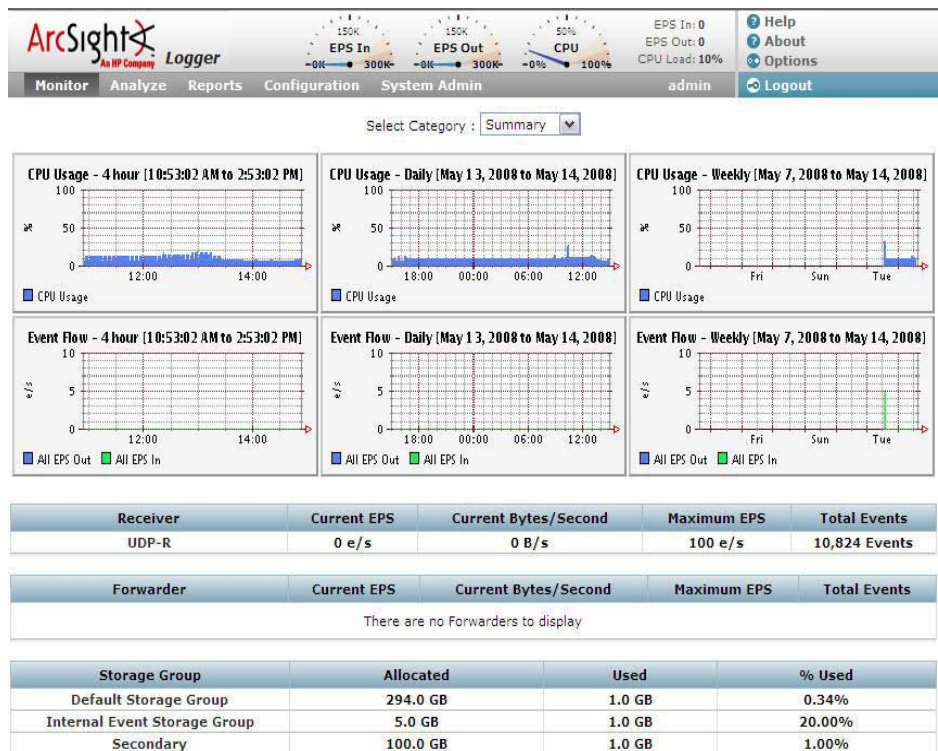


Figure 3-5 The Monitor tab displays summary information by default

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.



- The total space allocated for a storage group includes a certain amount that has been set aside to ensure that the group can receive new events when it is almost full. As a result, the percentage of used space for a storage group never reaches 100% (as displayed on the Monitor > Summary page). For software Loggers installed using the Minimal setting, the maximum % Used (On the Monitor > Summary page) for each storage group reaches up to 66.33%. (Two storage groups of 3 GB each; 1 GB is set aside for new events in each group. After 2 GB of space has been used and the new events are being written to the last 1 GB, Logger automatically triggers retention and reclaims 1 GB of the used space. Thus, the % Used field for each storage group only reaches up to 66.33%.)
 - The “Session Inactivity Timeout” setting on the Authentication Settings page (System Admin > Users/Groups > Authentication) does not apply to the user interface pages accessed through the Monitor menu. That is, if a user is on any of the user interface pages accessed through the Monitor menu and the session has been inactive for the number of minutes specified in the “Session Inactivity Timeout” setting, the user’s session will not time out.
-

Platform

The Platform monitor page, as shown in [Figure 3-6](#), 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 239](#).

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 246](#).

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 233](#).

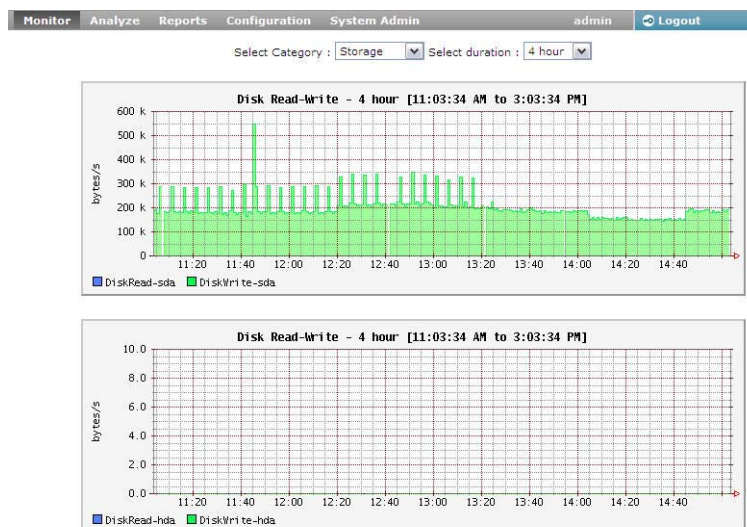


Figure 3-8 Monitor tab, Storage page

Chapter 4

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 71](#)
[“The Process of Searching Events” on page 71](#)
[“Elements of a Search Query” on page 72](#)
[“Syntax Reference for Query Expression” on page 98](#)
[“Using the Search Builder Tool” on page 102](#)
[“Search Analyzer” on page 106](#)
[“Regex Helper Tool” on page 108](#)
[“Searching for Events on Logger” on page 110](#)
[“Understanding the Search Results Display” on page 112](#)
[“Exporting Search Results” on page 116](#)
[“Indexing” on page 119](#)
[“Saving Queries \(Saved Filters and Searches\)” on page 123](#)
[“System Filters/Predefined Filters” on page 124](#)
[“Alerts” on page 130](#)

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 Search Helper facility provides suggestions and possible matches to quickly build a query expression. (See [“Search Helper” on page 64](#) for more information.)

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 72](#).

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 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 73](#).

Search Operators are described in [“Search Operators” on page 75](#).

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 98](#).

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 119](#). For discussion on field-based query performance, see [“Performance Optimizations for Indexed Fields in Search Queries” on page 107](#).)

Field search is case sensitive. Make sure you understand and follow other requirements and guidelines listed in [“Syntax Reference for Query Expression” on page 98](#).

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")
```



If a query includes the boolean operator OR and the metadata identifiers (discussed in ["Constraints" on page 97](#)), 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 119](#) 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
!=	destinationPort != 100
>	bytesIn > 100

Field Operator	Example
>=	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"] _Logger 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. Before you use any search operator, you need to use one of the two special search operators—**cef** and **rex**—that 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 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">1. <code>categorySignificance STARTSWITH "/Normal" AND agentType STARTSWITH "testalert" cef categorySignificance agentType</code>2. <code>logger cef deviceEventCategory name</code>

Search Operator	Description and Examples
chart	<p>Displays search results in a chart form of the specified fields. When multiple fields are specified, the chart function operates on the unique sets of all those fields, as illustrated in “Example 1 for Chart Operator, when multiple fields are specified:” on page 80.</p> <p>By default, a column chart is displayed. You can select from other types of charts, such as bar chart, line chart, pie chart, area chart, stacked column, or stacked bar.</p> <p>If an aggregation function such as count, sum, or avg is specified, a chart of the aggregated results is displayed along with the tabular results of the aggregation operation in a Results Table. For example, for the aggregation function <code>sum(deviceCustomNumber1)</code>, the <code>sum_deviceCustomNumber1</code> column in the Results Table displays the sum of unique values of <code>deviceCustomNumber1</code> field. If this field had two unique values 1 and 20, occurring 2 times each, the <code>sum_deviceCustomNumber1</code> column displays sum of those two values, as shown in the following figure:</p>

deviceCustomNumber1	sum_deviceCustomNumber1
1	2
20	40

Usage: | `chart <field>`

Old Usage: | `chart _count by <field1> <field2> <field3> ...`

New Usage: | `chart count by <field1> <field2> <field3> ...`
`[span [<time_field>]=<time_bucket>]`

Usage: | `chart {{sum | avg | min | max | stdev} (<field>)}+ by <field1>, <field2>, <field3> ...[span [<time_field>]=<time_bucket>]`

Usage: | `chart <function> as <new_column_name> by <field>`
`[span [<time_field>]=<time_bucket>]`

where

`<field>` is the name of the field that you want to chart.

`<time>` is the bucket size for grouping events. Use d for day, h for hour, m for minute, s for seconds. For example, 2h, 5d, 1m. (See Notes for details.)


`<function>` is one of these: `_count`, `count`, `sum`, `avg` (or `mean`), `min`, `max`, `stdev`

`<new_column_name>` is the name you want to assign to the column in which the count is displayed. For example, `Total`.

Notes:

- A `cef` or `rex` operator (to extract fields from matching events) must precede this operator.
- Only one field can be specified in the `chart` command without the `_count by` (or `count by`) option, and for mathematical operators such as `sum`, `avg` (or `mean`), `min`, `max`, and `stdev`. The specified field must contain numeric values.
- The mathematical operators `avg` and `mean` are identical.
- When you use the `chart <field>` operator, the query simply lists and charts each occurrence of the values of the specified field. For example, `chart sourcePort`. The field values are not aggregated. However, when you use an aggregation operator such as `_count by` (or `count by`), `sum`, `avg` (or `mean`), and so on, an aggregation of the specified fields is performed and charted.

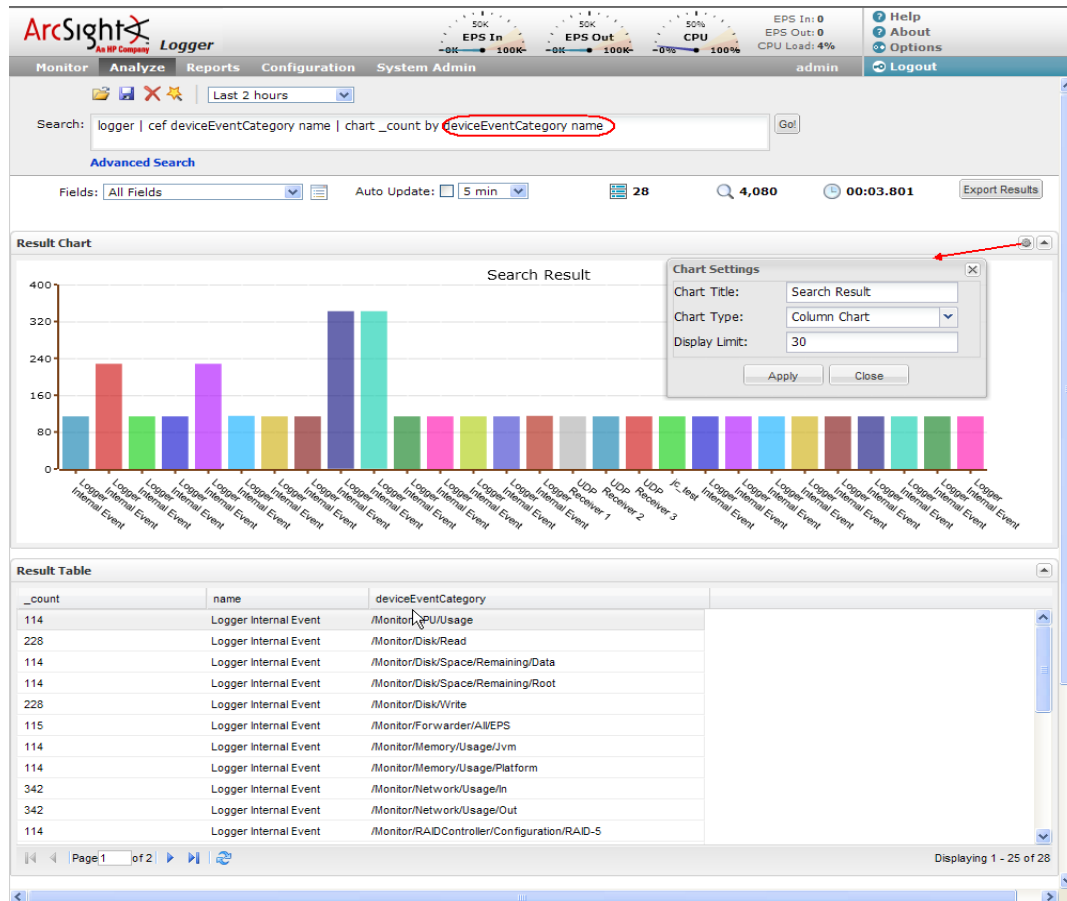
Search Operator	Description and Examples
chart	<ul style="list-style-type: none"> In addition to grouping events by the fields defined by the cef or rex operator, the span operator provides an additional way to group events by a time field (such as <code>EventTime</code> or <code>deviceReceiptTime</code>) and a time bucket. In the following example, <code>deviceReceiptTime</code> is the time field and <code>5m</code> (5 minutes) is the time bucket: <pre> chart _count by deviceEventCategory span (deviceReceiptTime) = 5m</pre> <p>If a time field is not specified for the span operator, <code>EventTime</code> is used as the default. For example, the following query uses <code>EventTime</code> by default:</p> <pre> chart _count by deviceEventCategory span = 5m</pre> <p>Grouping with span is useful in situations when you want to find out the number of occurrences in a specific time span. For example, if you want to find out the total number of incoming bytes every 5 minutes on a device, you can specify a span of 5m, as shown in this example:</p> <pre> chart sum(deviceCustomNumber3) span=5m</pre> <p>The above example assumes that <code>deviceCustomNumber3</code> field provides the incoming bytes information for these events.</p> <p>The span field can be used for grouping in conjunction with or without the cef or rex defined fields. When a span field is specified in conjunction with a cef or rex defined field, the unique sets of all those fields is used for grouping. The following example uses the cef defined fields, <code>deviceCustomNumber3</code> and <code>deviceAddress</code> in conjunction with span to find out the number of events from a specific source in one hour:</p> <pre> chart sum (deviceCustomNumber3) by deviceAddress span=1h</pre> <p>When span is included in a query, search results are grouped by the specified time bucket. For example, if <code>span=5m</code>, the search results will contain one row for each 5-minute span. If there are no events within a specific 5-minute span, that row will be empty.</p> <p>Additionally, the span operator assumes a 24-hour day, all year long. If <code>span=1d</code> or <code>24h</code>, on the day of daylight savings time change, the event time indicated by the <code>span_eventTime</code> field in the search results will be different from the previous day by one hour. On the day when there are 23 hours in a day (in March), the span bucket will still include events from the last 24 hours. Similarly, on the day when there are 25 hours in the day (in November), the span bucket will include events from the last 24 hours. The following example illustrates the <code>span_eventTime</code> field when the span time bucket is 1d and the daylight savings times occurs on March 14th, 2011 and November 7, 2011:</p> <pre>span_eventTime avg_logins 3/11/2011 12am 8 3/12/2011 12am 10 3/13/2011 12am 4 3/14/2011 1am 6 3/15/2011 1am 7 11/5/2011 1am 4 11/6/2011 1am 2 11/7/2011 12am 5 11/8/2011 12am 7</pre>

Search Operator	Description and Examples
chart	<ul style="list-style-type: none"> 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, Stacked column, Stacked Bar. The last two types create stacked charts in which multiple values are plotted in a stack form. These charts are an alternate way of representing multi-series charts, which are described below. - Display Limit: Number of unique values to plot. Default: 10 <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. Only one field can be specified for the aggregation function <code>sum</code>. That is, you cannot use the <code>sum</code> function to add two fields. You can include multiple functions in the same <code>chart</code> command. When doing so, separate each function with a comma, as shown in this example: <pre> chart count, sum(deviceCustomNumber3) by deviceEventClassId</pre> <p>When you include multiple functions, one column per function is displayed in the search Results Table. The Results Chart, however, plots the chart for the field specified in the “by” clause.</p> When you include multiple aggregation functions in a <code>chart</code> command, a multi-series chart is generated that plots the values of the specified aggregation functions along the Y-axis in a single chart. For the following query, unique groups of <code>deviceEventClassId</code> and <code>deviceEventCategory</code> are plotted along the X-axis, and the sum of <code>deviceCustomNumber1</code> and average of <code>deviceCustomNumber2</code> is plotted along the Y-axis, as shown in “Example 2 for a multi-series chart when multiple aggregation functions are specified in a chart command:” on page 81. <pre> chart sum(deviceCustomNumber1), avg(deviceCustomNumber2) by deviceEventClassId, deviceEventCategory</pre> <p>Multi-series charts can be represented in any of the chart settings described above except Pie charts. An example of a multi-series chart, represented as stacked column is shown in “Example 3 for multi-series chart, represented using Stacked Column chart setting:” on page 82.</p> When you export the search results of a chart operator, the newly defined column name (using the <code>chart function as new_column_name</code> command) is preserved. You can use the “as new_column_name” clause to name any column resulting from the aggregation functions, as shown in this example: <pre> chart sum(deviceCustomNumber3) as TotalStorage, avg(deviceCustomNumber3) as AverageStorage by deviceCustomNumber3</pre> Once defined, the newly defined column can be used in the pipeline as any other field. For example, <pre> cef deviceEventClassId deviceCustomNumber3 chart sum(deviceCustomNumber3) as TotalStorage, avg(deviceCustomNumber3) as AverageStorage by deviceCustomNumber3 eval UpdatedStorage = TotalStorage + 100</pre>

Search Operator Description and Examples

Example 1 for Chart Operator, when multiple fields are specified:

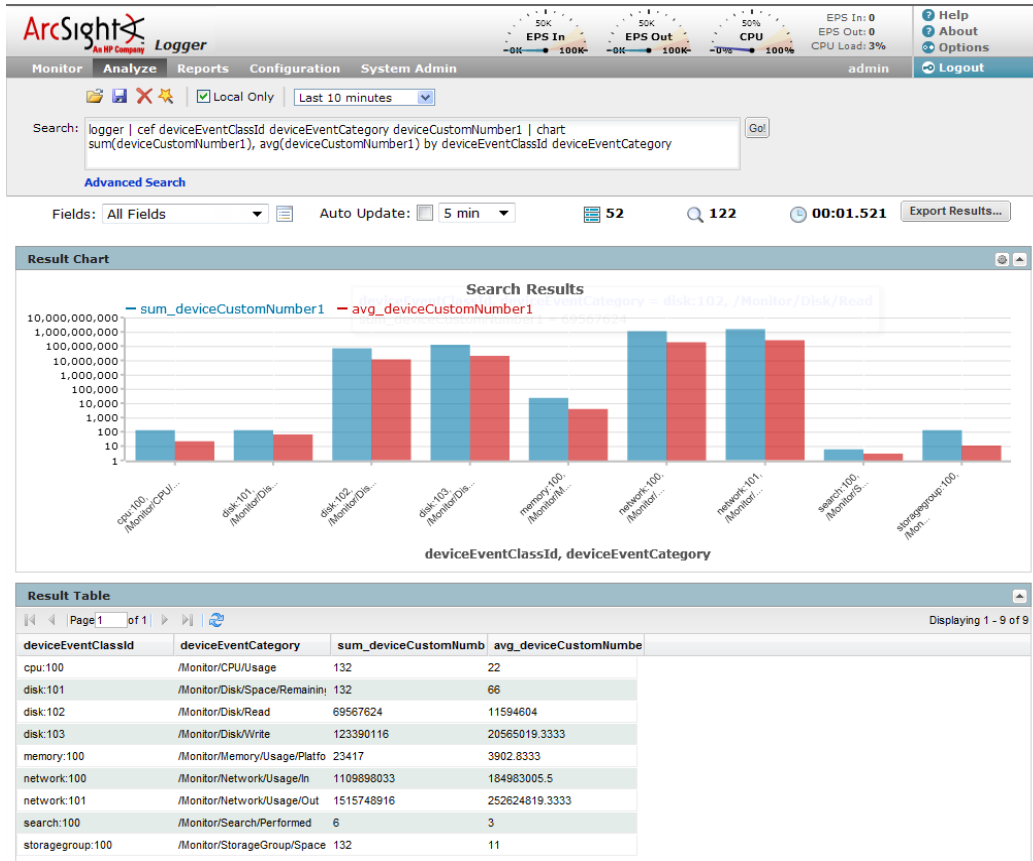
logger | cef deviceEventCategory name | chart _count by deviceEventCategory name



Search Operator Description and Examples

Example 2 for a multi-series chart when multiple aggregation functions are specified in a chart command:

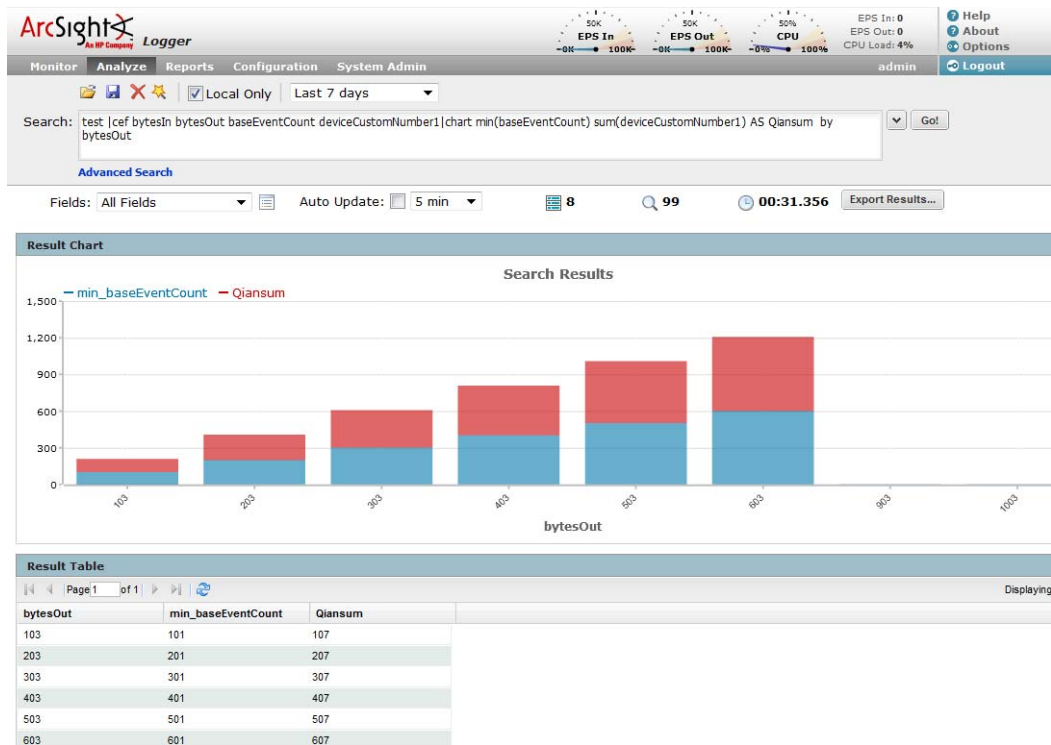
```
logger | cef deviceEventClassId deviceEventCategory deviceCustomNumber1 | chart
sum(deviceCustomNumber1), avg(deviceCustomNumber1) by deviceEventClassId
deviceEventCategory
```



Search Operator Description and Examples

Example 3 for multi-series chart, represented using Stacked Column chart setting:

test | cef bytesIn bytesOut baseEventCount deviceCustomNumber1 | chart min(baseEventCount) sum(deviceCustomNumber1) AS Qiansum by bytesOut



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: eval <expression></p> <p><expression> is a mathematical, string, or boolean operation; for example, total_bytes=bytesIn + bytesOut.</p> <p>Note: Typically, a cef or rex operator (to extract fields from matching events) precedes the eval operator, as shown in the examples below. However, you can use the eval operator on a field that has been defined by a previous eval operator in a query.</p> <p>Examples:</p> <ol style="list-style-type: none"> 1. _storageGroup IN ["Default Storage Group"] cef categoryBehavior eval cat=if(categoryBehavior=="Communicate", "communicate", "notCommunicate") If the Category Behavior is "Communicate", then assign the value "communicate" to a new field "cat"; otherwise, assign the value "notCommunicate" to it. 2. logger cef msg name eval fullname=name + "END" 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. _storageGroup IN ["Default Storage Group"] cef bytesIn bytesOut name eval Plus=bytesIn + 100 sort Plus 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.

Search Operator	Description and Examples
extract	<p>Extract key value pairs from raw events. The key represents a field in the raw event and its value is the characters that appear after the key until the next key in the event. The following raw event is used to illustrate the concept:</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>To extract the URL from the above event, you can define key value pairs that are delimited by greater than sign (">") and square bracket "[" and keys that delimited by equal to ("="). Thus,</p> <p><i>The key value pairs in the event will be:</i> [a href='function.memcache-pconnect']></p> <p><i>The key in the event will be:</i> <a href</p> <p><i>The extracted URL will be:</i> 'function.memcache-pconnect'</p> <p>Usage: <code>extract [pairedelim="<delimiters>"] [kvdelim="<delimiters>"] maxchars=<n> fields="key1,key2,key3..."</code></p> <p><code>pairedelim</code> is a delimiter (or a list of delimiters) that identifies the key value pairs in an event. By default, semi colon, pipe, and comma (; ,) are used. In the above example, ">" and "[" are used as delimiters.</p> <p><code>kvdelim</code> is a delimiter (or a list of delimiters) that separates a key from its value. By default, "=" is used. In the above example, the default delimiter "=" is used.</p> <p><code>maxchars</code> is the maximum number of characters in an event that would be scanned for extracting key value pairs. By default, 10240.</p> <p><code>fields</code> is a key (or a list of comma-separated keys) whose values you want to display in the search results. For example, if you extracted the "<a href" key in the above example and want a column in the search results to display those values, specify "<a href" in the <code>fields</code> list.</p> <p>Notes:</p> <ul style="list-style-type: none"> This operator only works on raw events. That is, you cannot extract key value pairs from CEF events or fields defined by <code>cef</code> and <code>rex</code> operators. You can specify the <code>pairedelim</code> and <code>kvdelim</code> delimiters in the <code>extract</code> operator command to extract keys and their values. However, if you want to determine the key names that these delimiters will generate, use the <code>keys</code> operator as described in "keys" on page 86. The <code>keys</code> operator can only be used to determine keys; you cannot pipe those keys in the <code>extract</code> operator. That is, <code> keys extract fields=field1</code> is incorrect. The keys specified in the <code>fields</code> list can be used further in the pipeline operations. For example, <code> extract pairedelim=" " kvdelim=":" fields="count" top count</code> If none of the specified <code>pairedelim</code> characters exist in an event, the event is not parsed into key value pairs. The whole event is skipped. Similarly, if the specified <code>kvdelim</code> does not exist, values are not separated from the keys. To specify double quotes (") as the delimiter, enter it within the pair of double quotes with backslash(\) as the escape character. For example, <code>"=\" "</code>. Similarly, use two backslashes to treat a backslash character literally. For example, <code>"\\"</code>. <p>Example: To extract only URLs from events in the format shown above, specify this command:</p> <pre> extract pairedelim=">[" kvdelim="=" fields="<a href"</pre>

Search Operator	Description and Examples
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"> categorySignificance STARTSWITH "/Normal" AND agentType STARTSWITH "testalert" <code>cef categorySignificance agentType fields - agentType + categorySignificance</code> logger <code>cef name rex "\sInternal(?<eventName>.*\d{1,3}.\d{1,3}.\d{1,2}.*)" fields - name</code>
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>

Search Operator	Description and Examples
keys	<p>Identifies keys in raw events based on the specified delimiters. Although this operator is not required to determine keys, it is recommended that you use it to first determine the keys whose values you want to obtain using the extract operator. This operator returns aggregated results. Therefore, the search results list the keys found in the matching events and their counts.</p> <p>Usage: <code> keys [pairedlim="<delimiters>"] [kvdelim="<delimiters>"] limit=<n></code></p> <p>pairedlim is a delimiter (or a list of delimiters) that identifies the key value pairs in an event. By default, semi colon, pipe, and comma (; ,) are used.</p> <p>kvdelim is a delimiter (or a list of delimiters) that separates a key from its value. By default, "=".</p> <p>limit is the maximum number of key value pairs to find. There is no default or maximum number for this parameter.</p> <p>Notes:</p> <ul style="list-style-type: none"> This operator only works on raw events. That is, you cannot identify key value pairs from CEF events or fields defined by the cef and rex operators. The keys operator can only be used to determine keys; you cannot pipe those keys in the extract operator. That is, <code> keys extract fields=field1</code> is incorrect. If a key value is blank (or null), it is ignored and not counted toward the number of hits. For example, for the following event data: Date=3/24/2011 Drink=Lemonade Date=3/23/2011 Drink= Date=3/22/2011 Drink=Coffee <i>Search Query:</i> <code>keys pairedlim=" " kvdelim="="</code> <i>Search Result:</i> Date, 3 hits and Drink, 2 hits If none of the specified pairedlim characters exist in an event, the event is not parsed into key value pairs. The whole event is skipped. Similarly, if the specified kvdelim does not exist, values are not separated from the keys. To specify double quotes (") as the delimiter, enter it within the pair of double quotes with backslash(\) as the escape character. For example, <code>"="\ "</code>. Similarly, use two backslashes to treat a backslash character literally. For example, <code>"\\"</code>. <p>Example: <code> keys pairedlim=">[" kvdelim="="</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 cef or rex 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>

Search Operator	Description and Examples
regex	<p>Select events that match the specified regular expression.</p> <p>Usages:</p> <pre> regex <regular_expression></pre> <p>OR</p> <pre> regex field_name (= !=) <regular_expression></pre> <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"> 1. <code>_storageGroup IN ["Default Storage Group"] regex "failure"</code> 2. <code>logger cef deviceEventCategory regex deviceEventCategory != "fan"</code>

Search Operator	Description and Examples
rename	<p>Rename the extracted field.</p> <p>Usage: <code>rename <extracted_field_name> as <new_name></code></p> <p><extracted_field_name> is the name of the field extracted using the cef or rex operator.</p> <p><new_name> is the new name you want to assign to the extracted field.</p> <p>Notes:</p> <ul style="list-style-type: none">• An additional column is added to the search results for each renamed field. The field with the original name continues to be displayed in the search results in addition to the renamed field. For example, if you rename deviceEventCategory to Category, two columns are displayed in the search results: deviceEventCategory and Category.• You can include the wildcard character, *, in a field name. However, you must enclose the field that contains a wildcard character in double quotes (" "). For example: rename "*IPAddress" as "*Address" OR rename "*IPAddress" as Address• If a field name includes a special character (such as _, a space, #, and so on), it should be included in double quotes (" ") in the rename operator expression. For example: rename src_ip as "Source IP Address"• If the resulting field of a rename operation includes a special character, it must be enclosed in double quotes (" ") whenever you use it in the pipeline operator expression. For example, rename src_ip as "Source IP Address" top "Source IP Address"• The internal field names (that start with "_raw") cannot be renamed.• The renamed fields are valid only for the duration of the query.• The resulting field of a rename operation is case sensitive. When using such a field in a search operation, make sure that you the same case that was used to define the field.• When you export the search results of a search query that contains the rename expression, the resulting file contains the renamed fields.

Search Operator	Description and Examples
replace	<p>Replace the specified string in the specified fields with the specified new string.</p> <p>Usage: <code>replace <orig_str> with <new_str> [in <field_list>]</code></p> <p><orig_str> is the original string you want to replace. (See Notes for more details.)</p> <p><new_str> is the new string you want to replace with. (See Notes for more details.)</p> <p><field_list> is the optional, comma-separated list of fields in which the string will be replaced. If a field is not specified, all user-defined fields (fields created as a result of the cef, rex, and eval operators) are scanned.</p> <p>Notes:</p> <ul style="list-style-type: none"> An additional column of the same name is added to the search results for each field in which string is replaced. The column with the original value continues to be displayed in the search results in addition to the column with replaced values. For example, if you replace err with Error in the "message" column, an additional "message" column is added to the search results that contains the modified value. If you want to replace the entire string, specify it in full (as it appears in the event). For example, "192.168.35.3". If you want to replace a part of the string, include wildcard character (*) for the part that is not going to change. Examples: If the original string (the string you want to replace) is "192.168*", only the 192.168 part in an event is replaced. The remaining string is preserved. As a result, if an event contains 192.168.35.3, only the first two bytes are replaced. The rest (35.3) will be preserved. Similarly, if the event contains 192.168.DestIP, DestIP will be preserved. However, if the event contains the string 192.168, it will not be replaced. If both, the original and the new strings contain wildcard characters, the number of wildcard characters in the <i>original</i> string must match the number of wildcard characters in the <i>new</i> string. <code>replace "*.168.*" with "*.XXX.*</code> If the original or the new string includes a special character such as / or ?, enclose the string in double quotes (" "): <code>replace "/Monitor" with Error</code> You can replace multiple values for multiple fields in a single operation by separating each expression with a comma (.). Note that you must specify the field list after specifying the "with" expression for all values you want to replace, as shown in the following example: <code>replace "Arc*" with HP, "cpu: 100" with EPS in deviceVendor, deviceEventClassId</code> The original string is case-insensitive. Therefore, the string "err" will replace an event that contains "Err".

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: rex <regular_expression containing a field name></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.</p> <p>The substitution only occurs in the Search results. The actual event is not changed.</p> <p>Usage: cef <field> rex field = <field> mode=sed "s/<string to be substituted>/<substitution value>/g"</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 an 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 469. 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 108. 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>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./xxx/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 [<n>] field1 field2 field3 ...</code></p> <p><n> limits the matches to the top <i>n</i> values for the specified fields.</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. To limit the matches to the top <i>n</i> values for the specified fields, specify a value for <i>n</i>. For example, <code> top 5 deviceEventCategory</code> <p>Examples:</p> <ol style="list-style-type: none"> <code>arcsight cef deviceEventCategory top deviceEventCategory</code> <code>_storageGroup IN ["Default Storage Group"] cef deviceEventCategory eventId rex "deviceEventCategory=(?<categories>[^\]*)" top 5 categories</code>
where	<p>Display events that match the criteria specified in the "where" expression.</p> <p>Usage: <code> where <expression></code></p> <p><expression> can be any valid field-based query expression, as described in "Indexed and Non-Indexed Fields" on page 73.</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. <expression> can only be a valid field-based query expression. Arithmetic expressions or functions are not supported. When the <code>where</code> 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"> <code>_storageGroup IN ["Default Storage Group"] cef eventId where eventId is NULL</code> <code>_storageGroup IN ["Default Storage Group"] cef eventId deviceVersion where eventId=10006093313 OR deviceVersion CONTAINS "4.0.6.4924.1"</code> <code>_storageGroup IN ["Default Storage Group"] cef deviceEventCategory eventId rex "deviceEventCategory=(?<categories>[^\]*)" 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 93](#). The <units> portion, if given, consists of an integer and a single, case-sensitive letter, as shown in [Table 4-2 on page 94](#).

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)

Period	Description
\$CurrentMonth	Midnight on the first day of the current month
\$CurrentYear	Midnight on the first day of the current year

Table 4-2 Units

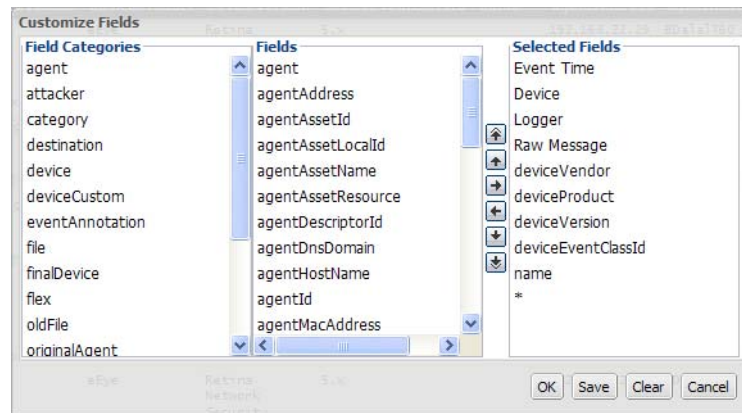
Unit	Description
m (lowercase)	Minutes do not confuse with 'M', meaning months)
h	Hours
d	Days
w	Weeks
M (uppercase)	Months (do not confuse with 'm', meaning minutes)

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	<p>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.</p> <p>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.</p>
All Fields (w/out raw messages)	
Minimal Fields	
Syslog Standard	
Categories	
Base Event Fields	

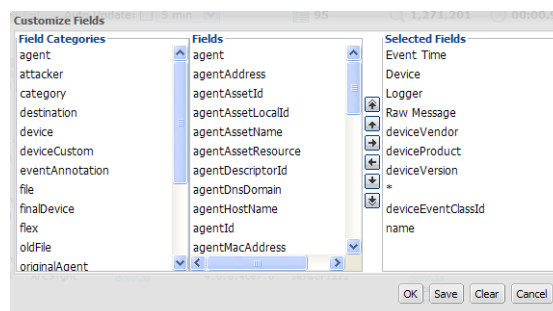
You can also create your own field sets. The Logger user interface offers a simple and intuitive interface to select and move event fields you want to include in a field set, as shown in the following figure.



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.



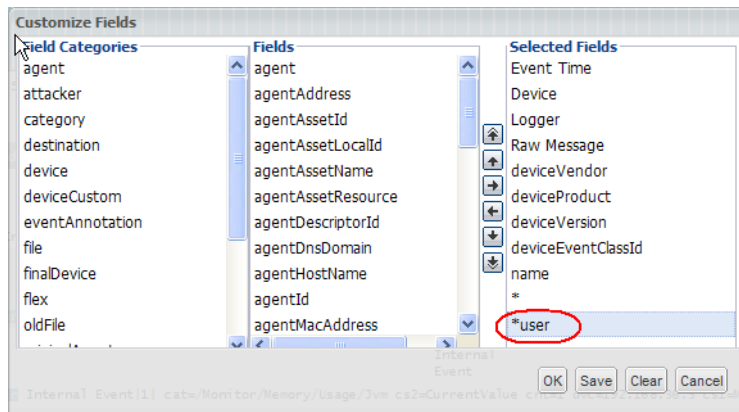
Field set selection is specific to a Logger user's interface. For example, UserA and UserB are connected to the same Logger and are using the default, All Fields, field set for search results display. UserA changes his selection to a custom field set. This change will only impact UserA's display; UserB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



Field sets are not included in the saved filter definition.

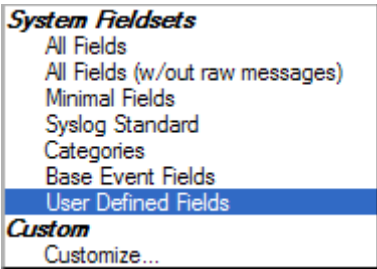
For information about deleting custom field sets, see [“Deleting Custom Field Sets” on page 280](#).

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 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>_Logger</code>	<code>_Logger 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 97), the expression to be evaluated with OR must be enclosed in parentheses, as shown in this example:

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

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

- When specifying multiple groups in a constraint, ensure that the group names are enclosed in 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 boxUse Logger’s Search Helper to enter a constraint in the Search text box. Once you type “_s” (for storage group), “_d” (for device group), or “_p” (for Logger) in the

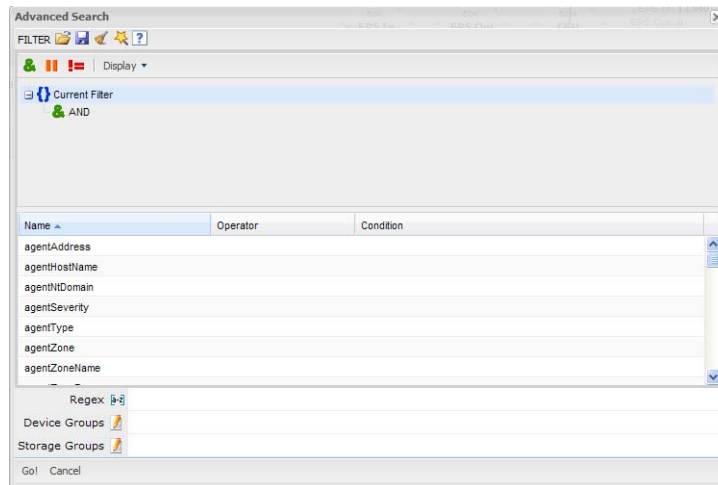
Search text box, Logger automatically provides a drop-down list of relevant terms and operators from which you can select.



Caution

If a search query contains constraints and a regular expression, make sure that the constraints are specified before the regular expression. For example, `_Logger 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 102](#).



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	<code>keyword1</code> <code>boolean_operator</code> <code>keyword2</code> <code>boolean_operator</code> <code>keyword3...</code>	<code>field_name operator</code> <code>field_value</code> (List of fields in the “Indexing” on page 119 section.) (List of operators in the “Indexed and Non-Indexed Fields” on page 73 section.)	<code> REGEX=" <REGEX1> "</code> <code> </code> <code>REGEX=" <REGEX2> "</code> <code> </code> <code>..</code>

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 97), 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 73 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 97), 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 92.</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 _Logger), 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 73 section to connect and nest field search expressions. Metadata identifiers (_storageGroup, _deviceGroup, and _Logger), 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</p> <p>(Cannot be changed.)</p>	<p>Sensitive*</p> <p>(Can be changed using Tuning options. See "Tuning Advanced Search Options" on page 279.)</p>	<p>Insensitive*</p> <p>(Can be changed using Tuning options. See "Tuning Advanced Search Options" on page 279.)</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> <p>name=*log (searches for ablog, blog, and so on.)</p> <p>name="*log"</p> <p>name=*log</p> <p>(both search for *log)</p>	<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> <p>log (matches log, logging, logger, and so on)</p> <p>"log" (matches only log)</p> <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> <p>message="failed login"</p>	<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
= . : / \ @ - ? # \$ & _ %	You can search for these characters in a keyword. However, enclose the keyword in double quotes. Example: "John="	You can search for these characters. Enclose value in double quotes if value contains any of these characters. Example: name="John="	<ul style="list-style-type: none"> Cannot contain ^ in the beginning and \$ at the end as a matching character unless the regular expression you specify must look for an event that contains only the pattern you are specifying; for example, <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	No specific format. The query needs to contain the exact timestamp string. For example, "10:34:35". Note: The string cannot contain spaces. For example, "Oct 19" is invalid.	Use this format to specify a timestamp in a query (including double quotes): <code>"mm/dd/yyyy hh:mm:ss"</code> OR <code>"yyyy/mm/dd hh:mm:ss timezone"</code> OR <code>"MMM dd yyyy hh:mm:ss timezone"</code> 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.	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,

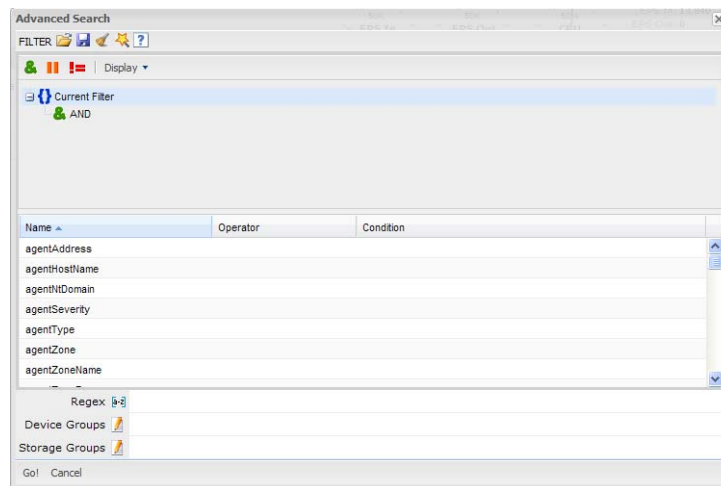
and regular expressions using this tool. In addition, the tool enables you to specify search constraints such as device groups and storage groups (see [“Constraints” on page 97](#)). This section describes how to use the tool.

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

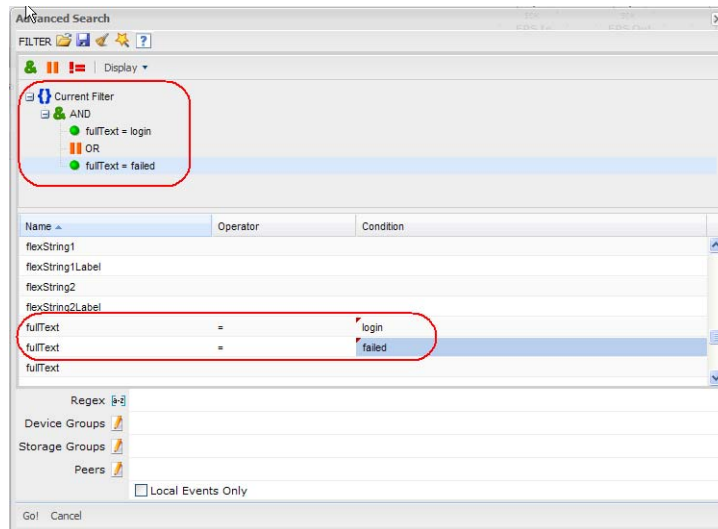
- 2 If you want to load a system or saved filter, or a saved search, click the icon. Select the filter or the saved search from the displayed list and click **Load+Close**.

For more information, see [“Saving Queries \(Saved Filters and Searches\)” on page 123](#) and [“System Filters/Predefined Filters” on page 124](#).

- 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 Loggers, click the  icon next to the constraint category. Select the relevant groups and Loggers. (To select multiple groups, hold the Ctrl-key down.)

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


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

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

- 7 Click **Go**.

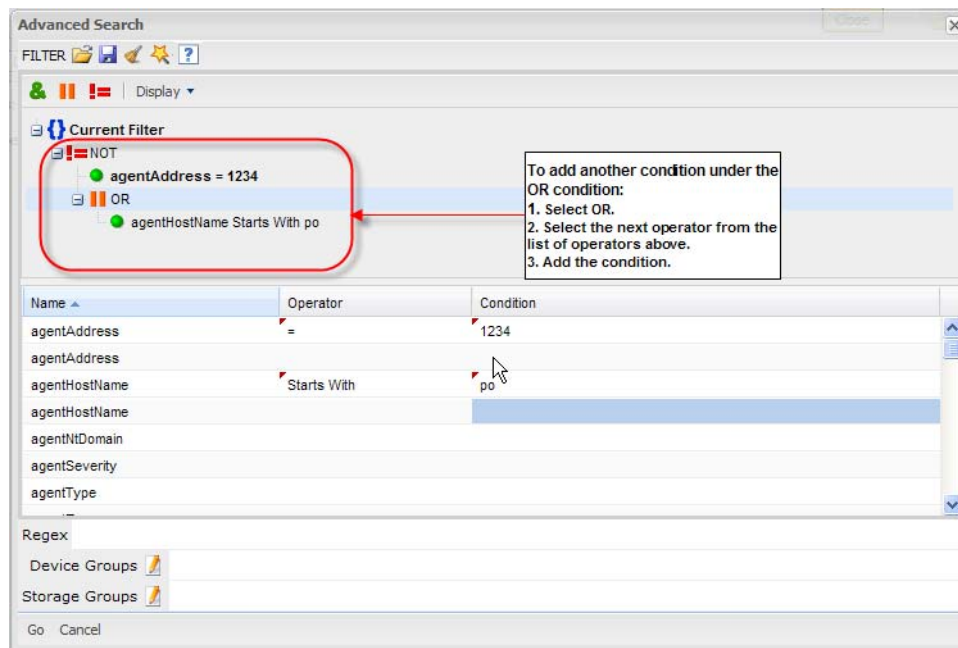
The query is automatically displayed in the Search text box and is ready to be run.

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 123](#).

Nested Conditions

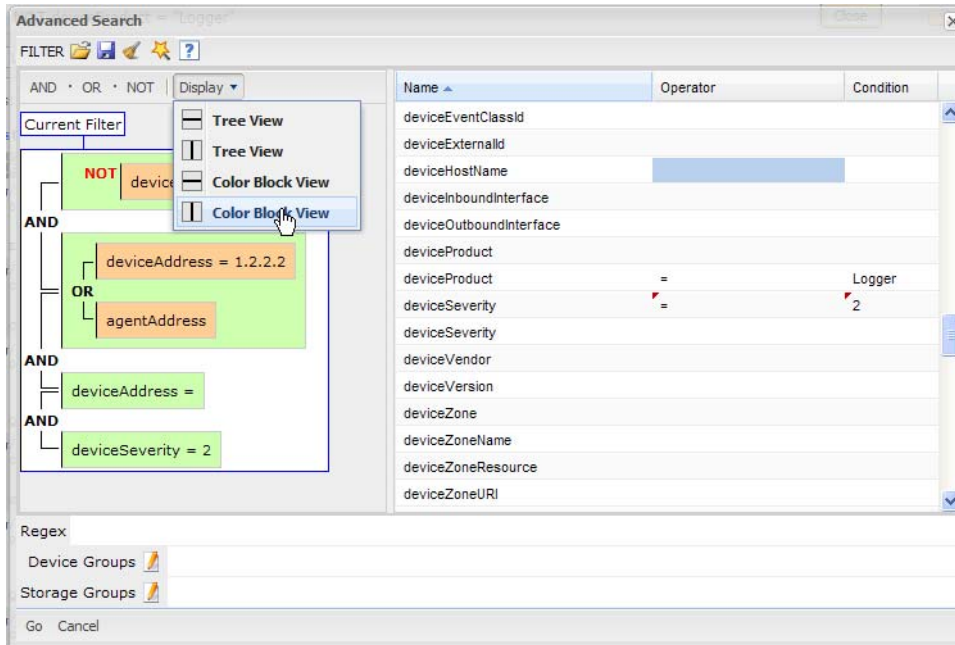
You can create search queries with nested conditions in Search Builder. To do so, click the operator under which you want to nest the next condition and add the condition as described in [“Accessing Search Builder” on page 103](#).



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.

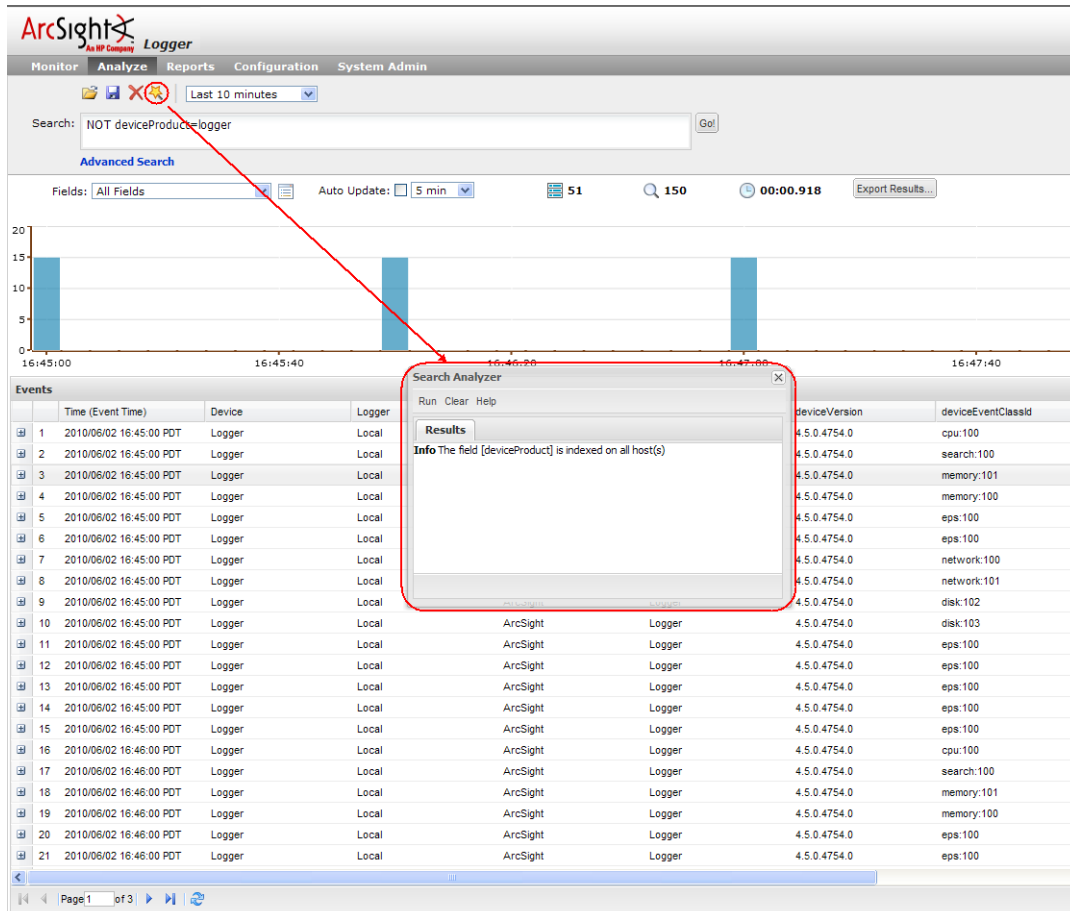
Search Analyzer

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 75](#) or [Appendix C, Using the Rex Operator, on page 469](#).) This tool not only simplifies the task of creating regular expressions for the `rex` operator but also makes it efficient and error free.

The tool, which is only available for non-CEF events, 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 110](#).)
- 2 Identify a syslog event that you want to analyze further. For example, in the shown figure, event #1 is the event we will analyze further.
- 3 Click the  icon (in the left-most column) for the identified event to expand it and display its raw event.
- 4 Click the  icon (next to the word **RAW**) to launch the Regex Helper tool.
- 5 Select the fields that you want to extract.
- 6 Click **OK**.

The screenshot shows the ArcSight Logger interface. At the top, there's a navigation bar with 'Monitor', 'Analyze', 'Reports', 'Configuration', and 'System Admin'. Below this is a search bar with 'tcp' entered. A bar chart shows event frequency over time. A table of events is displayed, with one event selected. The 'Raw Event' pane on the right shows the event details and a regex expression. The 'Extract Fields' pane shows the fields extracted from the event, including 'IPAddress_2' which is highlighted. The 'OK' button is also highlighted.

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.

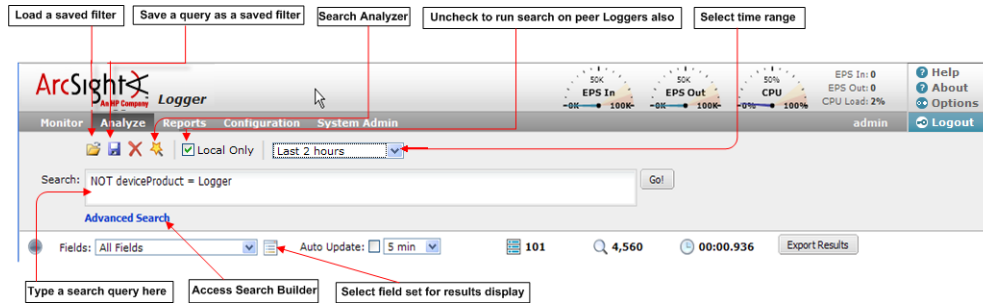
Search: `tcp | rex "[\S+ \S+ +\d+ \d+:\d+:\d+ \d+] [\S+?] [\S+?] (?<IPAddress_1>\d+.\d+.\d+.\d+) [\S+?] \S+?: \S+?(\S+?) [\S+?] \S+?=\S+?> \S+?<\S+?> \S+? (?<IPAddress_2>\d+.\d+.\d+.\d+).*" | top` **Go!**

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

Search: `tcp | rex "[\S+ \S+ +\d+ \d+:\d+:\d+ \d+] [\S+?] [\S+?] (?<IPAddress_1>\d+.\d+.\d+.\d+) [\S+?] \S+?: \S+?(\S+?) [\S+?] \S+?=\S+?> \S+?<\S+?> \S+? (?<IPAddress_2>\d+.\d+.\d+.\d+).*" | top` **Go!**

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 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 72](#) 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 72](#).


When you type a query, Logger’s Search Helper enables you to quickly build a query expression by automatically providing suggestions, possible matches, and applicable operators. See [“Search Helper” on page 64](#) for more information.

Use these guidelines to include various elements in a search query:

- For a complete list of fields in Logger schema, see [“Indexing” on page 119](#).
- Metadata terms (`_storageGroup`, `_deviceGroup`, `_Logger`)
Type “_s” (for storage group), “_d” (for device group), or “_p” (for Logger) in the Search text box to obtain a drop-down list of constraint terms and operators.

- Regular expression term (`|REGEX=`)


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 102](#) for more information.) Also, use this option to specify device groups, storage groups, and 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 123](#) and [“System Filters/Predefined Filters” on page 124](#).

- 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 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 92](#).
 - 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 94](#).
- 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 112](#).

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 123](#).

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 279](#).

Searching Peer Loggers (Distributed Search)

When you run a search query, by default, only your local Logger is searched for matching events. However, when specifying a query, you can select an option to run the search on the peer Loggers. You can also select the Loggers to which the search should be constrained, as described in [“Searching for Events on Logger” on page 110](#).

Follow these guidelines for searching across peers:

- Loggers can run different versions. However, these are the only supported paths for running a search across peers:
 - ◆ v5.x software Logger to v5.x software
 - ◆ v5.x software Logger to v5.x appliance
 - ◆ v4.5 software Logger to v5.x appliance
 - ◆ v4.5 appliance to v5.x appliance
 - ◆ v5.x appliance to v5.x appliance
- If you need to run a search using the pipeline operators across peers, make sure the peers are running Logger v5.1 or later. For information about pipeline operators, see [“Search Operators” on page 75](#).
- If Loggers do not have identical storage or device group names, a search query operation skips searching for events for those groups on those peers.
- A user needs to belong to these user groups with the listed permissions set to Yes to perform 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” user rights set to Yes.
- When a Logger becomes unavailable during a search operation, the one of the following errors might be displayed:

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

These error messages can occur when the Logger cannot be reached. Restore the 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 relationship is restored. However, ignore those messages as these go away when you run a new distributed search.

Tuning Search Performance

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


- Enable full-text and field-based indexing. When events are indexed, Logger can quickly and efficiently search for relevant data.
- The amount of time it takes to search depends on the size of the data set that needs to be searched through, the complexity of the query, and whether the search is distributed across peers. To limit the data set, ensure that time range within which the events must be searched does not result in a query that needs to scan multi-millions of events. Additionally, limiting search to specific storage groups typically results in better search performance than when the storage groups are not specified.
- Reduce other load on the system when your query needs to run, such as scheduled jobs, large number of incoming events, multiple reports being run.



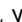

Understanding the Search Results Display

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.)

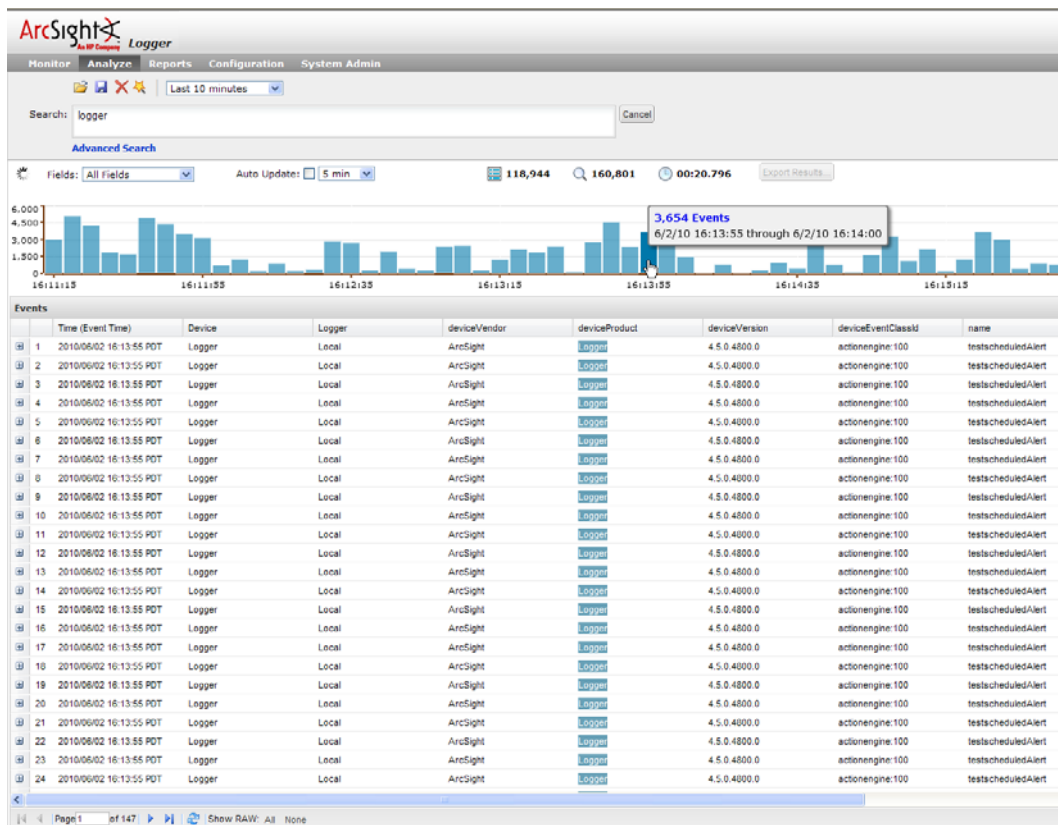
A search operation can take time when millions of events need to be searched. When the first screen of events that match the specified conditions is available, Logger automatically pauses the search and displays the matched events. By default, 25 events are displayed on one screen. Event data is categorized 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.

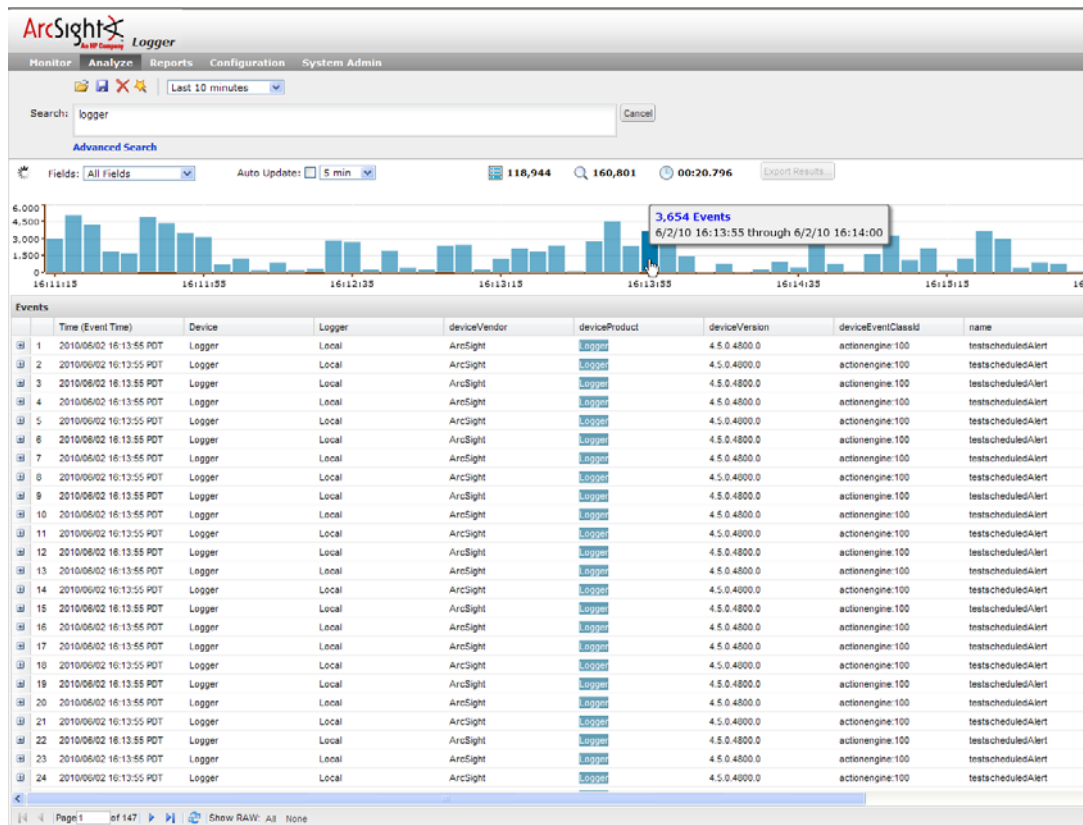
As you roll the mouse over other terms in the events table, they highlight in green. Click a green-highlighted term to add it to the current query. For example, if you search for "login" and roll over the word "fail" in the search results, "fail" will highlight in green. Click the

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.
- Histogram is based on the Logger receipt time of the events (similar to search queries that also use the Logger receipt time to search for events).
- The time distribution on the X-axis is determined automatically.
- You can mouse-over any histogram bar to view the number of matching events and the date and time period that the bar represents.
- You can drill-down to events in a specific time period by clicking the bar on the histogram that represents that time period. The selected section is highlighted and the events matching that time period are listed below the histogram. The histogram continues to display the distribution of all of the matching events, as shown in the following figure. For example, if you select a bar that represents 11,004 events on 2/22/2010 from 12:25:49 a.m. to 12:26:49 a.m. in the following histogram, the details of those events are listed below the histogram; however, the histogram displays all time units and the associated bars. You can also select multiple consecutive bars on the histogram to view matching events in all of the selected time units. To deselect a selected bar, click it.



- A histogram is progressively built and displayed as events match a search query. If the search query needs to scan a large amount of data or a large time period, the

histogram displayed initially might refresh multiple times while the query is running. To view the complete (and final) histogram of a search query, wait until the query has finished running (that is, the screen does not display the circular “waiting” icon anymore).

- The time range on the X-axis might not match the time range specified in the search query because the start and end times on the X-axis are determined by the event times of the first and last matching events of the search query.
- The first one million matching events are plotted on the histogram. If a search query matches more than one million events, an informational message is displayed on the screen.

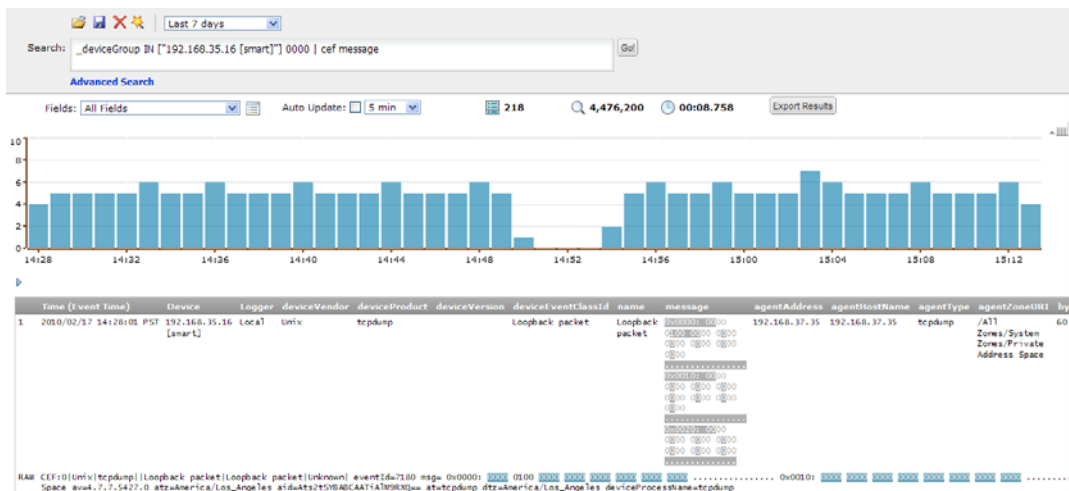
If you need to use the histogram view for event analysis for a search query that matches more than one million events, ArcSight suggests that you adjust the time range specified in your search query such that less than one million are matched to obtain a complete and meaningful histogram or use a pipeline operator such as `top`, `head`, or `chart` to further refine search results such that the total number of hits is under one million events.

Multi-line Data Display

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

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



- 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.
- 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.


Starting with Logger v5.1, the data for the following time fields is exported in human-readable format: deviceReceiptTime, startTime, endTime, agentReceiptTime. For example, 2011/03/21 20:22:09 PDT.



Note

When you export search results on a Logger 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).

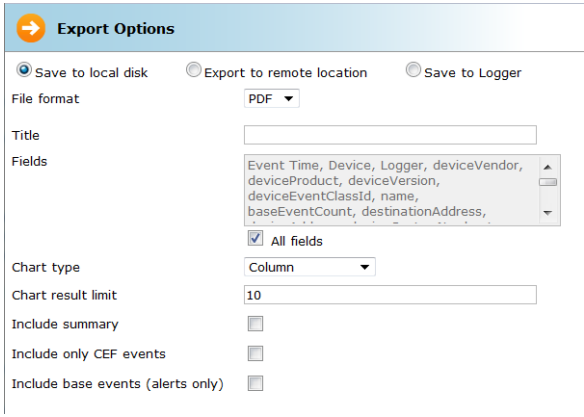
The following is an example of a quick report generated in PDF format. The chart is displayed first, followed by a table of matched events (not shown in this example). All generated charts (including stacked charts) can be exported.



To export search results:

- 1 Run a search query.
- 2 Click **Export Results** in the top right-hand side of the search results screen.
- 3 Select from the following export options.

Option	Description
Save to local disk	The file is saved to a local system from which you are accessing Logger or is it sent to the browser for viewing or saving.
Export to remote location	<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 NFS or CIFS.</p>
Save to Logger	The file is written to the Logger's local storage.

Option	Description
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 chart, top, and so on.</p> 
Export file name	<p>(Available only when the “Export to remote location” option is selected)</p> <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>
Title	<p>(Optional) A meaningful name that appears on top of the PDF file. If no title is specified, “Untitled” is included.</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>
Chart Type (for PDF only)	<p>(Available only when a chart is available in search results)</p> <p>Type of chart to include in the PDF file. You can select from: Column, Bar, Pie, Area, Line, Stacked Column, Stacked Bar.</p> <p>Note: If the Chart Type is different from the chart displayed on the Search Results screen, the value selected for this option overrides the one shown in the screen. Therefore, the exported PDF contains the chart you specify for this option and not the one shown on the screen.</p>
Chart Result Limit (for PDF only)	<p>(Available only when a chart is available in search results)</p> <p>Number of unique values to plot. Default: 10</p> <p>If the configured Chart Result Limit is less than the number of unique values for a query, the top values equal to the Chart Result Limit are plotted. That is, if the Chart Result Limit is 5 and 7 unique values are found, the top 5 values will be plotted.</p>
Include Summary	<p>Include an event count in the exported search results.</p>
Include Only CEF Events	<p>Only include CEF events in the exported search results.</p>

Option	Description
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 275](#).

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 119](#).
- Field-based indexing—Event fields are indexed based on a predetermined schema. See [“Field-based Indexing” on page 120](#).

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 [“Initializing the Logger Appliance” on page 27](#)). 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 122](#).

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 [“Initializing the Logger Appliance” on page 27](#)). 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.



Note

- 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
agentType	deviceCustomNumber2	flexNumber1Label
agentZone	deviceCustomNumber2Label	flexNumber2
agentZoneName	deviceCustomNumber3	flexNumber2Label
agentZoneResource	deviceCustomNumber3Label	flexString1
agentZoneURI	deviceCustomString1	flexString1Label

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

Guidelines for Field-based Indexing

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

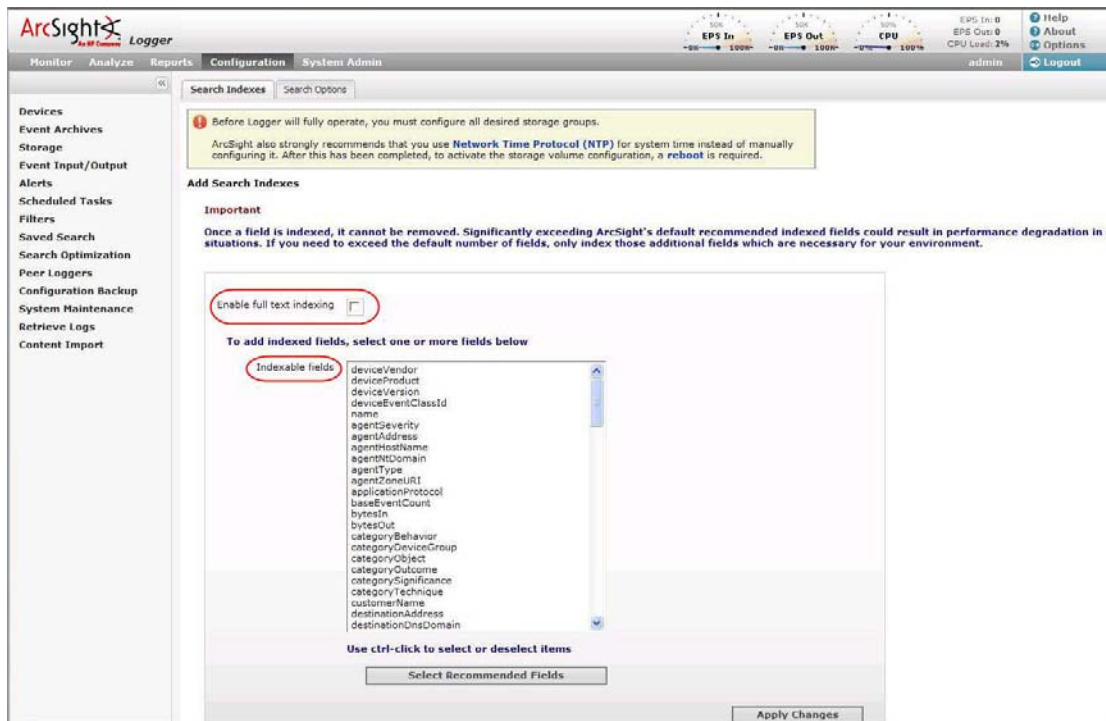
- Events are indexed by any fields you add to the index and the default event metadata fields—event time, Logger receipt time, and device address.

- 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 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 [“Initializing the Logger Appliance” on page 27](#)), 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 260](#).

Saving a Query

To save a query:

- 1 Define a query as described in [“Searching for Events on Logger” on page 110](#) or [“Using the Search Builder Tool” on page 102](#).
- 2 Click the Save icon (📌) and enter a name for the query in the Name field, as shown in the following figure.

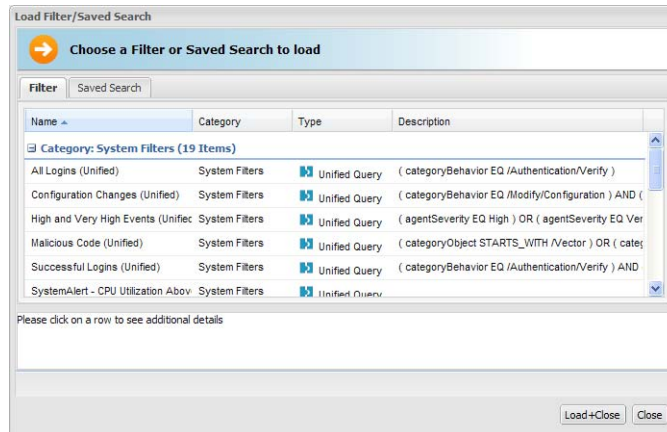
- 3 Select whether you want to save this query as a filter or as a saved search.
If you select to save as a Saved Search, you can either keep the saved query as Saved Search or change it to a Scheduled Alert by specifying a schedule based on which the query runs periodically and generates alerts.
If you choose to schedule the Saved Search, you can either specify the schedule in the following screens or skip it for now.

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



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

The following is a list of all the system filters.

Category	Unified Query Filters	Regular Expression Query Filters
Login Status use case	All Logins	All Logins (Non-CEF) All Logins (CEF format)
	Unsuccessful Logins	Unsuccessful Logins (Non-CEF) Unsuccessful Logins (CEF format)
	Successful Logins	Successful Logins (Non-CEF) Successful Logins (CEF format)
	Failed Logins	
Configuration	Configuration Changes	System configuration changes (CEF format)
Events use case	High and Very High Severity Events	High and Very High Severity CEF events
	Event Counts by Source	
	Event Counts by Destination	
		All CEF events
Intrusion use case	Malicious Code	Malicious Code (CEF format)
Firewall use case	Deny (Firewall Deny)	
	Drop (Firewall Drop)	
	Permit (Firewall Permit)	
Network use case	DHCP Lease Events	
	Port Links Up and Down	
	Protocol Links Up and Down	
Connector System Status use case	CPU Utilization by Connector Host	
	Disk Utilization by Connector Host	
	Memory Utilization by Connector Host	
UNIX Server use case	CRON related events	
	IO Errors and Warnings	
	PAM and Sudo Messages	
	Password Changes	
	SAMBA Events	
	SSH Authentications	
	User and Group Additions	
	User and Group Deletions	

Category	Unified Query Filters	Regular Expression Query Filters
Windows Events use case	Account Added to Global Group	
	Audit Policy Change	
	Change Password Attempt	
	Global Group Created	
	Logon Bad User Name or Password	
	Logon Local User	
	Logon Remote User	
	Logon Unexpected Failure	
	New Process Creation	
	Pre-Authentication Failure	
	Special Privileges Assigned to New Logon	
	User Account Changed	
	User Account Password Set	
	Windows Events (CEF)	

Category	Unified Query Filters	Regular Expression Query Filters
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 128.</p> <p>NOTE: Although these filters are displayed on the software version of Logger, these do not apply to it.</p>	
	CPU Utilization Above 90 Percent	CPU Utilization Above 90 Percent
	CPU Utilization Above 95 Percent	CPU Utilization Above 95 Percent
	Disk Failure	Disk Failure
	Root Partition Below 10 Percent	Root Partition Below 10 Percent
	Root Partition Below 5 Percent	Root Partition Below 5 Percent
	Device Configuration Changes	Device Configuration Changes
	Filter Configuration Changes	Filter Configuration Changes
	High CPU Temperature	High CPU Temperature
		Bad Fan
	Power Supply Failure	Power Supply Failure
	RAID Status Battery Failure	RAID Status Battery Failure
	RAID Status Disk Failure	RAID Status Disk Failure
	Storage Configuration Changes	Storage Configuration Changes
	Storage Group Usage Above 90%	Storage Group Usage Above 90%
	Storage Group Usage Above 95%	Storage Group Usage Above 95%
	Zero Events Incoming	Zero Events Incoming
	Zero Events Outgoing	Zero Events Outgoing

Using a System Filter

To use a predefined system filter, follow instructions in [“Using a Saved Filter or a Saved Search” on page 124](#).

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 124](#). The predefined system health filters are based on the system health events listed in [“System Health Events” on page 128](#).
- Searching for system health events in Logger's Internal Storage Group, as listed in [“System Health Events” on page 128](#). 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 [“Static Routes” on page 308](#)) or create an SNMP Destination (see [“SNMP Destinations” on page 266](#)) or Syslog Destination (see [“Syslog Destinations” on page 267](#)).

- 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 255](#)).
- 3 Enable the new Alert.

System Health Events

The following table lists the system health events that Logger generates. These events are also referred as Logger Internal Events because they are stored in Logger's Internal Storage Group. The predefined System Filters that provide system health status are based on some of these events. If a predefined filter does not suit your needs, create an alert using one of these events.

Starting with Logger v5.1, the format in which system health events are generated has changed to provide more meaningful information. Some of the changes are:

- Addition of new events. For example, Current and Voltage events.
- Instead of referring to all system health events as Logger Internal Event in the `name` field, meaningful names are used. For example, Fan OK, Temperature OK.
- Three severity levels for each event have been added to the `agentSeverity` field—1 (OK), 5 (Degraded), and 8 (Severe).
- The `deviceCustomString` and `deviceCustomStringLabel` field mappings have changed. Refer to a specific event to identify changes.
- Device Event Class ID (`deviceEventClassId`) and Device Event Category (`deviceEventCategory`) of the events have changed. The updated list is available in the following table.
- All hardware-related events are classified as `hardware:nnn` events, where `nnn` is a three-digit number that identifies the hardware component. For example, `hardware:13x` identifies the fan events.

When you upgrade to Logger v5.1, any existing filters or queries based on earlier events will not work on the events collected after the upgrade. However, those filters and queries will continue to work on the events collected prior to the upgrade. Note that the predefined System Filters are compatible with the new and the old format.



Note

The sensor names in each event are hardware specific; therefore, they are not consistent across various Logger platforms. Use the event name (Name) and status (CustomString3) fields to determine the status of a sensor. The raw status (CustomString4), location (CustomString5), and sensor name (CustomString6) fields are for informational use when diagnosing a hardware problem and are not consistent across appliance types.

Group	Device Event Category	Device Event Class ID
System Health Events for appliance and software Loggers		
CPU	/Monitor/CPU/Usage	cpu: 100
Disk	/Monitor/Disk/Read	disk: 102
	/Monitor/Disk/Write	disk: 103

Group	Device Event Category	Device Event Class ID
EPS	/Monitor/Receiver/EPS/All	eps: 100
	/Monitor/Receiver/EPS/Individual	eps: 102
	/Monitor/Forwarder/EPS/All	eps: 101
	/Monitor/Forwarder/EPS/Individual	eps: 103
Memory	/Monitor/Memory/Usage/Platform	memory: 100
Network	/Monitor/Network/Usage/In	network: 100
	/Monitor/Network/Usage/Out	network: 101
Search	/Monitor/Search/Performed	search: 100
Storage Group	/Monitor/StorageGroup/Space/Used	storagegroup: 100
	Note: The size of the storage group, indicated by the “fsize” field is in GB.	
System Health Events for appliance Loggers only		
Battery	/Monitor/Sensor/Battery/OK	hardware: 121
	/Monitor/Sensor/Battery/Degraded	hardware: 122
	/Monitor/Sensor/Battery/Failed	hardware: 123
Current (Electrical)	/Monitor/Sensor/Current/OK	hardware: 101
	/Monitor/Sensor/Current/Degraded	hardware: 102
	/Monitor/Sensor/Current/Failed	hardware: 103
Disk	/Monitor/Disk/Space/Remaining/Data	disk: 101
Fan	/Monitor/Sensor/Fan/OK	hardware: 131
	/Monitor/Sensor/Fan/Degraded	hardware: 132
	/Monitor/Sensor/Fan/Failed	hardware: 133
Power Supply	/Monitor/Sensor/PowerSupply/OK	hardware: 141
	/Monitor/Sensor/PowerSupply/Degraded	hardware: 142
	/Monitor/Sensor/PowerSupply/Failed	hardware: 143
RAID	/Monitor/RAID/Controller/OK	raid: 101
	/Monitor/RAID/Controller/Degraded	raid: 102
	/Monitor/RAID/Controller/Failed	raid: 103
	/Monitor/RAID/BBU/OK	raid: 111
	/Monitor/RAID/BBU/Degraded	raid: 112
	/Monitor/RAID/BBU/Failed	raid: 113
	/Monitor/RAID/Disk/OK	raid: 121
	/Monitor/RAID/Disk/Rebuilding	raid: 122
	/Monitor/RAID/Disk/Failed	raid: 123

Group	Device Event Category	Device Event Class ID
Temperature	/Monitor/Temperature/OK	hardware: 151
	/Monitor/Temperature/Degraded	hardware: 152
	/Monitor/Temperature/Failed	hardware: 153
Voltage	/Monitor/Sensor/Voltage/OK	hardware: 111
	/Monitor/Sensor/Voltage/Degraded	hardware: 112
	/Monitor/Sensor/Voltage/Failed	hardware: 113

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.

Starting with v4.5, audit events for alerts are only written to the Internal Storage group and not forwarded to ESM by default. If you need to forward these audit events to ESM, please contact ArcSight Customer Support for assistance. Please note that this change only applies to audit events generated for alerts; other audit events are unaffected.

Viewing Alerts

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

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

When you create Alerts (see [“Alerts” on page 255](#)), 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 92](#) for more detail.

Receiving Alerts for Events

To receive alerts:

- 1 Configure the Logger's SMTP with the desired e-mail address destination (see [“Static Routes” on page 308](#)) or create an SNMP Destination (see [“SNMP Destinations” on page 266](#)) or Syslog Destination (see [“Syslog Destinations” on page 267](#)).



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 123.)
- 3 Create an Alert that uses the new Filter and specify match count and threshold (see [“Alerts”](#) on page 255.) Enable the new Alert.

The screenshot shows the ArcSight Alerts interface. At the top, there are tabs for Monitor, Analyze, Reports, Configuration, and System Admin. The user is logged in as 'admin'. The interface displays a list of alerts with columns: Time (Event Time), Alert Name, Base Event Count, Time Threshold, and Matched Events. Below the alert list, there are detailed views for each alert, showing the Base Event (1 found) and a table of event details including Time (Event Time), Device, Logger, Name, Severity, Receipt Time, Device Vendor, and Device Product.

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 110, [“Understanding the Search Results Display”](#) on page 112, [“Viewing Alerts”](#) on page 130, and [“Advanced Search Options”](#) on page 111.

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 133](#)
[“Report Groups” on page 134](#)
[“Reports Home Page” on page 139](#)
[“Using the Dashboard” on page 140](#)
[“Running, Viewing, and Publishing Reports” on page 154](#)
[“Designing Reports” on page 164](#)
[“Scheduling Reports” on page 215](#)
[“Deploying a Report Package” on page 219](#)
[“Report Server Administration” on page 220](#)
[“Backup and Restore of Report Content” on page 222](#)

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.

The screenshot shows the 'My Reports' interface with tabs for 'Design' and 'Preferences'. Under 'Recent Reports', there is a 'Report Execution Status' section. It includes a 'Filters' dropdown and a table of reports. The table has columns for Report Name, Action, Execution Type, Status, View, Completion Time, Pages, Cancel, and Delete. One report is listed: 'SANS Top 5/SANS Top 5 -5- Alerts from IDS' with a status of 'Success' and a completion time of '10/21/2009;10:31'.

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 140](#).

Report Groups

Logger supports the following report groups:

- [“Foundation Reports” on page 135](#)—This report group contains ready-made reports that address common security use cases. This report group is displayed by default.
- [“Solution Reports” on page 137](#)—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 137](#)—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 138](#)—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 219](#).)

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 219.](#))

The Foundation Reports include [“SANS Top 5 Reports” on page 135](#), [“Network Monitoring Reports” on page 136](#), [“Intrusion Monitoring Reports” on page 136](#), [“Configuration Monitoring Reports” on page 137](#).

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 215](#).

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 154](#) and [“Task Options on Available Reports” on page 155.](#)

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 154](#).

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 154](#).

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 154](#).

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 219](#).)

For information on deploying Solutions Packages, see [“Deploying a Report Package” on page 219](#). 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 215](#).

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 154](#) and [“Task Options on Available Reports” on page 155](#).)

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 138](#), [“Cross Device Reports” on page 138](#), [“Database Reports” on page 138](#), [“Firewall Reports” on page 138](#), [“Identity Management Reports” on page 138](#), [“IDS-IPS Reports” on page 138](#), [“Network Reports” on page 138](#), [“Operating System Reports” on page 138](#), [“VPN Reports” on page 138](#).

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 215](#).

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 154](#) and [“Task Options on Available Reports” on page 155](#).)

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	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 154](#).

For information on using the Report Designer to create reports, see [“Designing Reports” on page 164](#).

For information on deploying Custom Packages, see [“Deploying a Report Package” on page 219](#).

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 141](#)).

To get started by creating a dashboard to show as your default Reports Home page, see [“Using the Dashboard” on page 140](#), [“Designing Dashboards” on page 141](#), and [“Setting Dashboard Preferences” on page 152](#).

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. A report run on ad-hoc basis is listed on the Report Execution

Status page for 24 hours, however reports run in the background are listed for a longer period of time

My Reports Design Preferences

Recent Reports

Report Execution Status

Filters Category Name [[All]] | Execution Type [[All]] | Status [[All]] | Report Name [[All]] | User [[All Users]]

Category: (All) Execution Type: (All) Status: (All)

Report: (All) Date From: 10/15/2009 To: 10/21/2009

Select User: (All Users) Refresh

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 154](#) and [“Scheduling Reports” on page 215](#).

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.

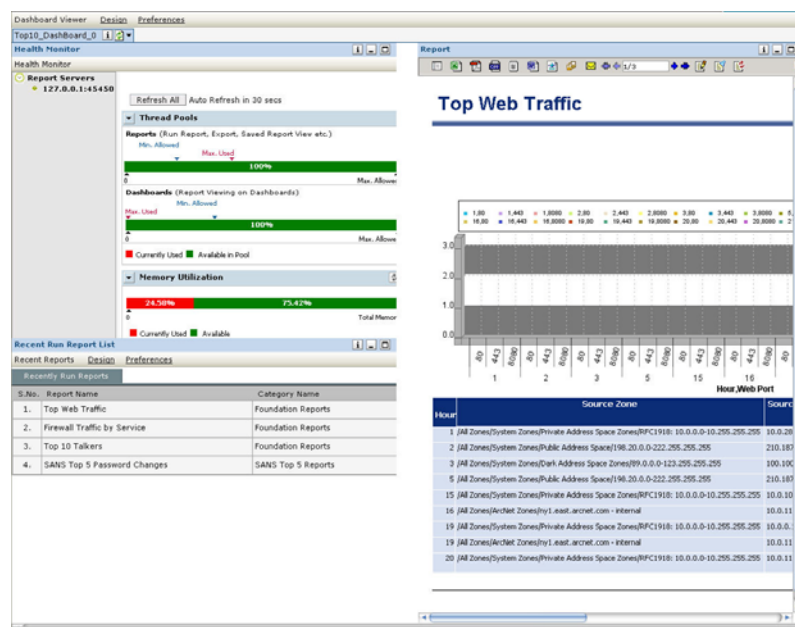


Figure 5-4 Dashboard View with Health Monitor and Reports



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 141](#) and [“Scheduling Reports” on page 215](#).

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 154](#).

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.



Note

By default, a scroll bar is not available in the Dashboard for external links. To include a scroll bar, set the "Show Scrollbar" property to "Yes" in the Widget Properties section of "External Links" under Dashboard Items.

- 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 153](#) 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.

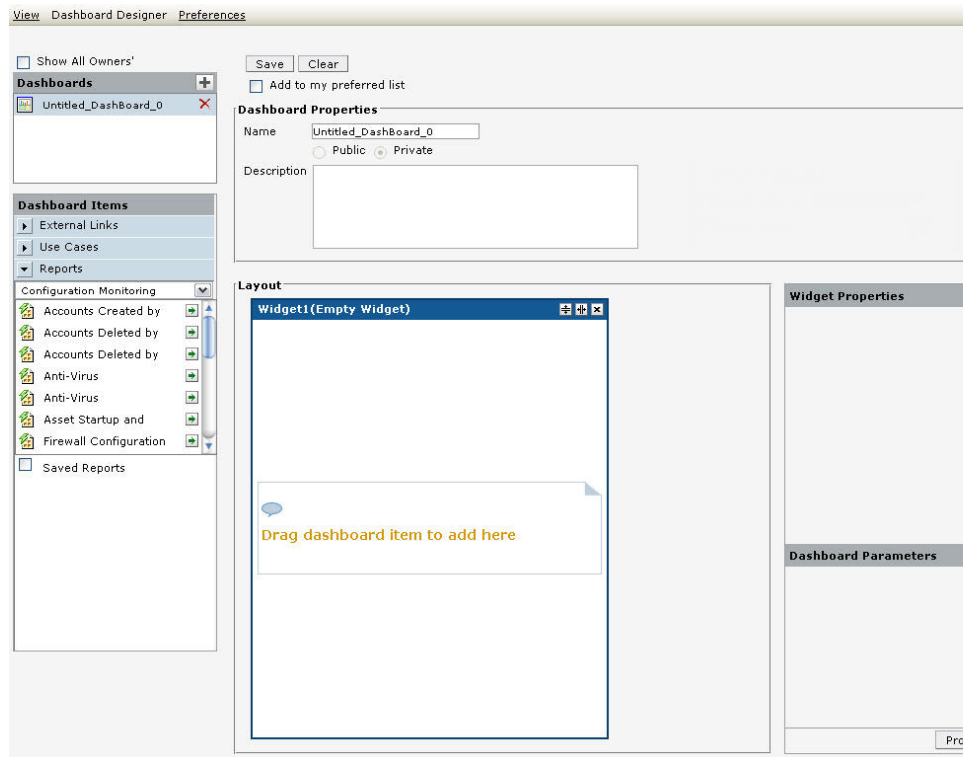


Figure 5-5 New Dashboard Layout

- 2 In the **Name** field, specify a unique name to identify the dashboard.
- 3 Specify Dashboard Properties. (For details, see [“Dashboard Properties” on page 145.](#))
- 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 153.](#))



Tip

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:

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.



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 145.](#))

You can also delete widgets you do not need. (See [“To remove a widget” on page 145.](#))

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 and 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.



Note

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.

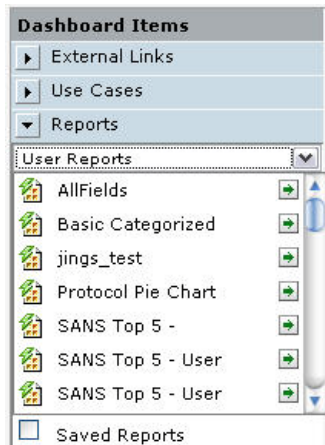


Figure 5-8 Reports under Dashboard Items


- 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 147.](#))

Widget Properties for Reports

Figure 5-9 Widget Properties for Reports on a Dashboard


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



By default, a scroll bar is not available in the Dashboard for external links. To include a scroll bar, set the “Show Scrollbar” property to “Yes” in the Widget Properties section of “External Links” under Dashboard Items.

Table 5-2 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 215.)</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

Property	Description
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 215.)</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.
Link Widgets	<p>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.</p> <p>See “Linking Widgets” on page 148.</p>
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 150.](#))

Widget Properties for Use Cases

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.
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 151.)

Widget Properties for External Links

 A screenshot of the 'Widget Properties' dialog box for an 'External Link' widget. The dialog has several fields: 'Name' (set to 'External Link'), 'Refresh Interval (in mins.)' (set to '15'), 'Auto Refresh' (set to 'YES'), 'Show Scrollbar' (set to 'NO'), 'URL' (set to 'www.arcsight.com'), and 'Description' (empty). There is a scrollbar icon next to the URL field.
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.

Property	Description
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.

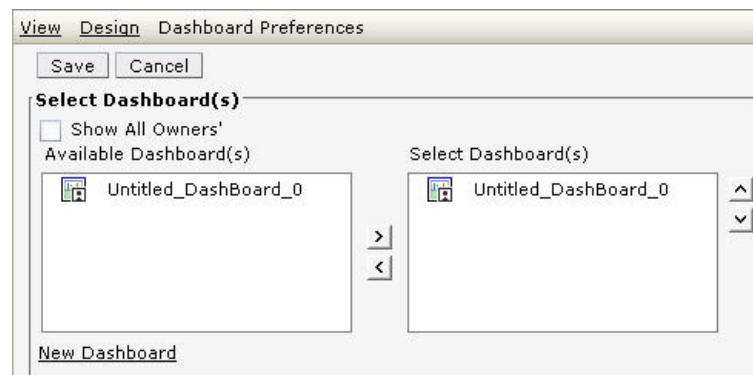


Figure 5-16 Dashboard Preferences

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 142](#).

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 142](#).

- 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 141](#).

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 141](#).

- 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 134](#))

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.

Tip

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 215](#).

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 134](#).)
- 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 149](#).) 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 215](#).)



The Search feature on the Logger “Analyze” page (described in [Chapter 4, Searching and Analyzing Events, on page 71](#)) 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.



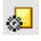





Your access to various reports and report options (view, publish, edit, etc.) depends the access rights associated with your user role and **Logger Report Group** affiliation. For example; depending on your access rights, you may have privileges to view some reports but not others, to view but not schedule 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 180](#), and [“Logger users are granted permissions by membership in a user group. A user group is a set of permissions and a set of users.” on page 337](#) in [“System Admin” on page 301](#).

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 157 and "Quick Run / Run In Background Report Parameters" on page 158.</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 157 and "Quick Run / Run In Background Report Parameters" on page 158.</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 157 and "Run Report Parameters" on page 160.</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 164, "Quick Run / Run In Background Report Parameters" on page 158, and "To publish a report" on page 161.</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 179.</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 135](#), ["Solution Reports" on page 137](#), and ["User Reports" on page 138](#).

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.

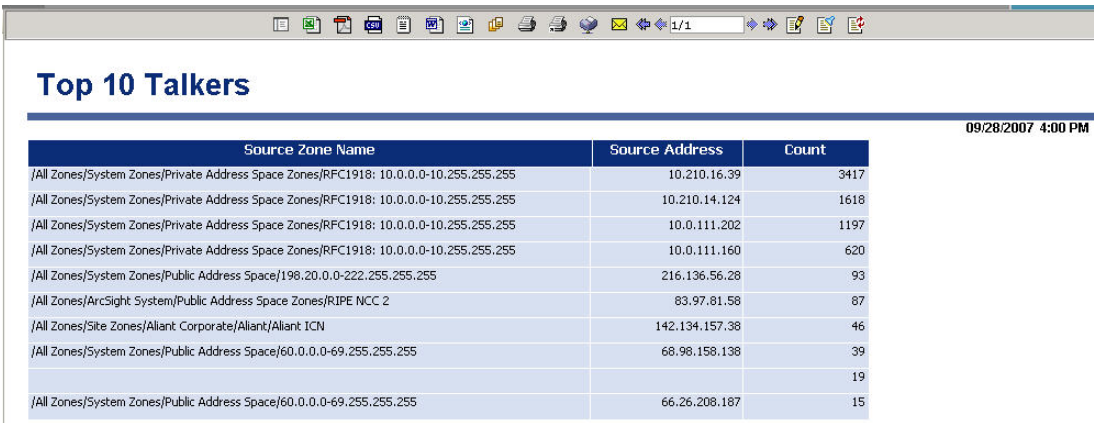
Currently, Logger limits the number of pages for horizontal pagination to 10. Consequently, if a report requires more than 10 pages to display all columns, complete report results may not be displayed. To view all columns of such reports, manually set the width of each column such that all columns fit 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).



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 161](#).

For information about other delivery options available to you at this point, see [“Report Delivery Options” on page 162](#).

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.

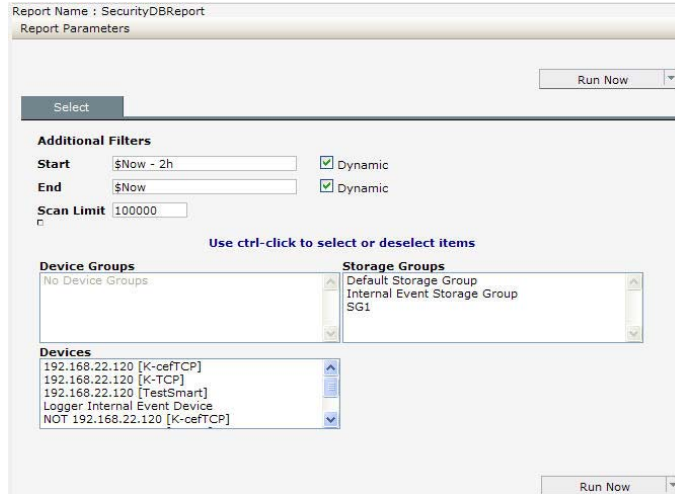


Figure 5-19 “Quick Run” / “Run in Background” Report Parameters

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>
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>

Option	Description
Scan Limit	Specify the number of events to scan. 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. The scan limit is 100,000 by default. If you set the scan limit to 0 (zero), all events are scanned. This setting does not apply to the scheduled reports.
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 159.)
Storage Groups	Select the storage group(s) on which to run the report query. (See “Selecting Device Groups, Storage Groups, or Devices” on page 159.)
Devices	Select the device(s) on which to run the report query. (See “Selecting Device Groups, Storage Groups, or Devices” on page 159.)

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.

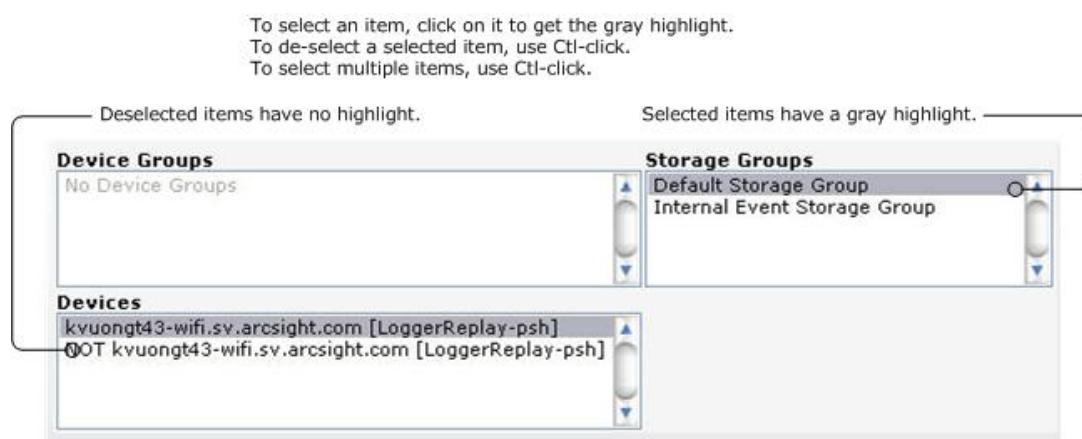



Figure 5-20 Selection Model for “Quick Run” or “Run in Background” Report Scope of Storage and Devices

- 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 Ctrl-Click.
- To de-select a currently selected item, use Ctrl-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.

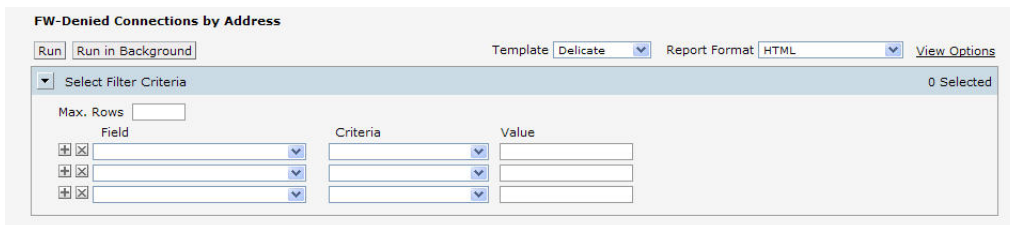


Figure 5-21 “Run” Report Parameters

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>Note: 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 161</p>
Select Filter Criteria	<p>Provides options to define filters, or modify default filters if any are already built in to the report.</p> <p>The filter expression is applied when the report runs, narrowing the focus of the report to the specified criteria.</p> <p>For example, you could set the filter criteria on a report on Top Password Changes to report only on password changes related to specified 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 170 in “Designing New Reports” on page 167.</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 158](#) 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 180](#) 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 156](#)), the output results for that run of the report are saved for subsequent.



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 215](#) and [“Scheduling Reports” on page 215](#) and [“Add Report Job Settings” on page 218](#).


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 164](#).

Report Delivery Options

When you run a report (as described in ["Running and Viewing Reports" on page 156](#)), many options are available for delivering the generated output.

The most common next step is to publish the resulting report (described in ["Publishing Reports" on page 161](#)), 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 e-mail options on *scheduled reports*, as described in ["Scheduling Reports" on page 215](#) and ["Add Report Job Settings" on page 218](#).

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
Report <%MENU_NAME%> has been generated.
Please click the following link to view the report in
<%REPORT_FORMAT%> format.

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	Choose one of these: <ul style="list-style-type: none"> To provide a link to the report in the body of the e-mail, select Link. To send the report as an attachment to the e-mail, click Attachment, and select a format for the attachment file.
To and CC	Specify e-mail addresses to which to send the report.
Subject	Provide e-mail Subject header.
Message	For the body of the e-mail, 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

- Click the  (Export) button or click one of the file formats on the published report top-level menu bar (     )
- In the Export Options dialog, specify the Export Format and associated settings you want in the Export Options dialog.

Export Options

Export Formats : MS EXCEL

☐ MultiSheet

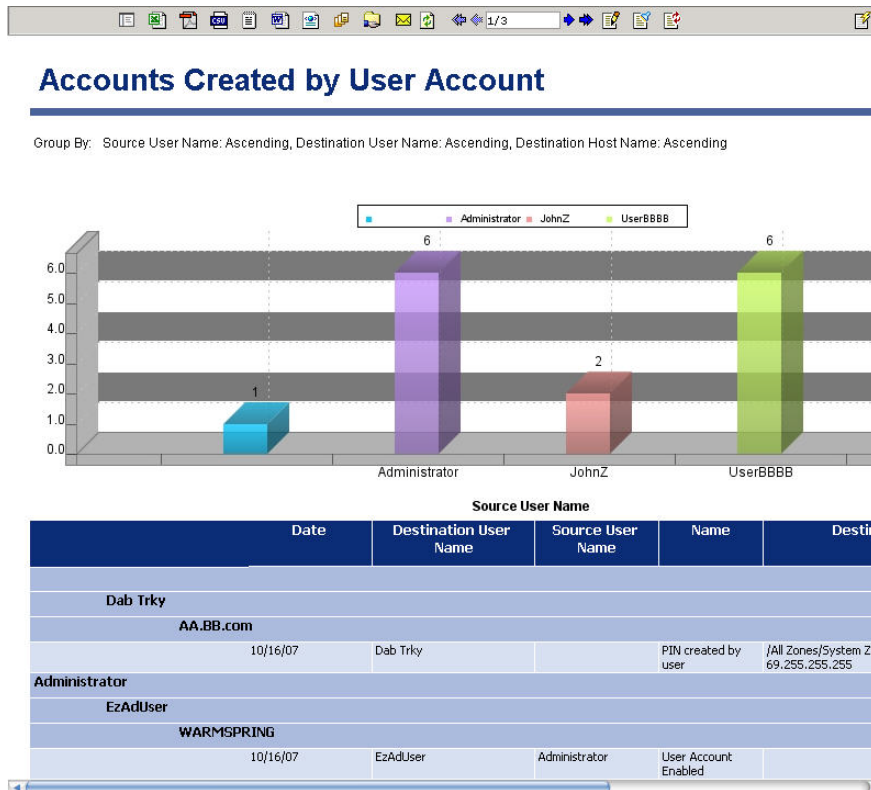
☐ Remove Blank Rows

☐ Repeat Page Header and Footer


Generate Cancel

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 155 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.

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		

Buttons: Refresh, Show All Owners'

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 134](#) and [“Task Options on Available Reports” on page 155](#) 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 165](#).

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 167](#).

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 134](#) 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							

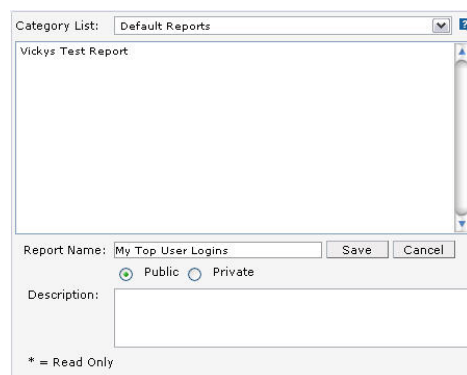
- 3 In the Report Designer for the selected report, click **Save As**.

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 shows the 'Category List' set to 'Default Reports'. The report name is 'My Top User Logins'. The 'Public' radio button is selected. The description field is empty. There are 'Save' and 'Cancel' buttons.

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.

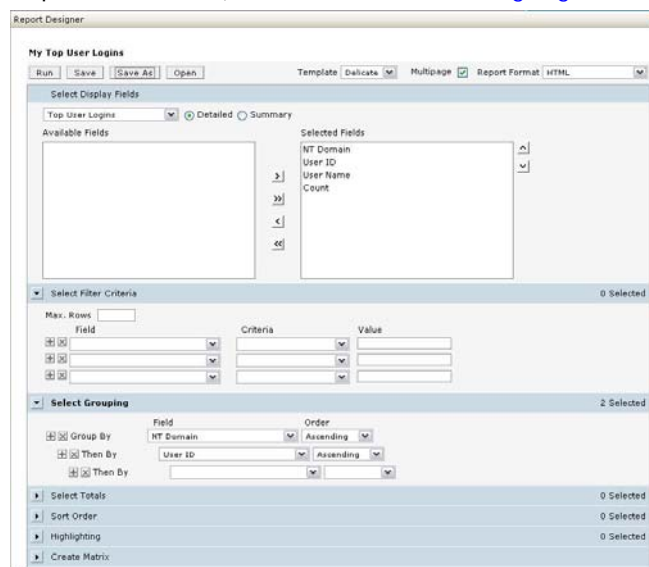
- 6 Click **Save** to save the report.

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 167.](#))



The Report Designer interface shows the 'My Top User Logins' report. It includes tabs for 'Run', 'Save', 'Save As', and 'Open'. The 'Template' is set to 'Default'. The 'Report Format' is set to 'HTML'. The 'Select Display Fields' section shows 'Top User Logins' selected. The 'Available Fields' list includes 'NT Domain', 'User ID', 'User Name', and 'Count'. The 'Selected Fields' list includes '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 'Then By' set to 'User ID'. 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 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. These templates define the look and feel, arrangement, orientation, and so on, of the report output. To include the start time, end time, scan limit, device group, storage group, and devices information (used to run a report) in a report, choose the "BlankWithHeader" template.</p> <p>See "Applying Report Template Styles" on page 214 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 161 .
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 181](#).)



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 119](#).




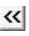


Under **Select Display Fields**, enter a meaningful title for the report (in the Report Title field) and select whether the report contents should be Detailed or Summarized (in the Report Contents field). The report title is the text that appears as the title on top of a report.

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.



Note

- 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.



Tip

For information on how to create query objects for use in reports, see [“Setting up Queries” on page 181](#). 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
 		<input type="text"/>
 		<input type="text"/>
 		<input type="text"/>

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



Tip

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 160](#).





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.

The screenshot shows a panel titled "Select Filter Criteria*". At the top, there is a "Max. Rows" input field. Below it is a table with three columns: "Field", "Criteria", and "Value". The first row in the table has "Time" in the "Field" column, "Is" in the "Criteria" column, and a red asterisk (*) in the "Value" column. There are also plus and minus icons to the left of each row in the table.

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 191](#) under [“Setting up Queries” on page 181](#).

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 172), 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 191 under "Setting up Queries" on page 181.
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

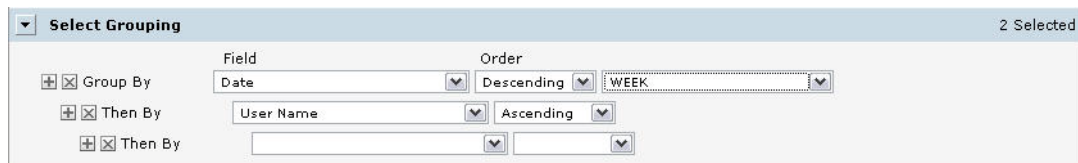


Figure 5-31 Grouping Items by Field in a Report

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.





- If you select set Max. Rows under **Select Filter Criteria** (as described in ["Select Filter Criteria" on page 170](#)) and also specify grouping, you may get a different result than if you just specified Max. Rows without grouping.
- A report that has a group defined can only display up to 100,000 lines.

To define a group

- 1 Select a field by which you want to group (as described in [Table 5-20 on page 218](#)).
- 2 Select the order of arrangement of group (as described in [Table 5-20 on page 218](#)).

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 174](#).

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



Note

A report that has a sort order defined can only display up to 100,000 lines.

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”).

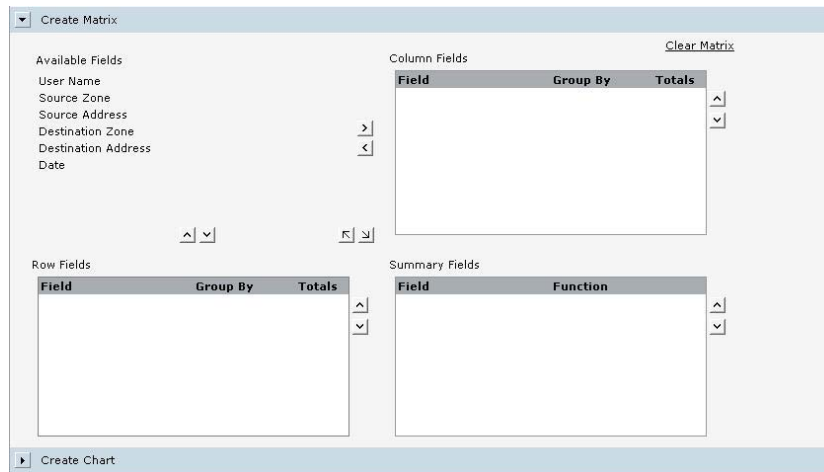


Figure 5-35 Adding a Matrix to a Report

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:

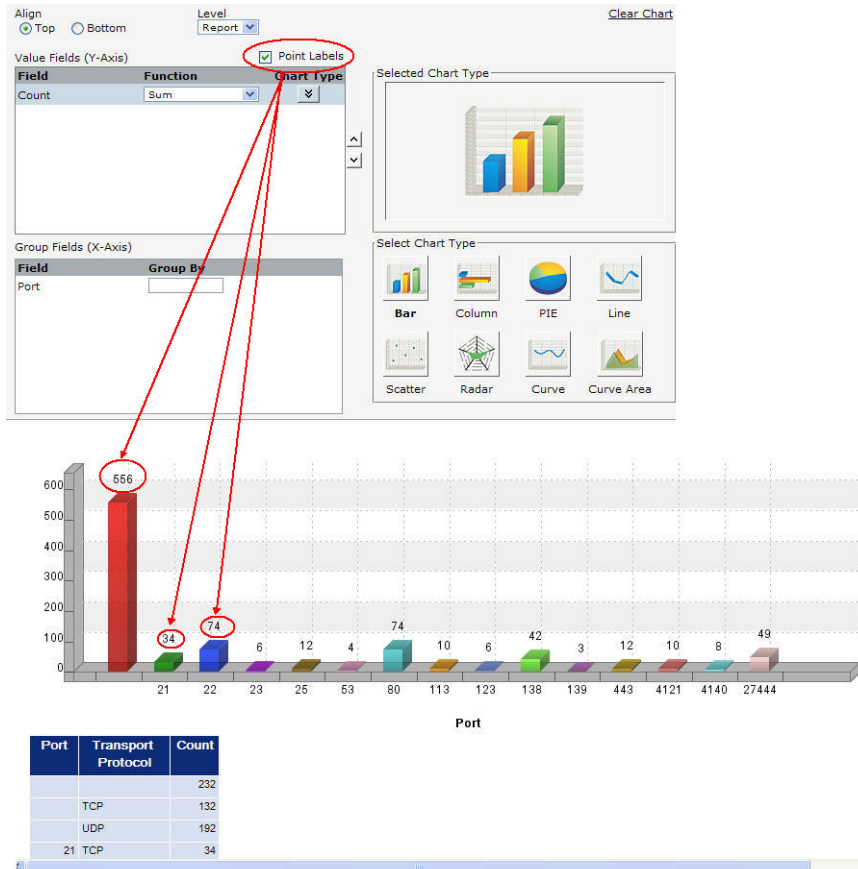
- Select **Top** to place the chart above other components.
- Select **Bottom** to place the chart below other components.
- In **Level**, select **PAGE** to plot chart having page level data. Select **REPORT** to plot chart from data that has come from entire report.

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 165](#)).
- 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 179](#).)
- 4 Click **Save**.



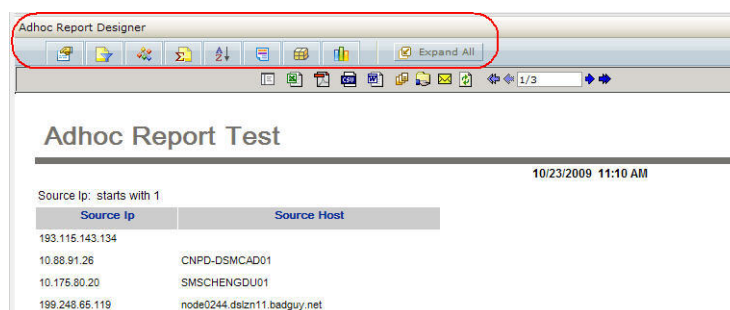
If a user has access rights to “view, run, and schedule all reports”, you can create **private** reports. If you do not have permissions to edit a **public** report that you want to modify but you do have permissions to create private reports, then you can save the public report as a private one and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

For more about publishing a report as “public” or “private”, see [Table 5-8 on page 162](#). For more about “access rights” on reports, see [“Setting Access Rights on Reports” on page 180](#).

See also [“Quick Start: Base a New Report on an Existing One” on page 165](#).

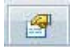







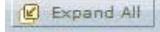
Adhoc Report Designer

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 169 for more information.
	Specify filter criteria. See "Select Filter Criteria" on page 170 for more information.
	Specify grouping. See "Select Grouping" on page 172 for more information.
	Specify the summary (total) fields. See "Select Totals" on page 174 for more information.
	Specify sort order. See "Sort Order" on page 174 for more information.
	Set up highlighting. See "Highlighting" on page 175 for more information.
	Include a matrix. See "Create Matrix" on page 176 for more information.
	Create a chart. See "Create Chart" on page 177 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 155.](#))

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 ["Logger users are granted permissions by membership in a user group. A user group is a set of permissions and a set of users." on page 337 in Chapter 7, System Admin, on page 301.](#)

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 205](#).

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 181](#), 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%"
```

Fields

- User Name
- Source Zone
- Source Address
- Destination Zone
- Destination Address
- Date

Field User Name

Caption User Name HyperLink ... Hidden

Data Type CHAR Group Label (Select to add group label)

Format

Width 10 Output Format ...

Align Left Input Format ...

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 134](#)).

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 110](#) 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.)

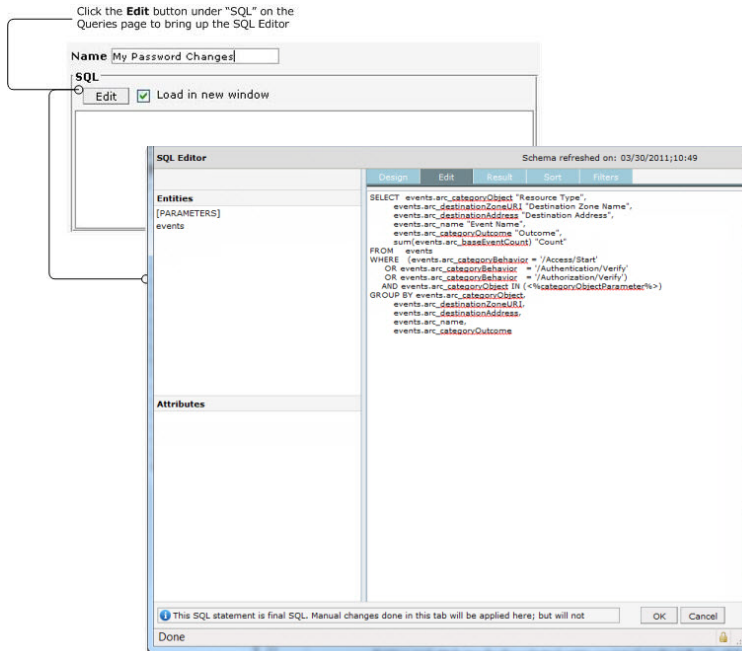


Figure 5-38 Query SQL Editor

- 4 Use the SQL Editor to define the query statement. (See [“Defining SQL in the Editor”](#) on page 198.) 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.

The **Data Format** dialog box has a left sidebar with a list of format categories: General, Number, Currency, Date, Time, Percentage, Scientific, Text, and Network Id. The **Network Id** category is selected and highlighted. To the right, under the **Formats** heading, there is a list box containing 'IP Address (IPv4)' and 'MAC Address'. Below the list box, a **Sample:** 127.0.0.1 and a **Format:** x.x.x.x are displayed. At the bottom right are **OK** and **Cancel** buttons.




The default date/time in reports does not include the time of day. You must choose a date format that includes HH:MM:SS to include the time.

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

Characters for the selected format are displayed in the Output Format entry field.

Specifying Input Format for a Field

If you specify the input format for a query field here, at run-time the report containing this query will accept data only in the format specified.

- 1 From Fields list, click (select) the field for which you want to define an input format. (The selected field is bold.)
- 2 Click  button next to the Input Format field to launch the Data Format dialog.

This is a duplicate of the Data Format dialog box shown in the previous image. It shows the **Network Id** category selected in the sidebar, with 'IP Address (IPv4)' and 'MAC Address' in the formats list. The sample is 127.0.0.1 and the format is x.x.x.x.

- 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.


The **Format** configuration panel includes a **Width** field with the value 15 and an **Align** dropdown menu set to **Right**. On the right side, there are two fields: **Output Format** and **Input Format**, both containing the text 'x.x.x.x'. Each of these fields has a small button with three dots (a menu icon) to its right.

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.

Group Label  (Select to add group label)

- 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.

- Specify the formula and click **OK** to save the formula and close the dialog. (See [“Syntax for Formulas” on page 188.](#))



- ◆ 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 188.](#))
- 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 prefixed 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.



Note

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.

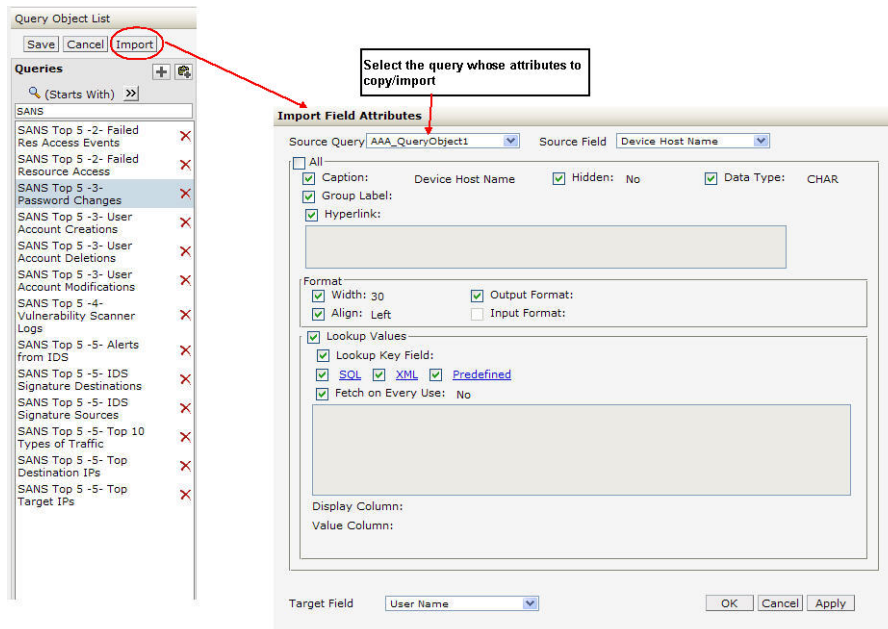


Figure 5-40 Importing Attributes from One Query into Another

- 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 170](#) under [“Designing Reports” on page 164.](#))

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 URLProvide 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 194. <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.)

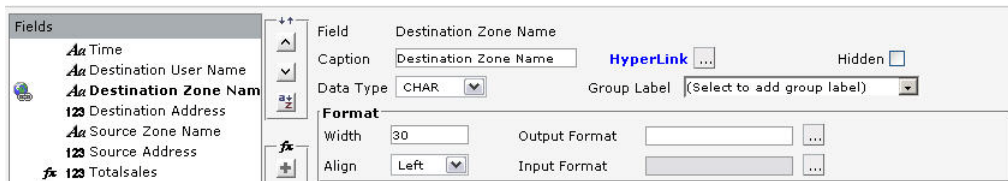


Figure 5-43 Hyperlink Options

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 161)

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.

The screenshot shows a 'Lookup Values' dialog box. At the top, there is a checked checkbox labeled 'Lookup Values'. Below it, the 'Lookup Key Field' is set to 'Destination User Name'. There are two checkboxes: 'SQL' (unchecked) and 'Predefined' (checked). Below these are two text input fields, one labeled 'Display' and one labeled 'Value'. At the bottom of the dialog is a large list box labeled 'Values'.

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 183](#).


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 181](#)) 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.

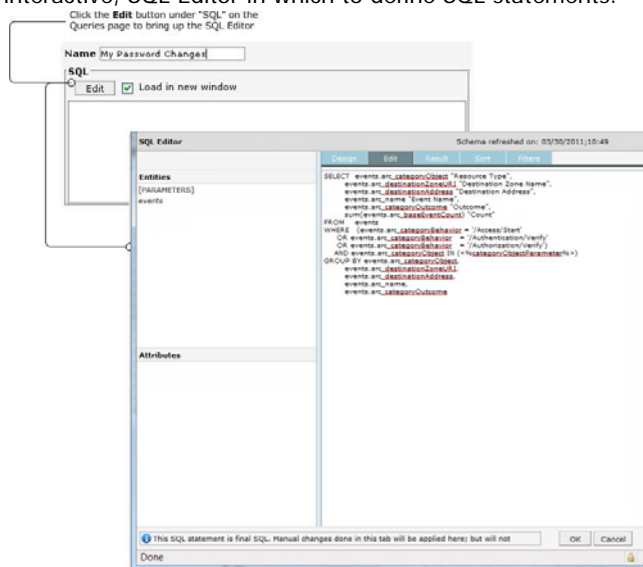


Figure 5-46 Accessing SQL Editor on the Reports | Queries page

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.



Note

The Attributes list shows a few attributes that are internal to Logger. They should not be used in queries because the resulting report will not contain expected results. All attributes listed after `arc_sourceZoneResource` are internal, including `arc_eventTime`, `arc_deviceName`, `arc_rowId`, and `arc_others`.

Table 5-14 SQL Editor Tabs

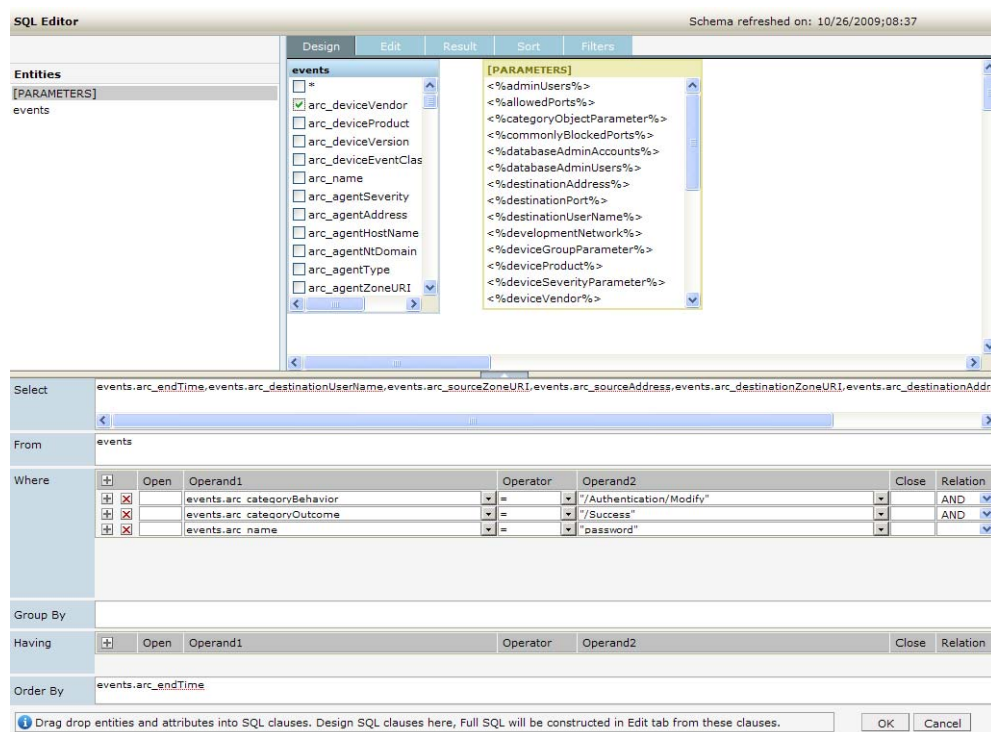
Option	Description
Design	Graphical SQL query designer. Use options on this tab to design relatively simpler queries using drag and drop method.
Edit	Shows the SQL statements. A query created on the Design Tab is represented as an SQL statement on this tab. You can also write or paste and SQL directly here.
Results	Displays rows received as a result of SQL execution.
Sort	Specify sorting preferences.
Filters	Add filters to set run-time filter criteria to be included in the query.

List of Database Objects

The SQL Editor shows the **Default Connection** to the database that provides the database objects list. 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.



Note

The Attributes list shows a few attributes that are internal to Logger. They should not be used in queries because the resulting report will not contain expected results. All attributes listed after `arc_sourceZoneResource` are internal, including `arc_eventTime`, `arc_deviceName`, `arc_rowId`, and `arc_others`.

- 3 Repeat these steps to select other attributes from different entities.



Tip




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 200](#).



Note

Be sure to include appropriate summary function in “Select” clause so that it can be used in the “Having” clause.

Order By

In the Order By clause you can provide sorting (ascending/ descending) criteria for the SQL statement. For a report with grouping, the "Order By" clause must have the columns in the same order as the respective sections in the Layout Editor.



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.

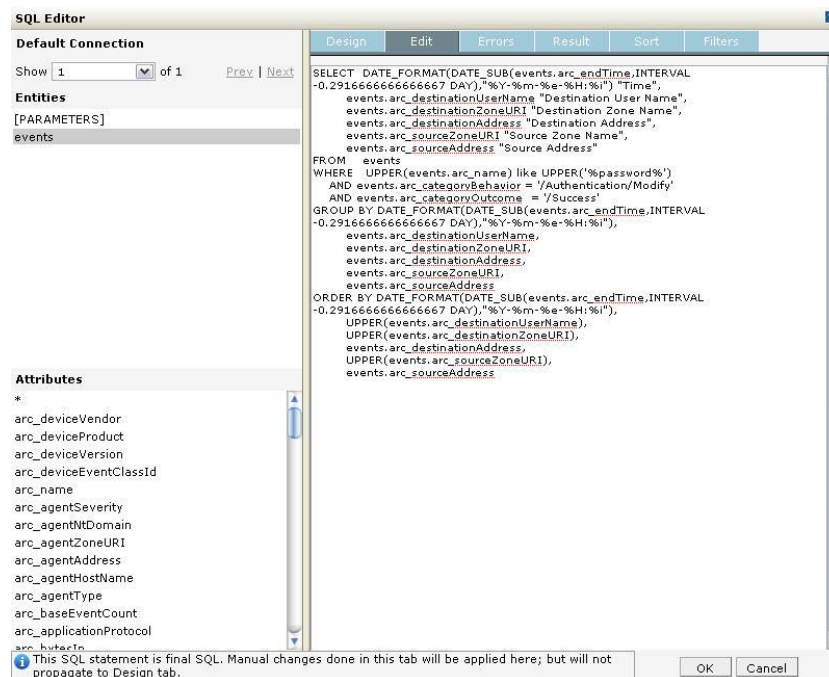


Figure 5-48 SQL Editor: Edit Tab

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.

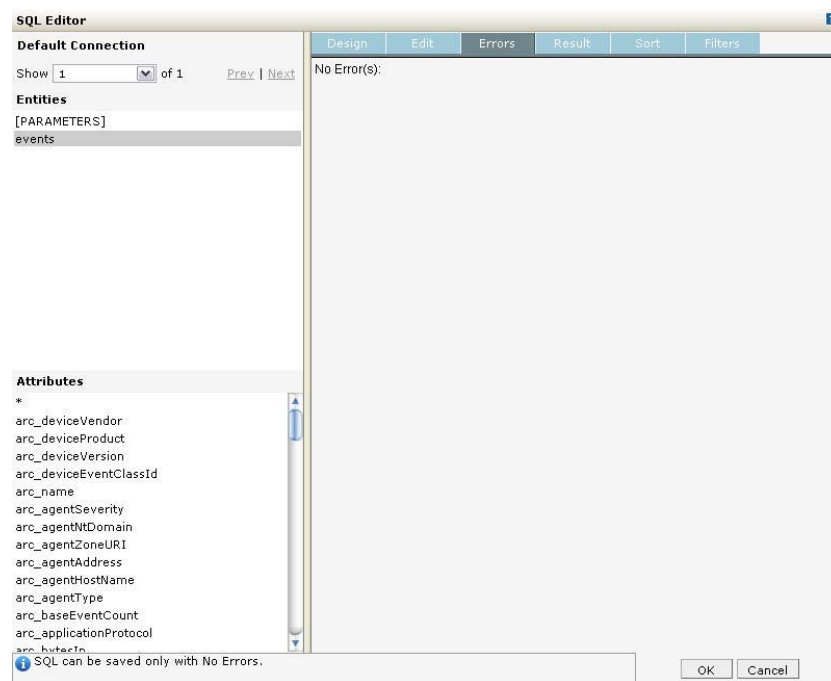


Figure 5-49 SQL Editor: Errors Tab

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.

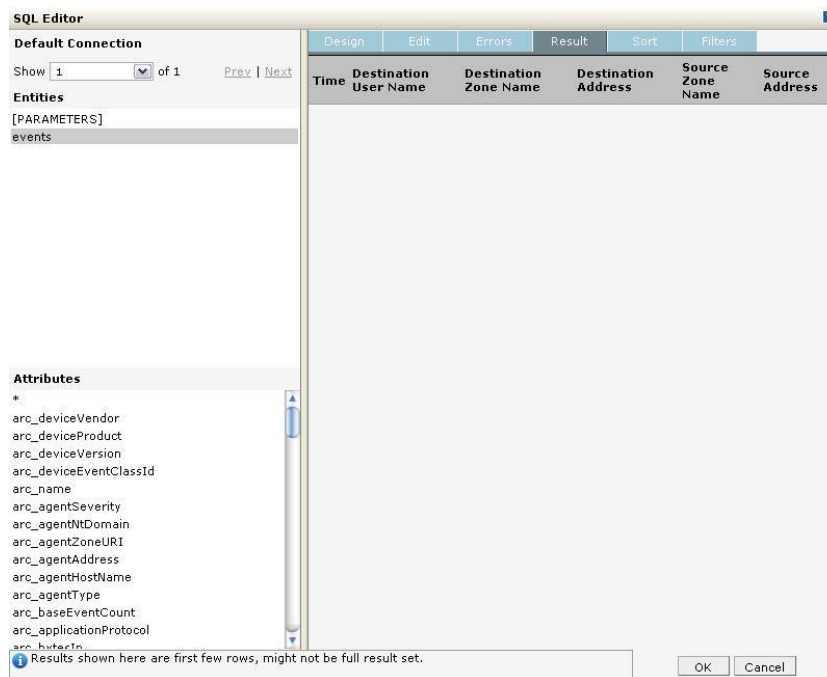


Figure 5-50 SQL Editor: Result Tab

Sort Tab

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

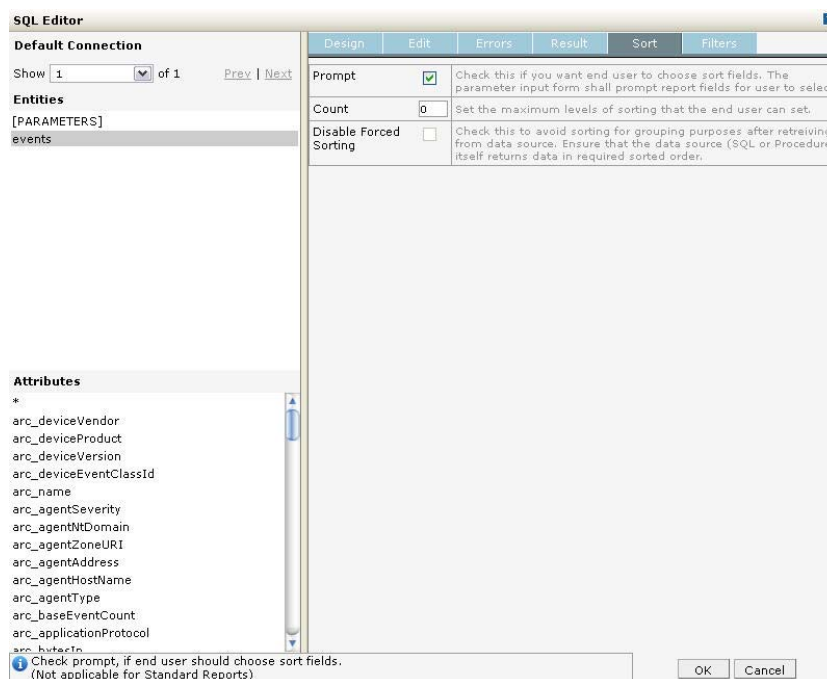


Figure 5-51 SQL Editor: Sort Tab

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.

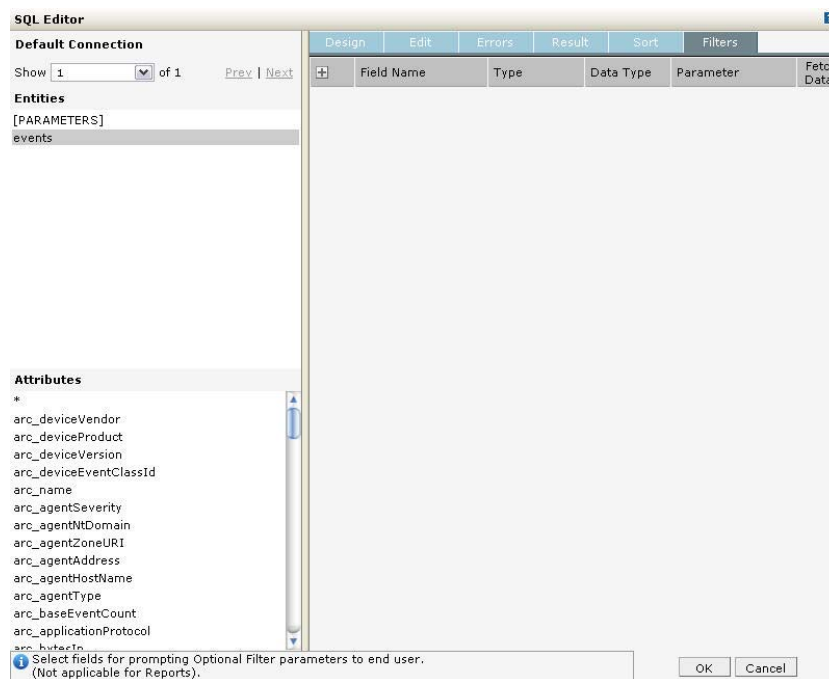




Figure 5-52 SQL Editor: Filters Tab


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 181.](#))

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.

Figure 5-53 Report Parameters Object List

Creating New Parameters



Note

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.

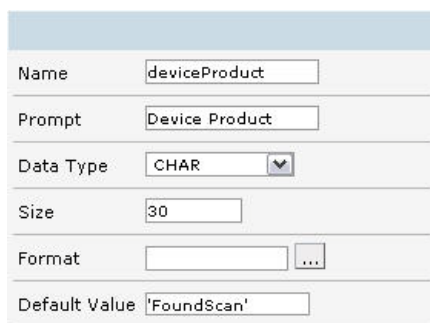
- 1 On the Parameter Object List, click  at the top right of the **Parameters** list box.

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




The screenshot shows a configuration form for a parameter. It includes the following fields:

- Name:** A text box containing 'deviceProduct'.
- Prompt:** A text box containing 'Device Product'.
- Data Type:** A dropdown menu with 'CHAR' selected.
- Size:** A text box containing '30'.
- Format:** A text box with a small '...' button to its right.
- Default Value:** A text box containing '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 209.)
Size	Specify number of characters or digits this parameter should accept. Note: This is only applicable to CHAR and NUMBER data types, not for Boolean or Date type parameters.
Format	Select the appropriate format in which user should provide value for this parameter. Click  to open a Data Format dialog box. Based on the format you have selected, a format string will appear in the entry box.
Default Value	Specify a default value that is appropriate in most cases to provide for this parameter at report run time. The default value will be automatically selected at report run time. The user can change the default value, if needed. If the user does not change it, the report will run using the default value you specify here for this parameter.

Default Value for Date Type Parameter

For a date type parameter, the **Default Value** field provides 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 rows: '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 to its right; and 'Default Value' with a dropdown set to 'CURRENT_DATE' and a '±Days' field with a small icon to its right.

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 control labeled 'Input Type' with three radio buttons: 'TextBox' (which is selected), 'Combo', and 'Option'.

Figure 5-55 Parameter Input Type

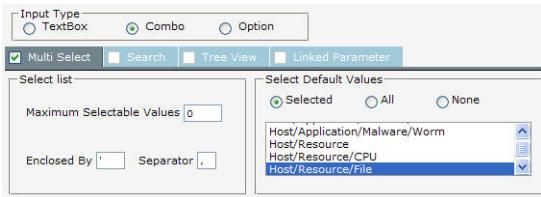
The parameter *input type* describes the style of interface provided to users at report run time in which to enter a value for this parameter. Choose from **Text Box**, **Combo**, or **Option** as described below.



In the Reports Designer, changing the parameter type TextBox to another type causes an error. If you need to change the parameter type to TextBox, do not edit an existing parameter, delete that parameter and add a new one.

Table 5-17 Input Type

Option	Description
Text Box	Select "Text Box" input type if you want the user to type the value for the parameter.

Option	Description
Combo	<p>Select “Combo” if you want the user to select one value or multiple values from a drop-down menu.</p> <p>Select the Multiselect checkbox so that user can select multiple values from the box.</p>  <p>See “Setting Multiple Default Values” on page 211 to configure other settings for this option.</p>
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>

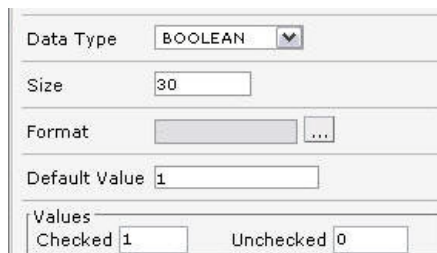
Setting up Boolean Parameters

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).



Setting Various Run Time Behaviors

You can specify a variety of options on how the parameter will look and act at report run time. These options are generally related to the input type, and further define acceptable user input values, whether the parameter will be displayed or hidden, provide searchable values, and so forth.

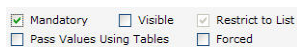


Table 5-18 Parameter Options

Option	Description
Mandatory	Select this checkbox if you want to require the user to specify a value for this parameter at report run time.
Visible	Select this checkbox if you want the parameter to be visible (displayed) on the input form at report run time. Keep this unchecked (clear) if the value for this parameter be populated from another report or if you want the parameter to use the default value in all cases.
Restrict to List	This setting is applicable for parameters with Input Type of Combo . Select (check) the Restrict to List checkbox here to force user input of a parameter value from the available run-time options only. If Restrict to List is <i>not selected</i> in the parameter definition you create here, the user can specify a value or can select value(s) from available options.
Pass Values Using Tables	This setting is applicable for Multi Select . Select this checkbox when you want to pass parameter values through a table. This is done especially when the number of values that can be passed (total number of bytes of selected values) as part of the SQL is more than allowed.
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/Dat
Host/Application/Service	/Host/Application/Service
Host/Application/Service/Email	/Host/Application/Service/Email
Host/Application/Service/Instan	/Host/Application/Service/Insta
Host/Application/Service/MMS	/Host/Application/Service/MMS
Host/Application/Service/Peer t	/Host/Application/Service/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 .)

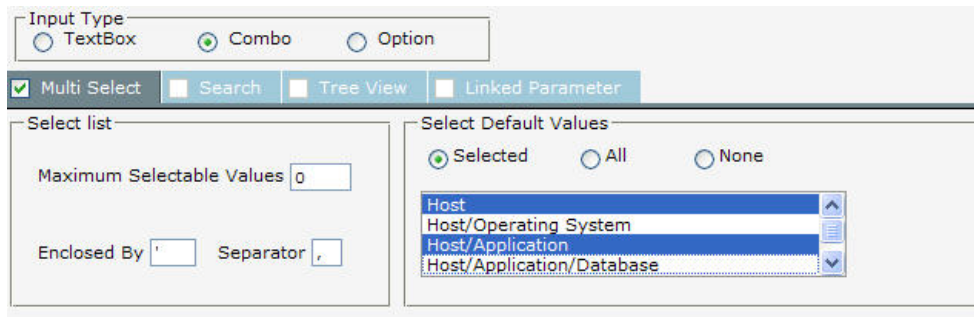
- 4 Repeat these steps for each option.

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 208](#)), 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 206](#)) 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.

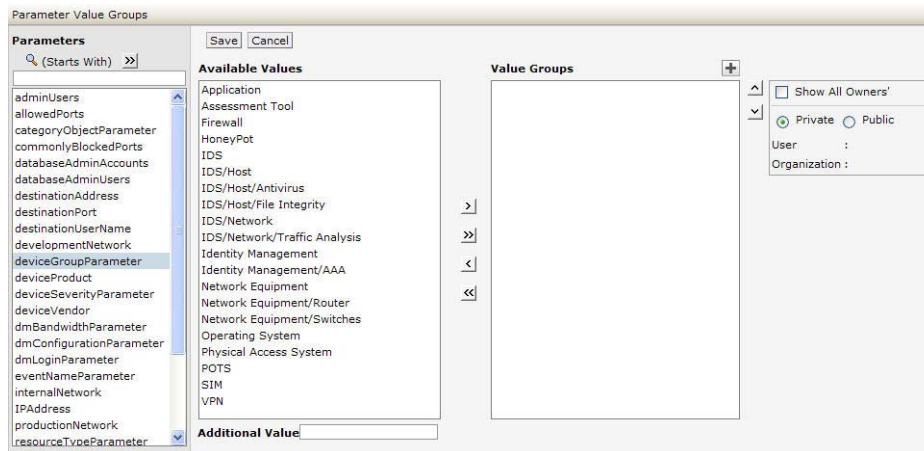


Figure 5-57 Parameter Value Groups

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.

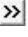


Table 5-19 Parameter Value Groups





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.
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).

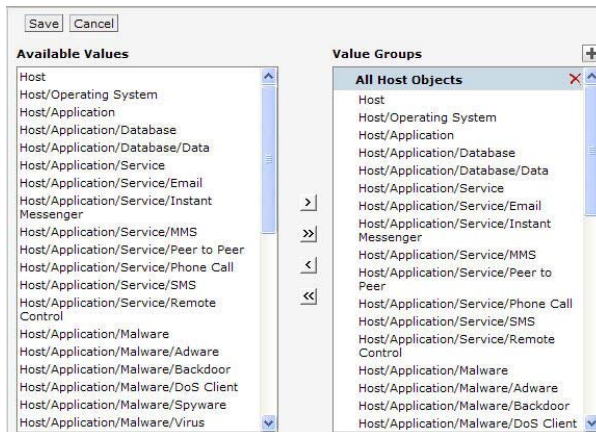
Click **Save** button to save the work.




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.

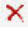
To create a tree view parameter

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 215](#) for more information.

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

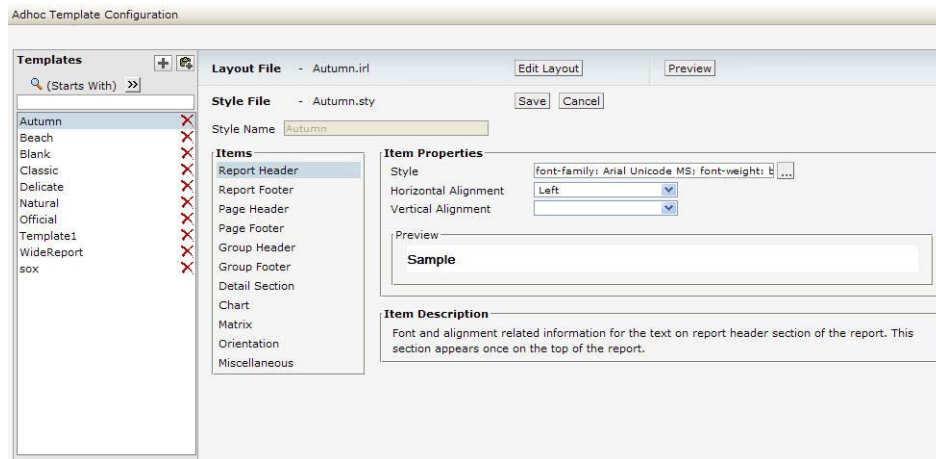


Figure 5-58 Report Template Styles Configuration


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 154](#).

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 307](#) 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.




To view scheduled reports, a user must belong to a Logger Reports Group, a Logger Search Group, and a Logger Rights Group.

Add				
Task	Type	Schedule	Next Run Time	
Password Changes	Report	Sunday at 23:00	Sun Sep 30 23:00:00 PDT 2007	 
Top 10 Talkers	Report	Saturday at 23:00	Sat Sep 29 23:00:00 PDT 2007	 
Top User Logins	Report	Daily at 23:00	Sat Sep 29 23:00:00 PDT 2007	 

Figure 5-59 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 216](#).



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 216](#).

To remove a scheduled report

Click  (Delete) next to the scheduled report job you want to remove.



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 307](#) 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](#).)

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

Add Report Job

Name:

Schedule: Hours

Report Name *  SANS Top 5/SANS Top 5 Vulnerability Scanner Log

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

kvuongt43-wifi.sv.arcsight.com [LoggerReplay-psh]

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

- 3 On the Add Report Job page, use the drop-down menu next to **Report Name** to select a report, and click **Go** to load the report.



You must click **Go** to load the selected report at the Report Name field before you save the scheduled report job. Attempting to save the scheduled job without first loading the report name will result in an error, and the report will not be saved.

- 4 Choose one or both delivery options under Delivery Option (**Email**, **Publish**). By default, only Publish is selected.

If you want to keep the Publish option, enter its associated parameters. Or uncheck it to unselect.

If you also want the Email option, click it and enter its associated parameters.

Both E-mail and Publish options for scheduled reports are the same as those provided after you run a report "on demand".

- ◆ **Email** - For details on setting e-mail delivery options, see [“E-mailing a Report” on page 162.](#)

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 161.](#)

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 218.](#)
- 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 162. • Publish - For details on setting publishing options, see "Publishing Reports" on page 161.
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 161.</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 158.</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**.
- 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.



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 137](#).

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 270](#) 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 284](#).

Chapter 6

Configuration

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

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

In this chapter:

- [“Devices” on page 223](#)
- [“Event Archives” on page 226](#)
- [“Storage” on page 233](#)
- [“Event Input/Output” on page 238](#)
- [“Alerts” on page 255](#)
- [“Scheduled Tasks” on page 269](#)
- [“Filters” on page 270](#)
- [“Saved Searches” on page 273](#)
- [“Search Optimization” on page 278](#)
- [“Peer Loggers” on page 280](#)
- [“Configuration Backup and Restore” on page 284](#)
- [“System Maintenance” on page 287](#)
- [“License Information” on page 295](#)
- [“Retrieve Logs” on page 296](#)
- [“Exporting and Importing Content” on page 297](#)

Devices

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

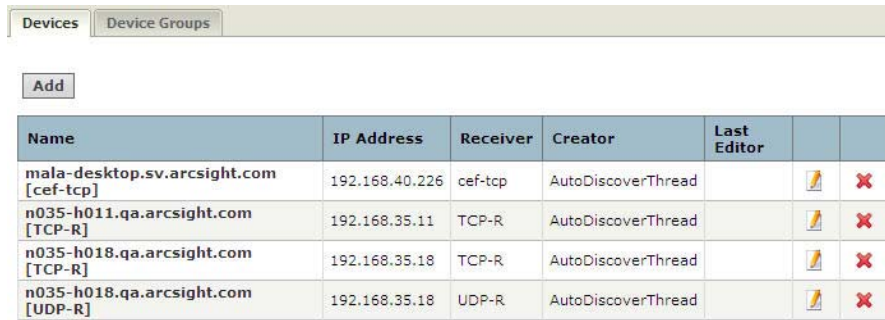
Devices

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. Event source is the device that directly sends the event to Logger. When an event is sent through a SmartConnector, the event source is the system

on which the SmartConnector is running and not the device that sent the event to the SmartConnector.

Devices can be added to Device Groups, and Device Groups can be referenced in filters and queries. Receivers perform *autodiscovery* by automatically creating a Device for each source IP address. Devices created by autodiscovery are named for their hostname, or if the hostname cannot be determined, their IP address.

Figure 6-1 shows the Devices page, which displays all defined Devices and includes controls to add, edit, or delete them.



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 235](#).


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.

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 enable you to save the events for any day in the past, *not including the current day*. Archive Storage Settings must be configured before Event Archives can be created. Archive Storage Settings specify the location to which event archives will be written.

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 Loggers**, the location is a directory (either local or a mount point that you have already established on the machine on which the Logger software is installed).

Starting with Logger v5.0 SP1, events are archived on a per storage group basis. That is, one archive file is created for each storage group, for each day. In addition, you can bulk archive events—that is, specify a range of dates to archive events in a single archive operation.

Archiving events from each storage group to a separate archive location enables you to keep data in specific storage groups longer than others. You need to specify these locations

when you configure the Archive Storage Settings before archiving any events, as shown in the following figure.



The above figure is from a Logger appliance. The Mount Location field is not available on a software Logger.

For Logger appliances, the path you specify in the Archive Path field is appended to the path specified in the Mount Location. On a software Logger, you need to enter a complete path where the archive file will be written in the Archive Path field. This path could be a local directory or a mount point that is already established on the machine on which the Logger software is installed. The Mount Location field is not available on a software Logger.

Logger uses the receipt time of an event to determine its archival day. For example, an event with timestamp of 11:55:00 p.m. on December 7 is received at 12:01:00 a.m. on December 8 on the Logger. This event is archived in the archive file created for December 8th and not December 7th. When an archive operation occurs, one archive file per storage group is created at the location specified in Archive Storage Settings. Each archive file contains events from 12:00:00 a.m. to 11:59:59 p.m. for a single storage group of any given day. When you specify a range of dates, one archive file per storage group, for each specified day is created.

You can archive events in two ways: manually and scheduled. When archiving events manually, you specify the start and end dates of the event archive, and the storage groups that should be archived. This operation occurs one-time, for the specified date range. When scheduling event archives, you specify the time at which the archive operation should occur everyday and select the storage groups that should be included.



You cannot set event archives to start at 1 a.m. for scheduled archives. This restriction is by design to account for the Daylight Savings Time (DST) changes.

When Logger starts archiving, it proceeds sequentially through the various storage groups, as listed on the Daily Task Settings page (for scheduled archives) or the Add Event Archives page (for manual archives).

Once the events have been archived, they are not deleted from the local storage until the events (and their related indexing information) age out due to the configured retention policy. These events continue to be included in search operations until they age out.

Once events that have been archived are deleted from Logger's local storage, they are not included in search operations. To include such events in search operations, you must load the archive in which those events exist back to the Logger. When an Event Archive is loaded, its events are included in searches, but the archive itself remains on the remote storage.

When events are archived, index information for those events is not archived. Therefore, when event archives are loaded, indices are not available. As a result, a search query that runs on archived events (that have been loaded on Logger) is slower than when the data was not archived because the index data for the archived data is not available.

Guidelines for Archiving Events

- 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.
- Starting with v5.0 SP1, multiple archiving operations such as loading, unloading, archiving, and deletion of archives can occur simultaneously. Therefore, you can initiate the loading of an existing archive, while an archive operation is in progress.
- Only one manual archive job can run at a time. However, a scheduled archiving operation can run in parallel with a manual job.
- You cannot re-archive the events that have been archived already. If you try to do so, the Logger reports an error.
- Do not move the archived files from their archive location. The archives that have been moved from the originally archived location cannot be loaded on to the Logger. If you need to delete the archives, use the Logger user interface to do so.
- If an archive job fails, you need to initiate it manually. To do so, delete the failed archive and archive it manually. To get notified of a failed archive, configure an alert for this audit event: Event Archive Failed. For more information about this event, see [“Logger Application Events” on page 490](#). For more information about configuring alerts, see [“Alerts” on page 255](#).
- If a Logger appliance goes down while an archive operation is in progress, you need to re-initiate the archive operation for only the storage groups that were not archived when the operation failed. The status of such storage groups is marked “Failed” in the Status column on the Event Archives page. For example, you archive the event data of 12/1/10, which consists of events from four storage groups “Default”, “Internal”, “Short-Term”, and “Long-Term”. The appliance goes down after the events from the “Default” and “Internal” groups have been successfully archived, and the events from “Short-Term” are being archived. The status of the “Short-Term” storage group on the Event Archives page will display “Failed”, while the status of the “Default” and “Internal” groups will display “Archived”. (The status of the “Long-Term” storage group

will not be displayed.) In this case, you need to manually re-initiate the archive for the “Short-Term” and “Long-Term” storage groups.



Note

In the above example, the status of the “Long-Term” storage group is not displayed on the Event Archives page after the failure occurs because archival of this group was never initiated during that archive operation.

If an archive operation fails, make sure you determine the storage groups that could not be archived and re-initiate the archival for all of those groups manually.

- You can cancel an in-progress archive operation that was manually initiated at any time using the Cancel link that displays on top of the Event Archives page.
- If you are upgrading from v5.0 Patch 2 or earlier, you need to note the following changes:
 - ◆ The existing event archives **cannot** be converted to make use of the storage-group level granularity that v5.0 SP1 offers. However, any data archived after you upgrade to v5.0 SP1, will be archived using the storage-group level granularity. Therefore, **after upgrading to v5.0 SP1**, specify the archive locations for each storage group on the Archive Storage Settings page. By default, the location you had configured before the upgrade is used for all storage groups.
 - ◆ Starting with v5.0 SP1, the archive locations specified on the Archive Storage Settings page can be changed anytime unlike the one-time configuration that was possible prior to this release.
 - ◆ If you change the archive location, the archives that were created on the previously configured location cannot be moved to the new location.
 - ◆ The Logger user interface for Event Archives has been updated to display relevant information for the archiving changes introduced in v5.0 SP1. For example, prior to v5.0 SP1, only the name and date of the archive was displayed. However, starting with v5.0 SP1, the name, date, and storage group name are displayed—the name of the storage group to which the archives pertains is displayed as a separate column on the Event Archives page, as shown in the following figure.

Archiving Events

To save events for a particular day, you need to add an Event Archive. The table in the Event Archives tab shows the current archives and their status. A one-time configuration for the archive storage location must be established on your Logger before you can add an event archive. To establish an archive storage location, see [“Archive Storage Settings” on page 231](#).

The screenshot shows the ArcSight Logger web interface. The top navigation bar includes links for Monitor, Analyze, Reports, Configuration, and System Admin. The Configuration tab is active, and the Event Archives sub-tab is selected. A sidebar on the left contains a tree view with links like Devices, Event Archives, Storage, and others. The main content area displays a table of event archives.

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 231](#) 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 Start and End dates in the format m/dd/yy, where m is month number, dd is the day of the month (with a leading zero if necessary), and yy is the two-digit year number.

When the Start and End dates are different, one archive file per storage group, for each specified day is created. For example, if you specify the following Start and End dates:

Start Date: 2/12/10

End Date: 2/13/10

And, you select two storage groups—Internal Event and Default. Then, four archive files will be created as a result of this archive operation—two files per storage group for the specified two days.

The Event Archives table (under the Event Archives tab) lists the archives by an alias in this format: <archive_name>[<yyyy-m-dd>][<storage_group_name>].

- 5 Select the names of storage groups that need to be included in the archive.
- 6 Click **Save** to start archiving events, or **Cancel** to quit.



You can cancel an in-progress archive operation at any time using the Cancel link that displays on top of the Event Archives page.

To delete an Event Archive

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Archives** in the left panel.
- 3 Click the checkboxes in the left-most column to select the event archives that you want to delete.
- 4 Click **Remove** from the top of the screen to delete the selected archives.
- 5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Event Archive.

Scheduled Event Archive

You can schedule a daily event archive and specify what hour of the day it should run. Scheduled event archives that have finished running appear on the archive list on the Event Archives tab. Only one scheduled event archive can run at a time; however, it can run in parallel with a manually scheduled archive.

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 307](#) 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 from the “Time For Daily Archive to Start” list. Scheduled archives must start on the hour, and midnight and 1:00 AM are not on the list to allow your Logger to receive all of the previous day’s events.
- 5 Select the storage groups whose events should be included in the scheduled archive.
- 6 Click **Save** to schedule daily event archive, or click on another tab or sub-menu to cancel.

Archive Storage Settings

On the **Logger appliance**, Event Archives are saved to a specific NFS or CIFS mount point, or SAN. For the **software Logger**, event archives are saved to the specified directory, which can be a path to a local directory or to a mount point on the machine on which the software Logger is installed. To establish a mount point, see your system’s operating system documentation.



Starting with v5.0 SP1, the archive locations specified on the Archive Storage Settings page can be changed anytime unlike the one-time configuration that was possible prior to this release.

To perform Archive Storage Setting setup

- 1 If you are using the Logger appliance, create the NFS or CIFS mount point. (See [“Storage” on page 314](#) and [“CIFS Settings” on page 314](#).)

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 Specify a mount location and an archive path for each storage group. You can specify a different path for each storage group, thus enabling the Logger to archive events to a different location for each storage group.

For Logger appliances, choose the name of an NFS mount, CIFS mount, or SAN mount point for the Mount Location field. This drop-down list contains the names you specified when creating the NFS, CIFS, or SAN mount points (System Admin > NFS/CIFS/SAN).

For software Loggers, the Mount Location field does not exist. You need to enter a complete path where the archive file will be written in the Archive Path field. This path

could be a local directory or a mount point that is already established on the machine on which the Logger software is installed.



You must configure settings for all storage groups on the following page even if you don't intend to archive all of them.

The above figure shows a screen from the Logger appliance. On a software Logger, the Mount Location field is not available.

- 6 Click **Save**.

Loading and Unloading Archives

Archived events must be loaded back on Logger before they can be included in a search operation. When an Event Archive is loaded, its events are included in searches, but the archive itself remains on the remote storage. When an Event Archive is unloaded, it is available for loading, but its events are not included in searches. You can unload a loaded archive if you no longer need to include it in your search operations.

To load or unload an Event Archive

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Archives** in the left panel.
- 3 Click the checkboxes in the left-most column to select the event archives that you want to load or unload.
- 4 Click **Load** or **Unload** from the top of the screen to load or unload the selected archives.

Storage

Different storage groups support the implementation of multiple retention policies. Each group can have a different policy, and storage rules determine which storage group is used for events from specific Device Groups. See [“Retention Policy” on page 27](#). 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 30](#)) 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.

Storage Groups can only be created during the Logger initialization phase, described in [“Initializing the Logger Appliance” on page 27](#). 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. If you are decreasing the size of the storage group and the new size is lesser than the currently used space on the storage group, you will need to delete data to achieve the new size. Logger UI guides you in this situation to delete sufficient data. See [“To edit \(including resizing\) a Storage Group” on page 234](#) to change the size of a Storage group.


Storage Groups Storage Rules Storage Volume					
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.

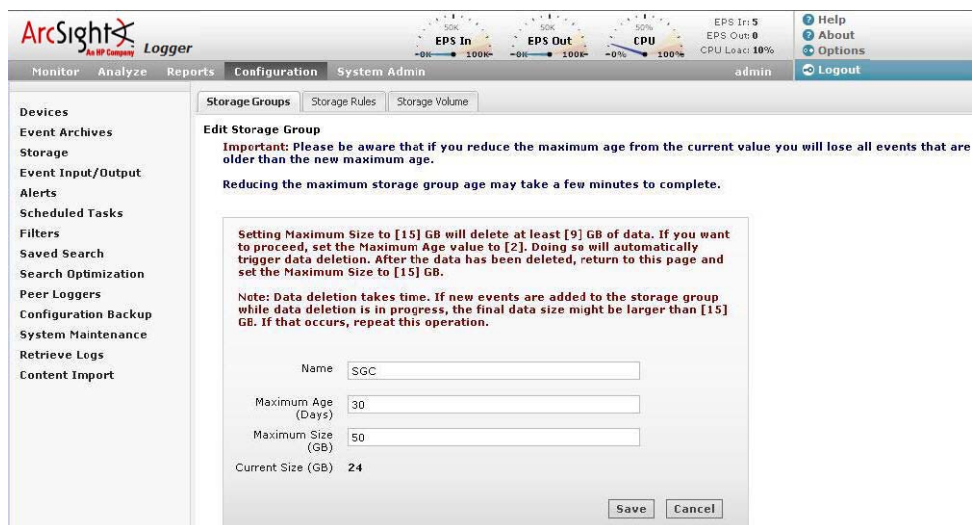
Once the Logger has been set up, the Storage Groups page, as shown in [Figure 6-3](#), does not allow adding or deleting Storage Groups.

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): 15

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 225](#).)

Logger allows you to create up to 40 storage rules. If you create additional rules, an error might be generated.

- 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	<p>Select a Storage Group from the drop-down list. The Storage Groups must already be set up before any Storage Rules are added.</p> <p>You can only add storage groups at the time of Logger initialization.</p>

Parameter	Description
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.

- 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

- Click the **Configuration > Settings** tab, click **Storage** in the sub-menu, and click the **Storage Rules** tab.
- Find the Storage Rule to be edited in the table.
- 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

- Click the **Configuration > Settings** tab, click **Storage** in the sub-menu, and click the **Storage Rules** tab.
- Find the Storage Rule to be deleted in the table.
- 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. Additionally, SAN is not supported for software Loggers.

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.



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 293](#).

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 314.) 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><install_directory>/data/logger</code> directory and cannot be changed.
Maximum Size	Enter the storage volume size, in GB. For the software version Logger , the maximum size is determined by the lower value of the following: <ul style="list-style-type: none"> Limit specified in the Logger license. (See “License Information” on page 295 for this information.) The storage volume partition size on the system on which Logger software is installed.

- 5 Click Save.

To increase the size of a storage volume:

See [“Storage Volume Size Increase” on page 293.](#)

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. These receivers can read single- or multi-line log files.

The SCP and SFTP protocols (for setting up File Transfer Receivers) are not FIPS compliant. For a FIPS-compliant configuration, configure the File Transfer receivers to use FTP.

- **File Receiver.** File Receiver-type receivers read log files from a network file system (NFS), CIFS, or Storage Area Network (SAN). These receivers can read single- or multi-line log files.
- **SmartMessage Receiver.** SmartMessage receivers listen for encrypted messages from ArcSight SmartConnectors.

About Multi-line File Receivers

A multi-line file receiver is a type of file or file transfer receiver that can read multi-line log files. This receiver is useful in reading events that span more than one line, such as exceptions in a server log. To create this type of receiver, create a file or file transfer receiver and specify a regular expression (in the Event Start field) to determine the start of a new event in the log file—each new event starts at the point where the regular expression is matched to the characters in the log file.

A multi-line file receiver is useful in finding exceptions in a server log. In such logs, the entire stack trace can be treated as one event instead of each line as a single event.

For example, in the following log file, each event starts with a timestamp embedded within square brackets ([]); therefore, you can use this regular expression to identify each event:

```
^\[\d+-\d+-\d+ \d+:\d+,\d+].*

[2010-12-06 13:11:26,824][INFO ][I18N]Locale has not been chosen by the user.
[2010-12-06 13:11:26,828][ERROR][DirectConnection$ReadChannel]
java.io.IOException: end of communication channel
    at com.arcsight.logger.distributed.DirectConnection.a(DirectConnection.java:39)
    at com.arcsight.logger.distributed.DirectConnection.access$200(DirectConnection.java:19)
    at com.arcsight.logger.distributed.DirectConnection$ReadChannel.run(DirectConnection.java:85)
    at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886)
    at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908)
    at java.lang.Thread.run(Thread.java:619)
```

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 241](#) for more information.
- 2 Add type-specific parameters. Receiver parameters are documented in [Table 6-1, “Receiver Parameters,” on page 242](#).
- 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.



Tip

Wait a few minutes after enabling a receiver before disabling it. Likewise, wait before enabling a receiver that has just been disabled. Background tasks initiated by enabling or disabling a receiver can produce unexpected results if they are interrupted.

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



Note

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 314](#).
- 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 54](#) and specify:

- Logger IP or hostname
- Port 443 (port must be 443) for Logger appliance; for software Logger, port you configured during its installation
- 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.

If you are creating a multi-line receiver, select File Transfer or File Receiver, depending on the protocol you want to use for reading the log file.

- 5 Click **Next** to edit receiver parameters listed in [Table 6-1 on page 242](#).


If you are creating a multi-line receiver, enter a regular expression in the Event Start field to determine the start of a new event in the log file, as shown in the following figure.

```
[2010-12-06 13:11:26,824][INFO ][I18N]Locale has not been chosen by the user.
[2010-12-06 13:11:26,828][ERROR][DirectConnection$ReadChannel]
java.io.IOException: end of communication channel
    at com.arcsight.logger.distributed.DirectConnection.a(DirectConnection.java:39)
    at com.arcsight.logger.distributed.DirectConnection.access$200(DirectConnection.java:19)
    at com.arcsight.logger.distributed.DirectConnection$ReadChannel.run(DirectConnection.java:85)
    at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886)
    at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908)
    at java.lang.Thread.run(Thread.java:619)
```

- 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.
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 304.</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	<p>For the Logger appliance:</p> <p>The default port is 514. (For SmartMessage receivers, configure the SmartConnector for port 443.)</p> <p>For the software version of Logger:</p> <p>If you installed software Logger as a root user, you can use any available port.</p> <p>If you installed software Logger as a non-root user, you can only use a port greater than 1024. Ports lesser than 1024 cannot be configured. Default 8514.</p>
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."

Parameter	Receiver Types	Description
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: <code>/tmp/SyslogData/syslog.log.gz, /security/logs/*/, /security/log?/admin/special/</code></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 307 before you schedule a file transfer.</p>
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 314. To mount CIFS shares, see "CIFS Settings" on page 314.</p> <p>For more information about SAN, see "SAN" on page 318.</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>

Parameter	Receiver Types	Description
Source Type	File Receiver, File Transfer	Select from the pulldown list of log file types, including: Apache HTTP Server Access Apache HTTP Server Error Juniper Steel-Belted Radius Microsoft DHCP Log IBM DB2 Audit
Wildcard	File Receiver	Regular expression describing the log files to read. Note: This is a regular expression, not a typical file wildcard like <code>"*. *"</code> 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> . The default is <code>.*</code> , meaning all files.
Mode	File Receiver	Mode is one of: 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	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. The default is 10 seconds. 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.
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 zone	File Receiver or File 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) .

Parameter	Receiver Types	Description
Date/time location regular expression	File Receiver or File Transfer	<p>A regular expression describing which characters represent the timestamp in the log file. For example:</p> <pre>.*\[(.*)\].*</pre> <p>This regular expression specifies that the timestamp is found inside the first set of square brackets on each line. The first capturing group (the part of the regex in parentheses) is that part that is then parsed using the Date/time format.</p> <p>The default is "" (no timestamp in log file).</p>
Date/time format	File Receiver or File Transfer	<p>Required if the log file contains timestamps in the same format for each event. If not specified (or if the Date/time location regex is blank), each event in the file will be stamped with the date that the file itself was first seen by Logger (not its file system timestamp).</p> <p>See Step Table 6-3 for a list of format specifiers.</p> <p>The default is "" (no timestamp in log file).</p>
Event Start (regex)	File Receiver or File Transfer	<p>A regular expression that specifies the start of a new event in a log file. Specify this expression to enable the receiver to read multi-line log files. Each new event starts at the point where the regular expression is matched to the characters in the log file. For example,</p> <pre>^\[d+-\d+-\d+ \d+:\d+,\d+\].*</pre> <p>This regular expression matches timestamps such as:</p> <pre>[2010-12-06 13:09:46,818]</pre> <p>When this field is left blank (""), each line in the log file is treated as a single event.</p> <p>The default is "" (each line in the log file is a single event).</p>

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 246](#).

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 245](#) 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

Source	Date and Time Pattern
0:08 PM, PDT	K:mm a, z
02001.July.04 AD 12:08 PM	yyyyy.MMMMM.dd GGG hh:mm aaa
Wed, 4 Jul 2001 12:08:56 -0700	EEE, d MMM yyyy HH:mm:ss Z
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: The number is limited by system resources—memory, CPU, disk input/output.



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 248. 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. You can only specify regular expression queries for forwarders. See “Searching for Events on Logger” on page 110. 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 270.) 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 \$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.</p>

Parameter	Forwarder Types	Description
Source Type	Connector	<p>Select from the pulldown list of log file types, including:</p> <ul style="list-style-type: none"> Apache HTTP Server Access Apache HTTP Server Error Juniper Steel-Belted Radius Microsoft DHCP Log IBM DB2 Audit <p>Note: Source Type must be the same in Receiver, Forwarder, and SmartConnector. See “Forwarding Log File Events to ESM” on page 254.</p>
Preserve Syslog Timestamp	UDP, TCP, ESM	<p>Set to true to preserve the syslog timestamp. The default is true--the timestamp is the original receipt time of the event.</p> <p>If set to false, original timestamp is replaced with Logger's receipt time.</p>
Preserve Original Syslog Sender	UDP, TCP, ESM	<p>Set to true to preserve the original sender. The default is true--the sender is the original sender.</p> <p>If set to false, the original sender information is replaced with Logger's information.</p>
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 251 .)

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 248](#), 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 455](#)) that are not normalized or categorized.

Maximum number of ESM destinations that can be configured: As many allowable on the SmartConnectors you are using.



Do not use basic aggregation for Logger's built-in SmartConnector because it is resource intensive. (Basic aggregation is set using the Enable Aggregation (in seconds) field from the ESM Console.) Instead, follow these steps on the ESM Console to configure field-based aggregation:

- 1 Ensure that Processor > Enable Aggregation (in seconds) is set to "Disabled" (to disable basic aggregation).
- 2 Right-click the connector and select **inspect/edit/**.

For additional details about configuring field-based aggregation, refer to the *ArcSight SmartConnector User's Guide*.

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 253](#).

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 252](#).
- 3 Create an ESM Forwarder that refers to this ESM Destination. (See ["Forwarders" on page 246](#)).

Add 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	••••••••

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 253](#) 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.



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
Name	The name for this ESM Destination.
Connector Name	The SmartConnector name.
Connector Location	The physical location of the SmartConnector machine. If you do not want to specify a location, enter “None.”
Logger Location	The physical location of the Logger. If you do not want to specify a location, enter “None.”
IP or Host	The ESM Manager to which the Forwarder will direct events. Note: Make sure the name or IP address you specify in this field is exactly the name or IP address configured on the ESM Manager. If the two names or IP addresses do not match, you will not be able to set up an ESM destination successfully.
Port	Typically 8443.
User Name	The name of an existing User of the 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 IP or Host field exactly matches the name configured on the ESM Manager.

There was a problem: Failed to add destination

Additionally, if the ESM Manager is configured using a host name instead of IP address, make sure you add the ESM Manager host name and IP address in the Logger’s hosts file (System Admin > Network > Hosts).


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 for forwarding events and alerts.

If your Manager does not have FIPS 140-2 mode enabled, you can obtain a certificate file for your Manager in these ways:

- From the Manager's keystore
- From the 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.



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.
- 4 In the **Certificates** tab, click **Add** to display the following screen.

- 5 Enter an alias for the certificate file. This name is used to easily identify a certificate file. For example, `arcsight_esm_manager1_cert`.
- 6 Click **Browse** to locate the Manager certificate file you copied.
- 7 Check the "Overwrite Certificate" box if you want this certificate to overwrite an existing certificate with the same alias.
- 8 Click **Save**.

Forwarding Log File Events to ESM

Logger can read events from a log file and forward those events to 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 242](#)).

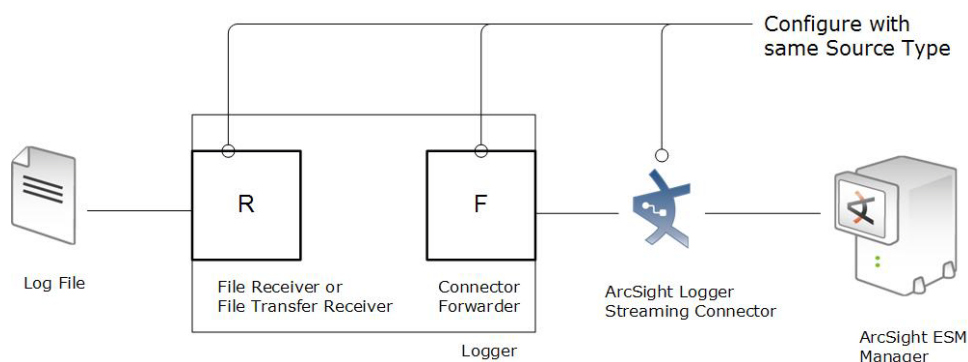


Figure 6-9 Configure the Receiver, Forwarder, and SmartConnector with the same Source Type to use Logger to forward log file events to ArcSight ESM.

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.	Any number of alerts can be defined. All defined alerts are enabled and effective, however, a maximum of 50 alerts can run concurrently.
A maximum of five alerts can be enabled at any time.	
No limit on the number of configured e-mail destinations; however, you can only set one SNMP, one Syslog, and one ESM destination.	No limit on the number of configured e-mail destinations; however, you can only set one SNMP, one Syslog, and one ESM destination.
Only regular expression queries can be specified for these alerts.	Queries for these alerts are defined using the flow-based search language that allows you to specify multiple search commands in a pipeline format, including regular expressions. Aggregation operators such as chart and top cannot be included in the search query.
Alerts are triggered in real time. That is, when specified number of matches occur within the specified threshold, an alert is immediately triggered.	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 .

Real Time Alerts

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

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. Audit events for alerts are only written to the Internal Storage group and not forwarded to ESM by default. If you need to forward the audit events generated for alerts to ESM, please contact ArcSight Customer Support for assistance.



To maintain an optimal size of an alert event, the event does not contain event IDs of all the triggering events if a **match count of 101 or higher is specified**. As a result, the `baseEventCount` field in the event does not reflect the true number of matching events for such alert events. Triggering events are truncated in multiples of 100. Therefore, if you specify a match count of 101, only one event is included in the alert event and the `baseEventCount` field value is 1. Similarly, if you specify a match count of 720, only 20 events are included and the `baseEventCount` field value is 20.

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]

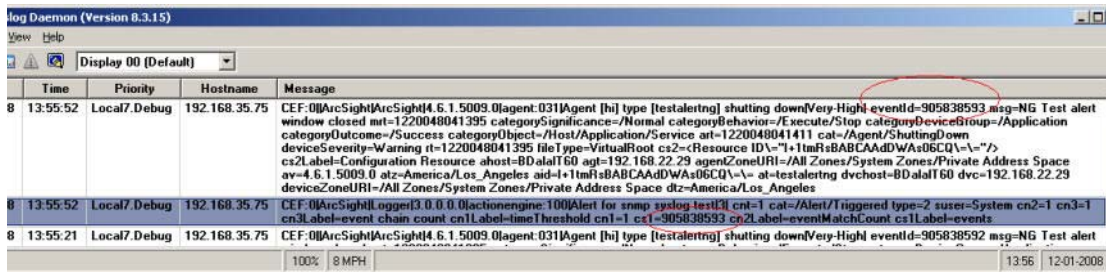
Event Receipt Time [Tue Nov 11 16:46:50 PST 2008]

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
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 mrt=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=111mRtBABCADWAs06CQ\> cs2Label=Configuration Resource ahost=BDalaf60 agt=192.168.22.29 agentZoneURI=/All Zones/System Zones/Private Address Space av=4.6.1.5009.0 atz=America/Los_Angeles aid=111mRtBABCADWAs06CQ\> at=testalertng dvchost=BDalaf60 dvc=192.168.22.29 deviceZoneURI=/All Zones/System Zones/Private Address Space dtz=America/Los_Angeles
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 cn1-1 cat=/Alert/Triggered type=2 suser=System cn2-1 cn3-1 cn3Label=event chain count cn1Label=timeThreshold cn1=1 cs1=905838593 cn2Label=eventMatchCount cs1Label=events
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.
- Logger supports SNMP v1.0.
- 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 [“Static Routes” on page 308](#). See also [“SNMP Destinations” on page 266](#), [“Syslog Destinations” on page 267](#), and [“ESM Destinations” on page 251](#).

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 259](#).

- 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 259](#) 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 71](#).



Tip

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.



Note

To maintain an optimal size of an alert event, the event does not contain event IDs of all the triggering events if a **match count of 101 or higher is specified**. As a result, the `baseEventCount` field in the event does not reflect the true number of matching events for such alert events. Triggering events are truncated in multiples of 100. Therefore, if you specify a match count of 101, only one event is included in the alert event and the `baseEventCount` field value is 1. Similarly, if you specify a match count of 720, only 20 events are included and the `baseEventCount` field value is 20.

- 6 Enter notification destinations. Enter any combination of:
 - ◆ One or more e-mail addresses, separated by commas
 - ◆ An SNMP Destination
 - ◆ A Syslog Destination
 - ◆ An ESM Manager
- 7 Click **Save**.

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 'Alert' sub-tab is selected. The 'Add Alert' dialog box is open, displaying 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 currently set to 'NONE'.
- Syslog Destination:** A dropdown menu currently set to 'NONE'.
- ESM Destination:** A dropdown menu currently set to 'NONE'.
- 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.




Note

A maximum of five alerts can be enabled at one time. To enable an additional alert, you will need to disable a currently enabled alert.

To Edit a Real Time Alert

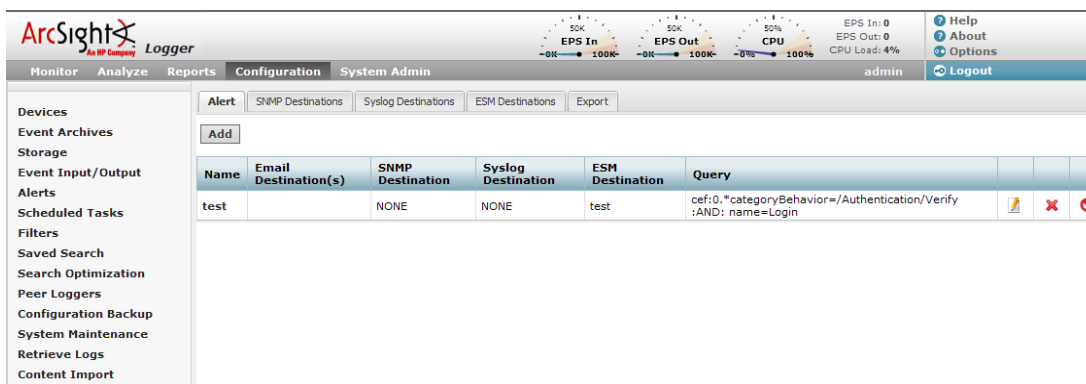
- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Locate the Alert that you want to edit.
- 4 Click the Edit icon (). A screen similar to that shown in [Figure 6-10 on page 259](#) appears. Remember that only alphanumeric characters can be used in an 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 273](#). 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).

**Note**

To ensure system performance, a maximum of 200 alerts are allowed per saved search alert job. Therefore, if a saved search alert job triggers more than 200 alerts, only the first 200 alerts are sent out for that job iteration; the rest are not sent. Additionally, the job is aborted so it does not trigger more alerts for that iteration and the status for that job is marked “Failed” in the Finished Tasks tab (Configuration > Scheduled Tasks > Finished Tasks). The job runs as scheduled at the next scheduled interval and alerts are sent out until the maximum limit is reached.


This limit does not exist on the real-time alerts.

Creating a Saved Search Alert

You can create a Saved Search Alert in two ways:

- From the search results page (Analyze > Search)
- From the Scheduled Searches page (Configuration > 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 110](#).
- 2 Click the Save icon () and enter the following settings.

Parameter	Description
Name	A name for the query you are saving.
Save as	Whether to save the query as a filter or as a Saved Search. To save the query as a Saved Search Alert, select “Saved search”.
Schedule it	Whether to schedule the alert right now or later. Click to schedule now, or leave blank to schedule later.
Schedule type	Whether the query is being saved as a scheduled search or as a scheduled alert. Scheduled searches run on a predetermined schedule and export results to a 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 273 .
----------------	--

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 273.</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 266](#).


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 263](#).
- 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
Name	The name for the destination.
Type	This is the value you entered in the previous screen. This value cannot be changed.
Ip/Host	Host name or IP address.
Port	Port (default is 514).
Connection Retry Timeout	(Only for TCP Syslog Destinations) The time, in seconds, to wait before retrying a connection. The default is 5 seconds.

- 7 Click **Save** to create the new Syslog Destination

To Edit a Syslog Destination

- 1 Click **Configuration** 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 253](#).

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 252](#).

Export

See [“Exporting and Importing Content” on page 297](#).

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 307](#) 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 275](#))
- File Receivers and File Transfer Receivers (See [“Receivers” on page 239](#))
- Event Archive (See [“Archiving Events” on page 229](#))
- Configuration Backup (See [“Configuration Backup and Restore” on page 284](#))

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

A shared filter is visible to all Logger users. Once created, any Logger user can use it to search for events. The query you specify for a shared filter can be a Unified query (that uses keywords, indexed, and non-indexed fields) or a Regex query (that specifies a regular expression). Creating Regex Query shared filters are useful for creating alerts and forwarders, which accept only regex queries.

■ Search Group

Search group filters provide an access control mechanism to limit the events that users in a particular user group can see. Search Group filters can also be used to limit the events processed by a category of reports (see [“Using Report Category Filters” on page 222](#)). Only users with administrative privileges can create these filters.

A set of pre-defined filters, also known as system filters, exist on your Logger as well. For more information about system filters, see [“System Filters/Predefined Filters” on page 124](#).

To create a filter

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

- 2 Click **Filters** in the left panel.
- 3 Click the **Filters** tab in the right panel to create a shared filter. OR click the **Search Group Filters** tab to create a search group filter. (See [“Filters” on page 270](#) for information about shared and search group filters.)
- 4 Click **Add** to display the following page.
- 5 Enter a name for the new filter in the Name field.
Filter names are case-sensitive.
- 6 If you are creating a shared filter, select **Unified** or **Regex Query**.
For Search Group filters, select **Search Group**. Note that non-administrator users cannot create Search Group filters.
- 7 Click **Next**.
- 8 If you selected Unified or Regex Query method in the previous step, enter the query for the new filter.
For Unified queries:

When you type a query, Logger’s Search Helper enables you to quickly build a query expression by automatically providing suggestions, possible matches, and applicable operators. See [“Search Helper” on page 64](#) for more information.

OR


Click **Advanced Search** to use the Search Builder Tool to create the query. For details about using the Search Builder Tool, see [“Using the Search Builder Tool” on page 102](#).
For Regex queries:

Enter the regular expression in the Query text box.
- 9 Click **Save**.




If you created a Search Group filter, make sure that you associate it to a user group, as described in [“Search Group Filters” on page 272](#).

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 222](#).

- **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.

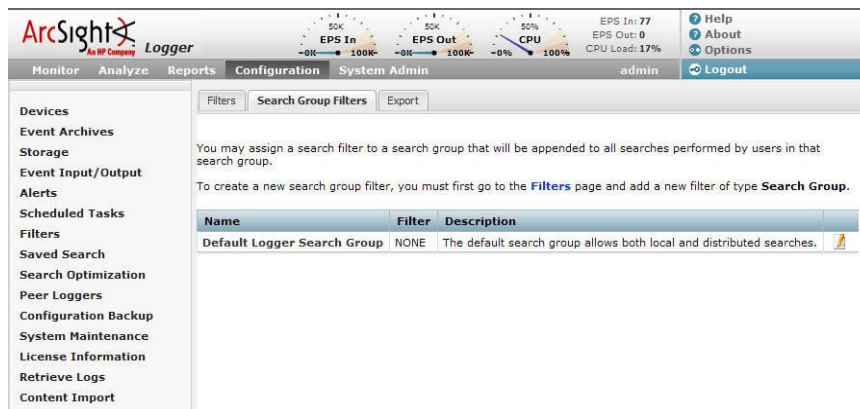


Figure 6-13 Search Group Filters Page




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 270](#). To create, edit, or delete User Groups, see [“Users/Groups” on page 333](#).



Only users that are members of a System Admin group can assign Search Group Filters. For more information, see [“Users/Groups” on page 333](#).

To associate a Search Group Filter with a User Group

- 1 If the User Group that you want to associate with the Search Group Filter does not exist, create a new User Group of type Search Group. For instructions, see [“Users/Groups” on page 333](#).
- 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 270](#). 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 297](#).

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 255](#).

You can schedule a saved search to run at a specific interval. For more information, see [“Scheduled Saved Search” on page 275](#).

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 307](#) 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**.

- 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. When you type a query, Logger's Search Helper enables you to quickly build a query expression by automatically providing suggestions, possible matches, and applicable operators. See "Search Helper" on page 64 for more information.
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.

- Click **Save** to add the new Saved Search, or **Cancel** to quit.

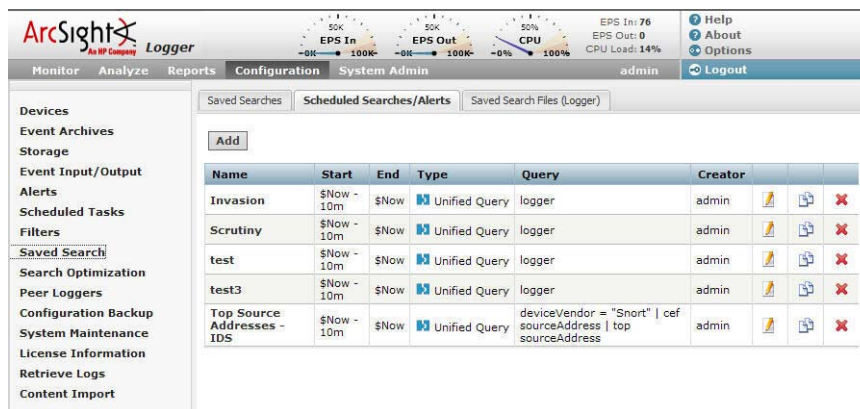


Figure 6-14 Saved Search page

To edit a Saved Search

- Click the **Configuration > Saved Search**.
- Find the Saved Search to be edited in the table.
- Click the Edit icon (). Change the information in the form and click **Save**.

To delete a Saved Search

- Click the **Configuration > Saved Search**.
- Find the Saved Search to be deleted in the table.
- 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 261](#).

The results of a scheduled Saved Search are written to a file, as described in [“Saved Search Files” on page 278](#).

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.

The screenshot shows the 'Add Scheduled Saved Search' configuration page. It includes the following elements:

- Name:** A text input field.
- Schedule:** A dropdown menu set to 'Everyday'.
- Hour of day:** A dropdown menu and a text input field for 'Hours (24 hour format)'.
- Saved Searches:** A list box containing 'Invasion', 'Scrutiny', 'test', 'test3', and 'Top Source Addresses - IDS'.
- Job type:** A dropdown menu set to 'Search'.
- Search Result Export Options:**
 - Export Options:** Radio buttons for 'Export to remote location' (selected) and 'Save to Logger'.
 - File format:** A dropdown menu set to 'PDF'.
 - Export directory name:** A text input field.
 - Title:** A text input field.
 - Fields:** A text area showing 'Event Time, Receipt Time, Device, Logger, Name, Version, Device Vendor, Device Product, Device Version, Signature ID, Severity'.
 - Chart type:** A dropdown menu set to 'Column'.
 - Chart result limit:** A text input field set to '10'.
 - Include summary:** A checkbox.
 - Include only CEF events:** A checkbox.
- Buttons:** 'Save' and 'Cancel' buttons at the bottom right.

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 273. When <i>multiple</i> Saved Searches are specified in one scheduled saved search job, the resulting file contains the number of hits for each Saved Search and not the actual events.</p>
Job Type	Select Search for a scheduled Saved Search.
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 278) 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>
File Format	<p>CSV, for comma-separated values file.</p> <p>PDF, for a report-style file that contains search results as charts and in tables. You must specify a title for the report in the Title field. If the search query contains an operator that creates charts such as chart, top, and so on, charts are included in the PDF file. In that case, you can also set the Chart Type and Chart Result Limit fields. These fields are described later in this table.</p>

Parameter	Description
Export Directory Name	<p>For the Logger appliance, select the directory where the search results will be exported from the pulldown menu.</p> <p>For the software version of Logger, enter the directory path in this field, which can be a path to a local directory or to a mount point on the machine on which the software version of Logger is installed.</p> <p>If a directory of the specified name does not exist, it is created. If a directory of the specified name exists and the Overwrite box is not checked, an error is generated. If the Overwrite box is checked, the existing directory contents are overwritten.</p>
Title	(Optional) A meaningful name that appears on top of the PDF file. If no title is specified, "Untitled" is included.
Fields	<p>A list of event fields that will be included in the exported file.</p> <p>By default, all fields are included.</p> <p>You can enter fields or edit the displayed fields by deselecting All Fields.</p>
Chart Type (for PDF only)	<p>(If the search query includes an operator that creates a chart, this field is meaningful; otherwise, it is ignored.)</p> <p>Type of chart to include in the PDF file. You can select from: Column, Bar, Pie, Area, Line, Stacked Column, Stacked Bar.</p> <p>Note: If the Chart Type is different from the chart displayed on the Search Results screen, the value selected for this option overrides the one shown in the screen. Therefore, the exported PDF contains the chart you specify for this option and not the one shown on the screen.</p>
Chart Result Limit (for PDF only)	<p>(If the search query includes an operator that creates a chart, this field is meaningful; otherwise, it is ignored.)</p> <p>Number of unique values to plot. Default: 10</p> <p>If the configured Chart Result Limit is less than the number of unique values for a query, the top values equal to the Chart Result Limit are plotted. That is, if the Chart Result Limit is 5 and 7 unique values are found, the top 5 values will be plotted.</p>
Include Summary	Check this box to include an event count with the Saved Search, or a total when more than one Saved Search is specified.
Include Non-CEF Events	Check this box to include all events. Uncheck the box to include only CEF (see "Common Event Format" on page 455) events in the output.

- 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.



Name	Last Modified	Size	State	Error Message	Retrieve
job_local_2007-06-22 17-00-00.csv	Fri Jun 22 17:00:02 PDT 2007	202 bytes	Exported		✖
job_local_2007-06-24 13-00-04.csv	Sun Jun 24 13:06:49 PDT 2007	202 bytes	Exported		✖
job_local_2007-06-22 07-00-00.csv	Fri Jun 22 07:00:02 PDT 2007	202 bytes	Exported		✖
job_local_2007-06-25 01-00-00.csv	Mon Jun 25 01:17:10 PDT 2007	202 bytes	Exported		✖
job_local_2007-06-24 11-00-00.csv	Sun Jun 24 11:12:35 PDT 2007	202 bytes	Exported		✖
job_local_2007-06-22 02-00-00.csv	Fri Jun 22 02:00:02 PDT 2007	202 bytes	Exported		✖
job_local_2007-06-24 23-00-01.csv	Sun Jun 24 23:17:38 PDT 2007	202 bytes	Exported		✖

Figure 6-16 Saved Search Files page

Search Optimization

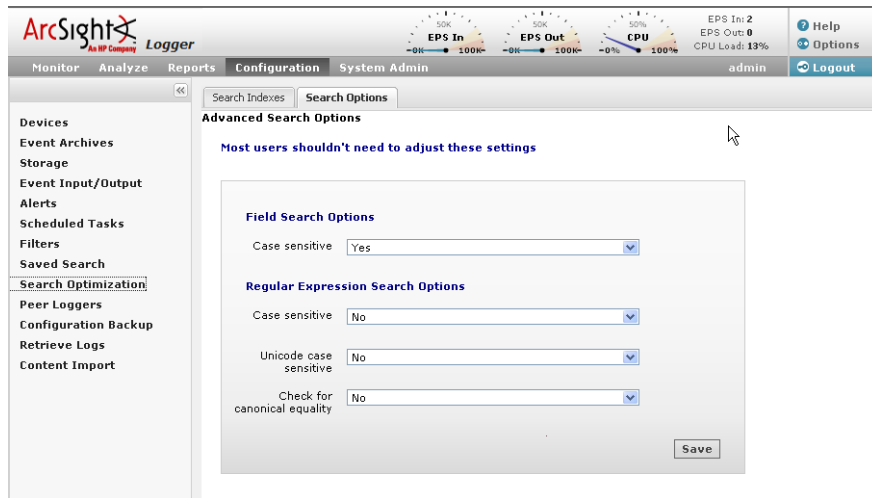
The search optimization option enable you to:

- Add search indexes for field query search operations
- Tune advanced search options
- Delete custom field sets

Add Search Indexes

See ["Indexing" on page 119](#) 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



Note

- You need to have the “Edit, save, and remove fieldsets” privilege to delete a custom field set.
- You can only delete the field sets you create, and not the predefined ones available on Logger.

To delete a custom field set:

- 1 Click **Configuration** 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:



Note

On a Logger using local or RADIUS authentication, you can **use either authentication method**, although peer authorization ID and code are recommended for reasons described below. However, if you are using SSL Client Authentication (CAC) on your Logger, **authorization ID and code is the only way to authenticate a peer**. You cannot use a user name and password.

FIPS-enabled Loggers are not limited to a specific authentication method. Therefore, you can use any listed below.

■ User name and password

A user name configured on the Logger is used for authentication

■ Peer Authorization ID and Code

Authorization ID and Code generated on a remote Logger are used by the initiator Logger to peer with it. The generated ID and Code are specific to the initiator Logger because the IP address of the initiator is used to generate the ID and code, and can be used only for peering from the initiator. Therefore, this method is more secure than using user name and password.

Guidelines

You should be aware of these guidelines when peering Loggers:

- You can peer a Logger to one or more remote Loggers.
- Peer Loggers can run different versions. A Logger appliance can peer with a software version of Logger. However, a **v5.x software** Logger cannot peer with a Logger (appliance or software) running v4.x. These are the supported paths for peering Loggers:
 - ◆ v5.x software Logger to v5.x software
 - ◆ v5.x software Logger to v5.x appliance
 - ◆ v4.5 software Logger to v5.x appliance
 - ◆ v4.5 appliance to v5.x appliance
 - ◆ v5.x appliance to v5.x appliance
- The time and date on the system on which the software Logger is installed must be set correctly with respect to its timezone to peer with other Loggers. ArcSight recommends that you configure the Logger system to synchronize its time with an NTP server regularly.
- If you need to run a search using the pipeline operators across peers, make sure the peers are running Logger v5.1 or later. For information about pipeline operators, see [“Search Operators” on page 75](#).
- 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.

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 and the Logger Rights Group with “View registered peers” privilege set to Yes. See [“Searching Peer Loggers \(Distributed Search\)” on page 111](#).
- 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/Groups" on page 333](#).
- 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 283](#).
- 5 Add Logger B's information, as described in ["To add a peer Logger" on page 282](#):
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 two tabs: 'Peer Loggers' and 'Peer Authorizations'. The 'Peer Loggers' tab is active, showing fields for 'Peer Host Name', 'Peer Port' (set to 443), 'Peer User Name', and 'Peer Password'. There are two radio buttons: 'Peer Login Credentials' (selected) and 'Peer Authorization Credentials'. Below these, a note states: 'Following fields are for local (currently connected) logger and are optional. This needs to be changed only seldomly.' This is followed by 'External IP Address' (192.168.36.42) and 'Local Port' (443). 'Save' and 'Cancel' buttons are at the bottom right.

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 for Logger appliance and pre-v5.0 software Loggers. Port you configured when installing software Logger. See "Guidelines" on page 281 .
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 326 .) 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 283 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 283 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 Logs Global Settings User and Group Information All Configuration Settings Existing Filters and Saved Searches Logger Monitor settings The following Reports content: <ul style="list-style-type: none"> • Queries, Reports, Parameters, Parameter Value Groups, Dashboard • Templates 	The following Report content only: <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 'Edit Configuration Backup' dialog in the ArcSight Logger interface. The left sidebar lists various configuration options, with 'Configuration Backup' selected. The main panel contains the following fields and settings:

- Protocol: SCP (dropdown)
- Port: 22 (text input)
- Ip/Host: (empty text input)
- User: (empty text input)
- Password: (empty text input)
- Remote directory: (empty text input)
- Backup content: All (dropdown)
- Schedule: ☒ One time only

Buttons for 'Save' and 'Cancel' are located at the bottom right of the dialog.

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 307.</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 285.

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

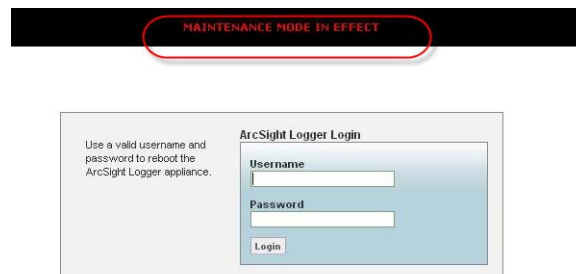
Logger users who will be performing operations that require it to be in maintenance mode must have the “Enable Maintenance Mode” privilege checked (System Admin > User Management > Groups tab > System Admin Group).



When a Logger is in maintenance mode, users with the “Enable Maintenance Mode” privilege can login but see this UI message:



All other users cannot login. The login screen displays this message:



Copyright © 2009 ArcSight Inc. All rights reserved.
ArcSight and the ArcSight logo are trademarks of ArcSight, Inc.

Entering Maintenance Mode

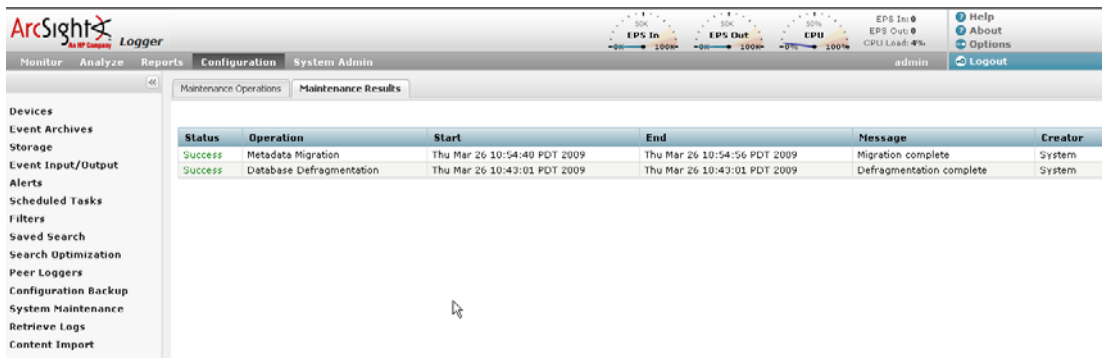
You cannot place a Logger in maintenance mode directly. A Logger can enter maintenance mode only when you perform an operation that requires it to be in that mode. For example, when defragmenting database, the user interface prompts you to enter Logger in maintenance mode, as illustrated in “Database Defragmentation” on page 288.

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**.



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:

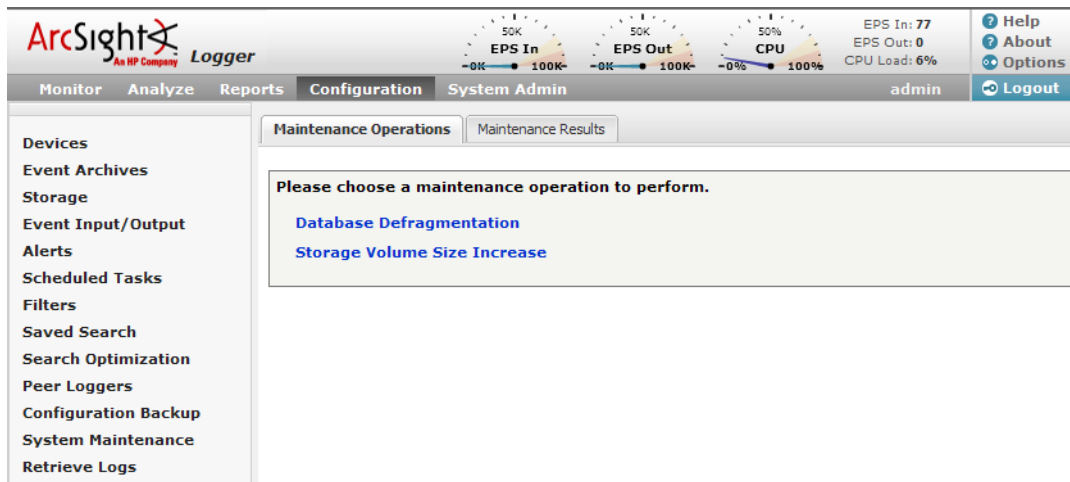
- Slow search and reporting
For example, even a search operation over the last two minutes of data is slow.
- Long pauses in the receiver and forwarder operations

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 287](#).
- 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 appliance** and restart the process from the beginning. **For the software Logger**, restart the Logger process as described in [“System - Process Status” on page 346](#).

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 appliance** and restart the process from the beginning. **For the software Logger**, restart the Logger process as described in [“System - Process Status” on page 346](#).

- 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 287](#).



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.



Note

On the software Logger, the following Database Defragmentation screens instruct you to click **restart** to resume normal operation when Logger is in maintenance mode. When you click restart, only the Logger service and its related processes are started on the machine on which the software Logger is installed.

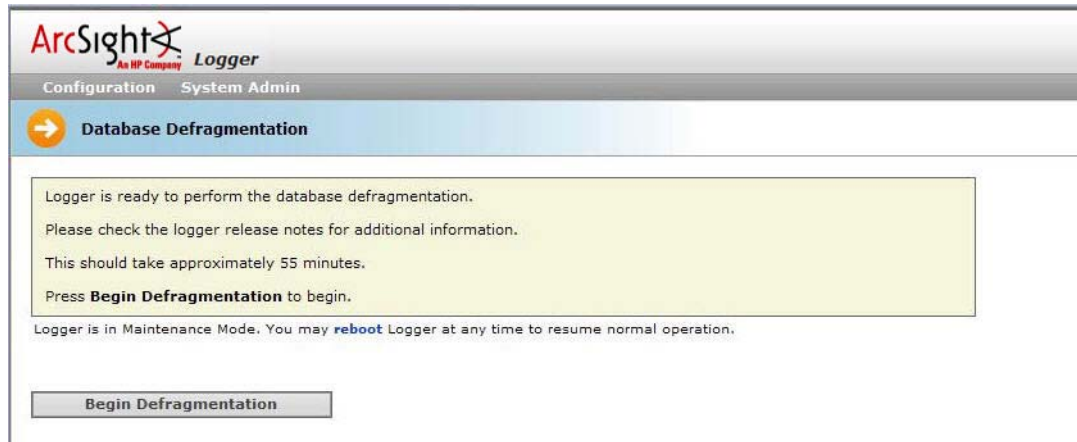



Figure 6-17 Begin Database Defragmentation

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.

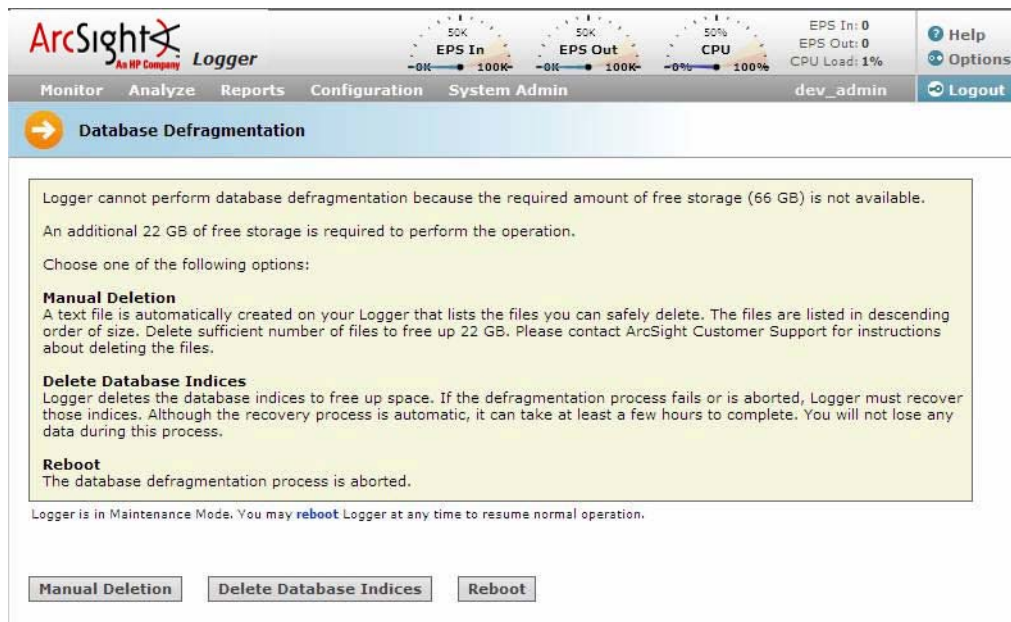


Figure 6-18 Required storage for Database Defragmentation is not available

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 291](#). 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 290](#) is displayed. Click **Begin Defragmentation** to proceed further.

If sufficient storage is still not found, the screen in [Figure 6-18 on page 291](#) is displayed. Choose from the listed options to create additional space. See [“You can choose from one of the following options:” on page 291](#) for more information.



If you need to exit the defragmentation process without creating sufficient storage, click **Reboot**.

Note

◆ **Delete Database Indices**

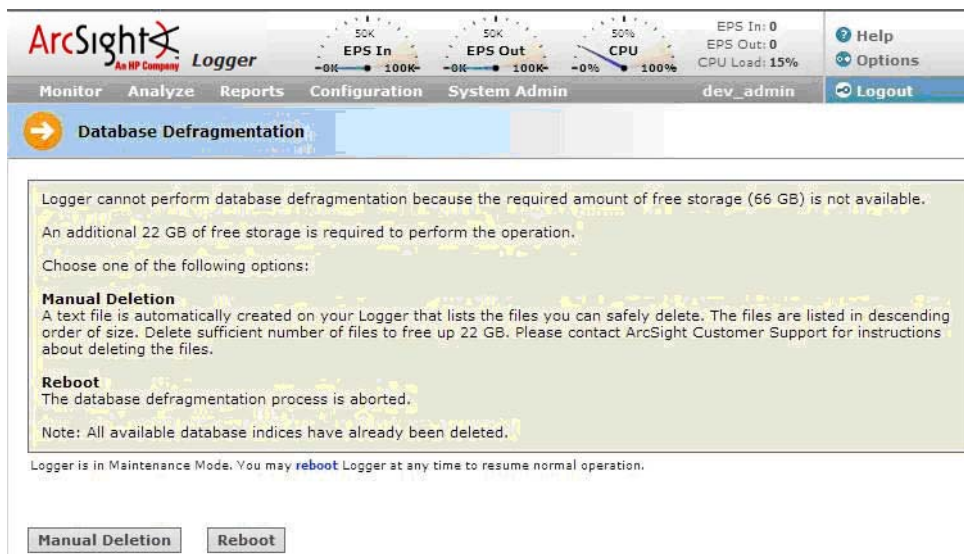
Logger automatically deletes a sufficient number of database indices, starting with the largest index, to free up the required amount of storage. If sufficient space becomes available after deleting database indices, defragmentation proceeds further automatically.

However, if sufficient storage is not available even after dropping database indices, the following screen is displayed.



The Manual Deletion option (shown in the following figure) is not displayed on L7100 Loggers as it is not applicable to those platforms.

Note



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 289](#).



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.



- For the “Storage Volume Size Increase” operation to show as an option under the System Maintenance operations (Configuration > System Maintenance), you need to belong to the System Admin group (with “Enable Maintenance Mode” privilege enabled) and the Logger Rights group.
- You can perform this process only if you have the “Enable Maintenance Mode” privilege enabled (System Admin > User/Groups > Manage Groups > System Admin Group).



About Increasing Storage Volume Size on a SAN Logger

Logger cannot detect a resized LUN. Therefore, if you change the LUN size after it has been mounted on a Logger, the new size is not recognized by Logger. As a result, you can only increase the size of a storage volume to the LUN size that was initially mounted on the Logger. Currently, Logger supports up to a 5.4 TB LUN.

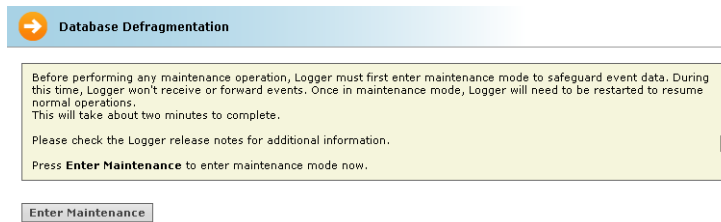
The following examples illustrate storage volume increase on a SAN Logger.

Initial LUN Size	LUN Resized	Current Storage Volume Size	Storage Volume Size Increase Allowed
4 TB	No	1 TB	Yes
4 TB	No	4 TB	No
4 TB	5 TB	1 TB	Yes, only up to 4 TB
2 TB	4 TB	1 TB	Yes, only up to 2 TB

To increase the size of a storage volume:

- 1 Click **Configuration > 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 287](#).

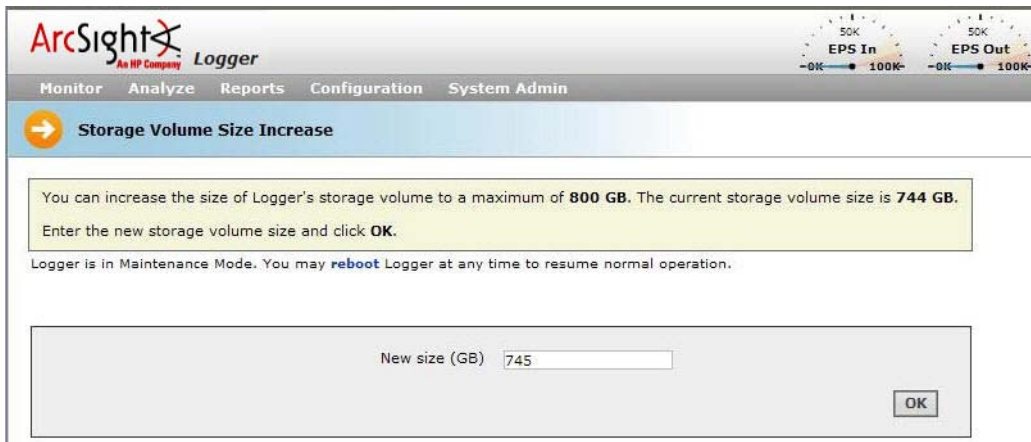


- 4 While entering the maintenance mode, Logger performs a check to determine if the storage volume size can be increased and by what amount.

If the storage volume can be increased, a message similar to the following is displayed. Enter the new size and click **OK**.



On the software Logger, the following Storage Volume Size Increase screens instruct you to click **restart** to resume normal operation when Logger is in maintenance mode. When you click restart, only the Logger service and its related processes are started on the machine on which the software Logger is installed.



ArcSight An HP Company **Logger**

Monitor Analyze Reports Configuration System Admin

Storage Volume Size Increase

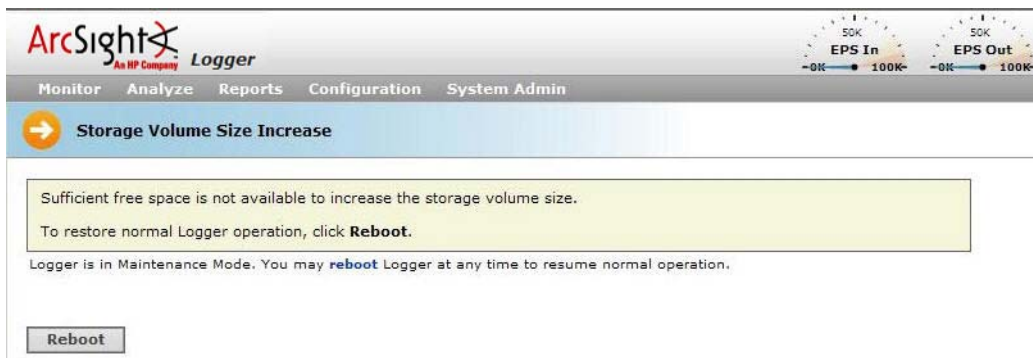
You can increase the size of Logger's storage volume to a maximum of **800 GB**. The current storage volume size is **744 GB**. Enter the new storage volume size and click **OK**.

Logger is in Maintenance Mode. You may **reboot** Logger at any time to resume normal operation.

New size (GB)

OK

If sufficient space is not found to increase the storage volume, the following message is displayed. Click **Reboot** to restart the Logger and exit the maintenance mode.



ArcSight An HP Company **Logger**

Monitor Analyze Reports Configuration System Admin

Storage Volume Size Increase

Sufficient free space is not available to increase the storage volume size. To restore normal Logger operation, click **Reboot**.

Logger is in Maintenance Mode. You may **reboot** Logger at any time to resume normal operation.

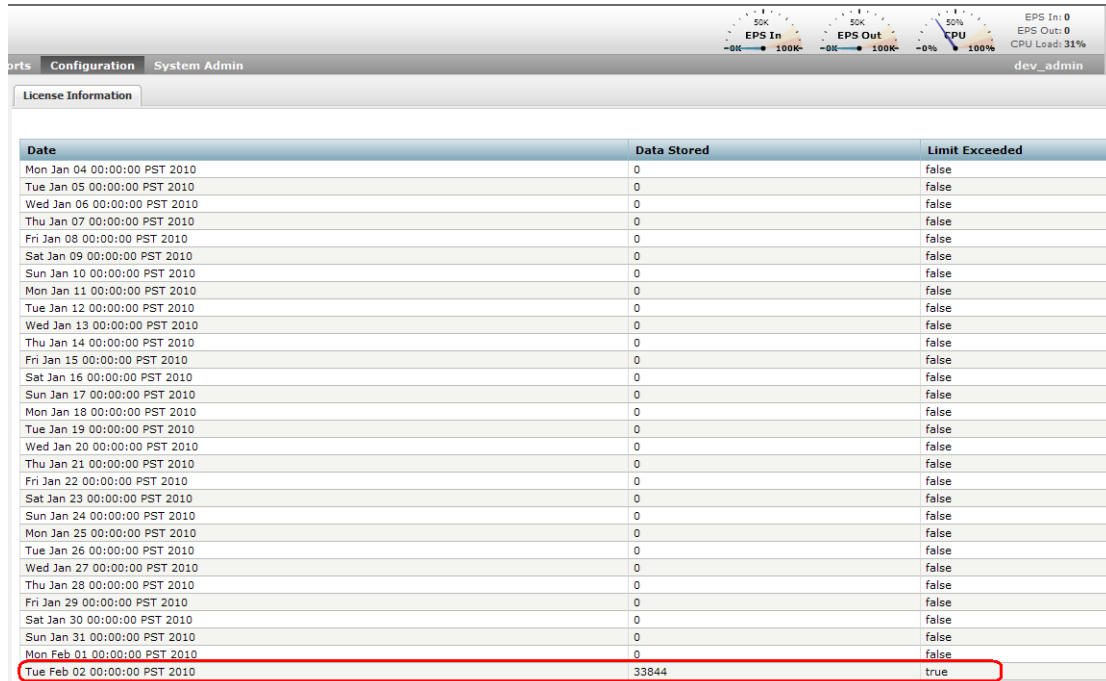
Reboot

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.

- 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.

Export Filters

Choose Items To Export

- Test1
- All Logins (CEF format)
- All Logins (Non-CEF format)
- All Logins (Unified)
- CEF
- Configuration Changes (Unified)
- High and Very High CEF Events
- High and Very High Events (Unified)
- Malicious Code (CEF format)
- Malicious Code (Unified)

Use ctrl-click to select or deselect items

Save to local disk ☒

Export

Export Filters

Choose Items To Export

- Test1
- All Logins (CEF format)
- All Logins (Non-CEF format)
- All Logins (Unified)
- CEF
- Configuration Changes (Unified)
- High and Very High CEF Events
- High and Very High Events (Unified)
- Malicious Code (CEF format)
- Malicious Code (Unified)

Use ctrl-click to select or deselect items

Save to local disk ☐

Export to remote file system

Mount Location

Specify file path without extension

Remote file path and name

Overwrite if file exists ☐

Export

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.

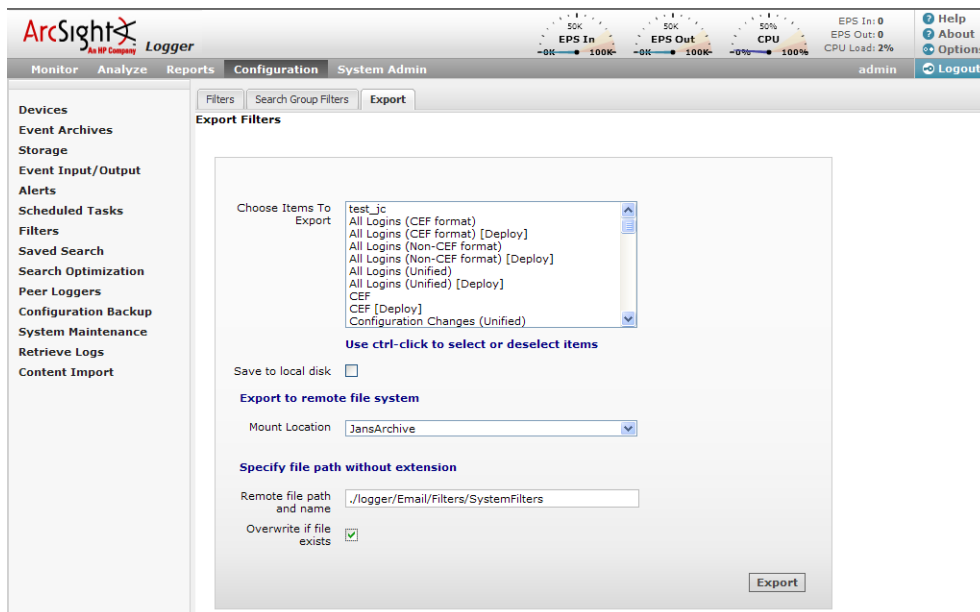
Importing Guidelines

- 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

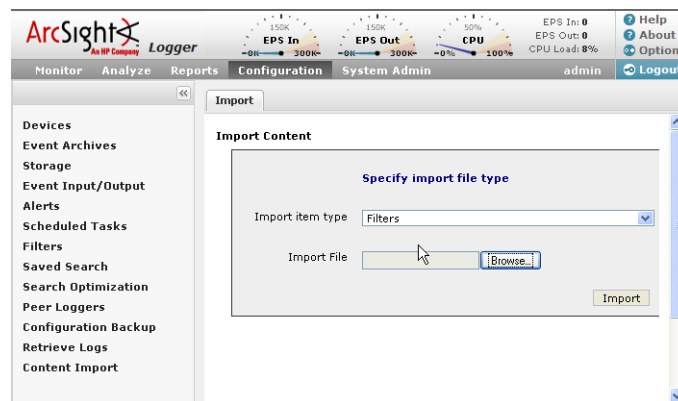


To export Alerts or Filters:

- 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 314](#).
- 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:

["System Locale" on page 302](#)
["Reboot" on page 302](#)
["DNS Settings" on page 303](#)
["Hosts" on page 303](#)
["Network" on page 304](#)
["Time/NTP" on page 306](#)
["Static Routes" on page 308](#)
["SMTP Settings" on page 308](#)
["License & Update" on page 309](#)
["Process Status" on page 310](#)
["SSH Access to Logger" on page 311](#)
["Logs - Audit Logs" on page 313](#)
["Logs - Audit Forwarding" on page 313](#)
["Storage" on page 314](#)
["SAN" on page 318](#)
["Security" on page 323](#)
["Users/Groups" on page 333](#)
["Users/Groups - Change Password" on page 345](#)

System Locale

The Locale setting ensures that the user interface displays information such as date, time, numbers, and messages in the format and language appropriate for the selected country.

On a new Logger appliance, Locale is generally configured during the Logger initialization process. However, if you skipped this step during the initialization process, you can configure it later. Once configured, Locale cannot be changed.

To view or configure System Locale:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **System Locale** from the System section.
- 3 To configure System Locale, select the new value.

You must reboot Logger for the changes to become effective. See ["Reboot" on page 302](#).

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 Network Settings

DNS | Hosts | Network | Time/NTP | 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 302](#).

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 Network Settings

DNS | **Hosts** | Network | Time/NTP | 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 localhosts.)
- 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 302](#).

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 302](#).)

ArcSight Network Settings

DNS | Hosts | **Network** | Time/NTP | Static Routes

Network Settings

Note: Settings take effect after reboot.

System Hostname

Default Gateway

☐ Automatically route outbound packets (interface homing)

<p>NIC 'ETH0'</p> <p>IP Address <input type="text" value="192.168.35.162"/></p> <p>Mask <input type="text" value="255.255.255.0"/></p> <p>Speed/Duplex <input type="text" value="Auto (recommended)"/></p>	<p>NIC 'ETH1'</p> <p>IP Address <input type="text" value="192.168.36.35"/></p> <p>Mask <input type="text" value="255.255.255.0"/></p> <p>Speed/Duplex <input type="text" value="Auto (recommended)"/></p>
---	--

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 324.</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 324.</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 benefit.</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 302](#)) 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.

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 Network Settings

DNS | Hosts | Network | **Time/NTP** | Static Routes

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 appliance, 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: 15:44:12

Date: 09 / 16 / 2010 **Time:** 15 : 33 : 39

Timezone Configuration

Local Timezone
US/Pacific

☒ **Enable appliance as NTP server** (Click for Status)

NTP Server List (Click to Test)
time.nist.gov

- 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 302](#).

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.

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.




Figure 7-4 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.

SMTP Settings

Alerts use Simple Mail Transfer Protocol (SMTP) to send e-mail.

To change SMTP configuration:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SMTP** from the System section.

- 3 On the **SMTP** page, enter new values for the following fields.

Parameter	Description
Primary SMTP Server	Enter the IP address or hostname of the SMTP server that will process outgoing e-mail.
Backup SMTP Server	Enter the IP address or hostname of the SMTP server that will process outgoing e-mail in case the primary SMTP server is unavailable.
Outgoing Email Address	The e-mail address that will appear in the From: field of outbound e-mail.

- 4 Click **Save** to make the changes. You must reboot the Logger for changes to take effect.

License & Update

You can update a Logger appliance or apply a license to it on this page. The procedures described in this section are only for Logger appliances and **not** for software Logger.

Updating system software requires uploading an upgrade file you downloaded from the ArcSight Customer Support web site. This page also displays the elapsed time since the appliance was last rebooted, license information, and the version of the Logger components. The Logger version and build number is found at 'arcsight-logger'.

To upload an upgrade file:

- 1 Download the update file to the computer from which you are accessing the Logger user interface.
- 2 Connect to the Logger user interface.
- 3 Click **System Admin** from the top-level menu bar.
- 4 Click **License & Update** from the System section.
- 5 Click **Browse** to locate the file.
- 6 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 302](#). A reboot is not required if you are only updating the license.

To apply a license file to a Logger appliance:

- 1 Download the license file from the ArcSight software download site at <https://arcsight.subscribenet.com> to a computer from which you can connect to Logger.
- 2 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.
- 3 Click the **System Admin** tab > **License & Update**.
- 4 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.

Process Status


The Process Status page lists all Logger processes, including the Logger service, that are running on Logger. This page enables you to view the details of those processes and start, stop, or restart them.

To view the Process Status page:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Process Status** from the System section to display a page similar to the one shown in the following figure.

In the process list, processors refers to Forwarders.


Process Status


 Refresh Status


System						
System	Status	Load	CPU Usage	Memory Usage	Data Collected	
n035-h027.qa.arcsight.com	running	[1.89] [1.65] [1.89]	16.8%us 1.9%sy 1.2%wa	68.0% [4076648 kB]	09/15/2010 15:01:23	











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


Processes

 Start


 Stop

 Restart

Process	Status	Uptime	CPU Usage	Memory Usage
 apache	running	4h 25m	0.0%	0.0% [3320 kB]
 aps	running	4h 25m	0.3%	5.0% [303676 kB]
 connector	running	4h 15m	0.0%	0.0% [608 kB]
 mysqld	running	4h 23m	0.0%	1.0% [61668 kB]
 postgresql	running	4h 25m	0.0%	0.1% [9260 kB]
 processors	running	4h 15m	0.1%	4.3% [261348 kB]
 receivers	running	4h 14m	0.0%	3.4% [204392 kB]
 reportengine	running	4h 15m	0.0%	4.8% [289668 kB]
 servers	running	4h 17m	0.0%	31.6% [1894552 kB]
 web	running	4h 15m	3.2%	4.8% [292716 kB]

To view the details of a process, click the  icon to the left of the process name, as shown in the following figure.



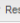
Process Status











 Refresh Status

System					
System	Status	Load	CPU Usage	Memory Usage	Data Collected
n035-h027.qa.arcsight.com	running	[1.50] [1.84] [1.90]	3.9%us 0.5%sy 2.1%wa	68.4% [4097992 kB]	09/15/2010 15:07:54

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

Processes

 Start  Stop  Restart

Process	Status	Uptime	CPU Usage	Memory Usage
 apache	running	4h 32m	0.0%	0.0% [3320 kB]
Children 0 CPU Percent 0.0% CPU Percent Total 0.0% Data Collected 09/15/2010 15:08:09 Memory Kilobytes 3320 Memory Kilobytes Total 96128 Memory Percent 0.0% Memory Percent Total 1.6% Monitoring Status monitored Parent PID 1 PID 28151 Status running Uptime 4h 32m				
 aps	running	4h 32m	0.2%	5.1% [311040 kB]
 connector	running	4h 22m	0.0%	0.0% [608 kB]
 mysqld	running	4h 30m	0.0%	1.0% [61668 kB]
 postgresql	running	4h 32m	0.0%	0.1% [9260 kB]
 processors	running	4h 21m	0.0%	4.3% [261496 kB]
 receivers	running	4h 21m	0.0%	3.4% [204392 kB]
 reportengine	running	4h 21m	0.0%	4.8% [289672 kB]
 servers	running	4h 24m	0.0%	31.7% [1901108 kB]
 web	running	4h 21m	3.1%	4.6% [281096 kB]

To start, stop, or restart a process, select the process and click **Start**, **Stop**, or **Restart** at the top of the process list.

SSH Access to Logger

When you report an issue to ArcSight Customer Support that requires them to access your appliance for troubleshooting and diagnostics in situations such as an upgrade failure, unresponsive appliance, and so on, they will direct you to enable SSH access on it.

By default, SSH access (known as Support Login in previous releases) to your Logger appliance is disabled; however, you can select one of these options in the Logger user interface to enable it:

- Enabled—SSH access is always enabled.
- Enabled, only for 8 hours—SSH access is automatically disabled eight hours after it was enabled.
- Enabled, only during startup/reboot—SSH access is enabled during the time Logger reboots and is starting up. It is disabled once all processes on Logger are up and running. This option provides a minimal period of SSH access for situations such as when Logger does not start successfully after a reboot.



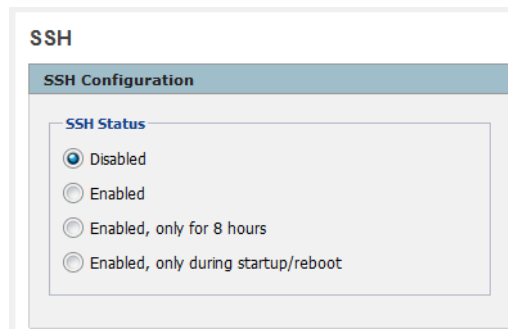
Note

Even if SSH is disabled on your Logger, you can access its console if you have it setup for remote access using the HP ProLiant Integrated Lights-Out (iLO) Advanced remote management card. For more information about setting up your Logger to access its console remotely, see [“Configure Logger for Remote Access” on page 33](#).

Enabling or Disabling SSH Access

To enable or disable SSH access to your Logger appliance:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSH** from the System section.
- 3 Select one of the following options.



Once you select an option, the user interface displays a message that requires you to confirm the action. Once you confirm it, the change takes effect.

Connecting to Logger using SSH

Once you have enabled the SSH access, follow these steps to connect to it using SSH:

- 1 Connect to the Logger appliance as "root" using an SSH client.
- 2 When prompted to enter a password, enter any text and press **Enter**.
You are prompted to enter a response to the challenge string that is displayed on your screen.
- 3 Call ArcSight Customer Support to obtain the challenge response string. Enter it at the "Enter response:" prompt and press **Enter**.

```
login as: root
root@192.168.36.29's password:
Last login: Thu Mar 17 01:50:38 2011 from 10.4.10.190
Challenge is 46024. Enter response: 184096
[root@logger ~]# pwd
/root
[root@logger ~]#
```

If the correct string is entered, you are connected to the Logger for the amount of time specific to the option you had selected in ["Enabling or Disabling SSH Access" on page 311](#).

Logs - Audit Logs

Logger audit logs are available for viewing.

Audit Logs

Search Audit Logs

Timestamp --

Description

User

Search Results

User	Description	Timestamp ▼
admin	Session expired	09/16/2010 11:00:14
admin	Session expired	09/16/2010 10:38:13
admin	Session expired	09/16/2010 10:16:13
admin	Session expired	09/16/2010 10:12:12
admin	Session expired	09/16/2010 10:02:12

To view Audit logs:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Audit Logs** from the Logs section.
- 3 Select the date and time range for which you want to obtain the log.
- 4 To refine the audit log search, optionally specify a string in the Description field and user name in the User field. When a string is specified, only logs whose Description field contains the string are displayed. Similarly, when a user is specified, only logs whose User field contains the username are displayed.
- 5 Click **Search**.



Audit logs, as Common Event Format (CEF) audit events, can be sent to ArcSight ESM directly for analysis and correlation.

For more information about audit event forwarding, see [“Logs - Audit Forwarding” on page 313](#). For information about audit events that you can forward, see [Appendix E, Logger Audit Events, on page 483](#).

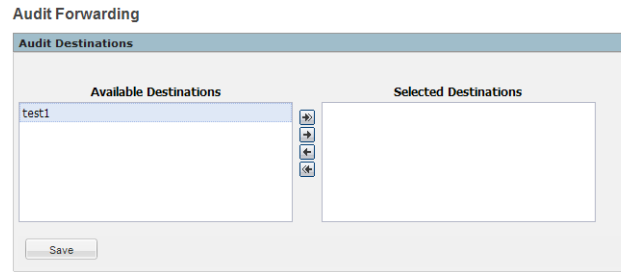
Logs - Audit Forwarding

For information about audit events that you can forward, see [Appendix E, Logger Audit Events, on page 483](#).

To forward audit events to specific destinations:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Audit Forwarding** from the Logs section.
- 3 Select destinations from the Available Destinations list and click the right arrow icon () to move the selected destination to the Selected Destinations list. You can select multiple destinations at the same time and move them over. Or you can move all available destinations by clicking the () icon.

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**

EPS In: 0 EPS Out: 0 CPU: 0% CPU Load: 2%

Help About Options Logout

System Admin

- System
 - Reboot
 - Network
 - SMTP
 - License Update
 - Process Status
 - Support Login
- Logs
 - Audit Logs
 - Audit Forwarding
- Storage
 - CIFS**
 - NFS
 - RAID Controller
 - Hard Disk SMART Data
- Security
 - SSL Server Certificate
 - SSL Client Authentication
 - FIPS 140-2
- Users/Groups
 - Authentication
 - User Management
 - 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:



Note

You can edit a CIFS mount only if it is not in use. The edit link is not displayed if the mount cannot be edited.

When you rename a CIFS mount, access to the archives that were made using the original name is lost until you revert the mount point name to the original name.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **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:



Note

You can delete a CIFS mount only if it is not in use. The delete link is not displayed if the mount cannot be deleted.

When you delete a CIFS mount that is in use for event archiving, access to the archives is lost until you add a mount point of the same name and same location.

- 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.

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.



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 284](#).

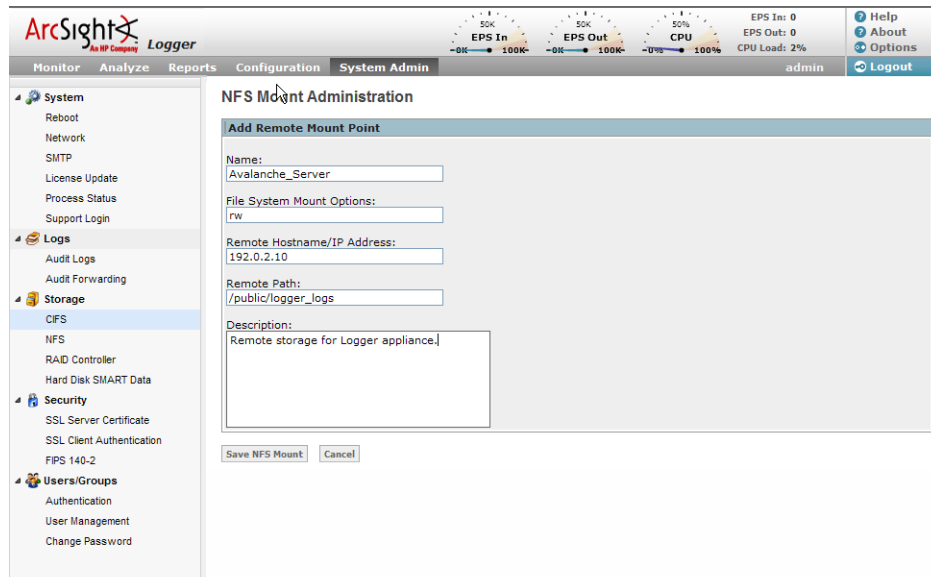


Figure 7-5 NFS Mount Administration page

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, ro for read-only from the remote host, rw for read-write, or hard to keep retrying until the remote host responds.
<p>Note: Even if you configure rw permission at your mount point, rw permission is not granted to the remote host if the host is configured to allow read-only access.</p>	

Parameter	Description
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:



Note

You can edit an NFS mount only if it is not in use. The edit link is not displayed if the mount cannot be edited.

When you rename an NFS mount, access to the archives that were made using the original name is lost until you revert the mount point name to the original name.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **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:



Note

You can delete an NFS mount only if it is not in use. The delete link is not displayed if the mount cannot be deleted.

When you delete an NFS mount that is in use for event archiving, access to the archives is lost until you add a mount point of the same name and same location.

- 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.

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 7-6](#), 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 7-6 SAN Logical Unit state diagram.



Caution

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 475 .
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.

SAN

SAN Configuration

Detach
 Refresh

<input checked="" type="checkbox"/>	LUN Name	Device	Type	Manufacturer	Mfr. Uni...	WWN	Size	Status
<input checked="" type="checkbox"/>	LoggerLun	/dev/mpath/mpath0	xfp	CLARIION	13CE016...	500601613ce01664:0000000000000000	100.00GB	Attached

HBA Information

HBA Model

HP 8Gb Dual Channel PCI-e 2.0 FC HBA

HBA FW Version

1.11A5 (U3D1.11A5), sl-3

HBA Driver Version

Emulex LightPulse Fibre Channel SCSI driver 8.2.0.63.3p

To attach a LUN:



Note

- Logger can attach to only one LUN (on SAN) at a time for primary storage. You can add an additional LUN for event archival, configuration backup, and export. If multipathing is enabled on your Logger, you cannot use an additional LUN for event archival, configuration backup, and export. For information about multipathing, see ["Multipath" on page 322](#).
- You can attach a LUN only if the storage volume is not set up on Logger yet and the LUN is in "Available" state on Logger.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN** under the Storage section in the left panel.
- 3 Locate and select the LUN in the LUN Name List.
- 4 Click **Attach** from the top left of the SAN Configuration page.

The LUN's Attachment Status will change to "Attached" when the LUN is ready for use.

Note: If you do not see the Attach menu option, there are no LUNs available that can be attached to the Logger at this time.

To detach a LUN:



You cannot detach a LUN if a storage volume is configured on it.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN** under the Storage section in the left panel.
- 3 In the LUN Name List, locate the LUN to be detached.
- 4 Click **Detach** from the top left of the SAN Configuration page.

Note: If you do not see the Detach menu option, there are no LUNs available that can be detached from the Logger at this time.

To re-attach a LUN:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN** under the Storage section in the left panel.
- 3 In the LUN Name List, locate the LUN to be re-attached. The LUN must be in the 'detached' state.
- 4 Click **Re-attach** from the top left of the SAN Configuration page.

Note: If you do not see the Re-attach menu option, there are no LUNs available that can be re-attached from the Logger at this time.

To destroy a LUN:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN** under the Storage section in the left panel.
- 3 In the LUN Name List, locate the LUN to be destroyed. The LUN must be in the 'detached' state.
- 4 Click **Destroy** from the top left of the SAN Configuration 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 SSH access to your Logger (see ["SSH Access to Logger" on page 311](#)). 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/pgsql/data
```

When Customer Support has finished these tasks, reboot Logger.

Multipath

When you multipath a LUN, you create two different network paths to it from the system to which the LUN connects. Doing so reduces the possibility of a single point of failure causing the LUN to become unavailable.

By default, the HBA card on your Logger has two ports. Prior to Logger v5.1, you could only connect one port to the LUN; however, starting with v5.1, you can connect both of those ports to the same LUN.

Although any SAN vendor that supports multipathing can work with Logger, ArcSight specifically tests with EMC Clariion SANs. Logger provides a default multipath configuration as a starting point. However, make sure that you consult your SAN documentation for information specific to your set up and environment.

Multipath user interface (UI) is available by default on Logger models that support SAN. However, you must connect the LUN to both HBA ports and configure multipath configuration in the UI for it to function. Once enabled, **multipath cannot be disabled on Logger**.



Note

- Multipath does not enable you to attach an additional LUN to Logger. Only one LUN can be attached at any given time.
 - If multipathing is enabled on your Logger, you cannot use an additional LUN for event archival, configuration backup, and export.
 - Connecting two LUNs to the HBA card ports without configuring multipath will result in unexpected results.
-

For new Logger installations, you must configure multipathing before attaching the LUN.

If you are upgrading from a version prior to Logger v5.1 and want to enable multipath on your Logger, see the release notes for this release for information.

Enabling Multipath

To enable multipath:

- 1 Ensure that a LUN is attached to the Logger, as described in [“SAN” on page 318](#).
- 2 Click **System Admin** from the top-level menu bar.
- 3 Click **Multipath** under the Storage section in the left panel.
- 4 Select a SAN Multipathing Configuration from the drop-down menu.
- 5 If you chose Custom or if the displayed configuration does not meet your needs, customize the parameters.
- 6 Click **Test** to ensure the configuration you chose or the changes you made are valid.

If the test fails, make additional changes or click **Reset** to start over again.

- 7 Click **Save** to save the configuration.

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 7-7](#).



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	: 11223334455667788
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
NRAM	: 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 7-7 RAID Controller Information page

This information 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 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 displays the ArcSight Logger web interface. At the top, there's a status bar with performance metrics: EPS In (50K), EPS Out (50K), CPU (50%), EPS In: 1,324, EPS Out: 0, and CPU Load: 72%. Below this is a navigation menu with tabs: Monitor, Analyze, Reports, Configuration, and System Admin (selected). The System Admin tab is active, showing a sidebar with categories: System (Reboot, Network, SMTP, License Update, Process Status, Support Login), Logs (Audit Logs, Audit Forwarding), Storage (CIFS, NFS, RAID Controller, Hard Disk SMART Data), Security (SSL Server Certificate, SSL Client Authentication, FIPS 140-2), and Users/Groups (Authentication, User Management, Change Password). The main content area is titled 'ArcSight SSL Settings' and has three sub-tabs: Generate CSR (selected), Install Cert, and View Results. The 'Generate CSR' tab contains the 'Generate Certificate Signing Request' form. The form prompts the user to 'Please enter the Certificate Settings' and includes the following fields: Country (2-letter code) set to 'US', State/Province set to 'California', City/Locality set to 'Cupertino', Organization Name set to 'ArcSight, Inc.', Organizational Unit set to 'Support Team', Hostname set to 'loggerA.arcsight.com', Email Address set to 'support@arcsight.com', and a Private Key Password field. A 'Generate CSR' button is located at the bottom left of the form.

Figure 7-8 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 7-8 on page 324](#).
- 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	<p>The host name or IP address of this Logger.</p> <p>When specifying the host name, make sure that this name matches the name registered in the Domain Name Service (DNS) server for the Logger. Additionally, this name must be identical to the host name specified in "Network" on page 304.</p> <p>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.</p>
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.
Private Key Length	The size to specify the security strength of your Private Key in bits.

- 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

Figure 7-9 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 7-9 on page 326](#)), 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 323](#) to load the certificate.
- 2 Enable client certificate authentication, as described in [“Client Certificate Authentication” on page 335](#).



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 327](#).

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/Groups” on page 333](#).
- 5 (Optional) Upload a certificate revocation list (CRL), as described in [“Uploading a Certificate Revocation List” on page 328](#).
- 6 (Optional) If this Logger is configured to use **only** SSL Client Authentication, make sure this Logger's Authorization ID and Code are appropriately configured on other Loggers that with it. For more information, see [“Peer Loggers” on page 280](#).

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 331](#) 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 251](#) to complete configuration of this setup.
- Loggers
Loggers running v4.0 automatically use FIPS 140-2 compliant algorithms. Therefore, no action is required on a 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 ed 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 323](#).
- 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.
- The File Transfer Receivers are set up using FTP. The SCP and SFTP protocols (for setting up File Transfer Receivers) are not FIPS compliant.

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 329.

- 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.

The screenshot shows the ArcSight Logger web interface. The top navigation bar includes Monitor, Analyze, Reports, Configuration, and System Admin (selected). The left sidebar shows a tree view with categories: System (Reboot, Network, SMTP, License Update, Process Status, Support Login), Logs (Audit Logs, Audit Forwarding), Storage (CIFS, NFS, RAID Controller, Hard Disk SMART Data), Security (SSL Server Certificate, SSL Client Authentication, FIPS 140-2), and Users/Groups (Authentication, User Management, Change Password). The main content area is titled 'Enable/Disable FIPS Mode'. It contains a 'Configure FIPS Mode' section with a warning: 'Do not perform any FIPS-related activity on the appliance while the FIPS mode change is in progress.' Below this is a 'Select FIPS Mode' section with radio buttons for 'Enable' (selected) and 'Disable', and a 'Save' button. At the bottom is a 'FIPS Status Table' with two columns: 'Name' and 'FIPS Enabled'. The table lists the following components: Apache Web Server, ESM Forwarder, Processor, Receiver, Server, and Tomcat.

Name	FIPS Enabled
Apache Web Server	
ESM Forwarder	
Processor	
Receiver	
Server	
Tomcat	

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 331.
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 331.
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 331 .

- 1 Follow device configuration steps provided in the SmartConnector's configuration guide (available from the ArcSight Download Center at <https://arcsight.subscribenet.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 .crt 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 Login (Global) 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 7-10](#), also lets you specify how many seconds of inactivity to allow before automatically ending the current session. The default is 900 (15 minutes).

Authentication Settings

Login Passwords Authentication

Global Settings

Max Simultaneous Logins per User
15

Session Inactivity Timeout in Seconds
900

Days After Which an Inactive User Account is Disabled
0

Save Settings

Figure 7-10 Login Settings page

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 7-10 on page 333](#).
- 3 Enter new values for the maximum simultaneous logins per user, the session inactivity time-out, or the number of days after which an inactive user account is disabled.



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

- 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 302](#).

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.

The screenshot shows the 'Authentication Settings' window with the 'Passwords' tab selected. The 'Password Settings' section includes the following options:

- Enable Password Lockout:** Radio buttons for 'Yes' (selected) and 'No'. Below are input fields for:
 - Number of failed attempts before lockout: 3
 - Maximum time between attempts (in seconds): 60
 - Lockout duration (in minutes): 15
- Enable Password Expiration:** Radio buttons for 'Yes' (selected) and 'No'. Below are input fields for:
 - Days until password expires: 90
 - Days before expiration to notify user: 5
- Enable Password Validation:** Radio buttons for 'Yes' (selected) and 'No'.
- Password Length Limits:** Input fields for:
 - Minimum password length: 10
 - Maximum password length: 20
- Minimum Requirements:** Input fields for:
 - Numeric characters [0-9]: 2
 - Uppercase characters [A-Z]: 0
 - Lowercase characters [a-z]: 0
 - Non-alphanumeric characters [!\$%^*...]: 2
 - Number of characters different from old password: 2

A 'Save Settings' button is located at the bottom of the form.

Figure 7-11 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 7-11 on page 334](#).
- 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 302](#).

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 336](#).)

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/Groups” on page 333](#) 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.



Caution

All SSL client certificates used for authentication must be FIPS compliant (that is, hashed with FIPS compliant algorithms) even if FIPS is not enabled on your Logger.

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/Groups" on page 333 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/Groups" on page 333](#)) 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.


User 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.

Groups are organized by type, as shown in [Figure 7-14](#). Each user group is one of the following types: System Admin, Logger Rights, Logger Search, or Logger Reports.

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

 Note	The user group privileges available on a system depend on the Logger model and the installed license. For example, SAN-related privileges are only displayed on Logger models that support SAN.
Section	Privilege
Reboot	Reboot Appliance. (See “Reboot” on page 302.)
Update	Update Appliance. Enable Maintenance Mode (See “System Maintenance” on page 287.)
System Information	Process Status. (See “Process Status” on page 310.) RAID Controller. (See “RAID Controller” on page 323.)
SSL Certificates	Generate SSL Certificate Signing Request (CSR). (See “Generating a Certificate Signing Request” on page 324.) Install new SSL Certificates. (See “Installing a Signed Certificate” on page 326.)

Section	Privilege
Platform Settings	<p>Configure DNS Settings. (See "DNS Settings" on page 303.)</p> <p>Configure Network Settings. (See "Network" on page 304.)</p> <p>Configure Time Settings. (See "Time/NTP" on page 306.)</p> <p>Configure SMTP Settings. (See "Static Routes" on page 308.)</p> <p>Configure Static Routes. (See "Static Routes" on page 308.)</p> <p>Configure Hosts File. (See "Hosts" on page 303.)</p> <p>Configure Security Settings. (See "Security" on page 323.)</p> <p>Set the Application Locale. (See "System Locale" on page 302.)</p>
External File Systems	Configure NFS, CIFS, and SAN settings. (See "Storage" on page 314 , "CIFS Settings" on page 314 , and "SAN" on page 318.)
Global Settings	<p>Configure Login Settings. (See "Authentication Settings" on page 333.)</p> <p>Configure Password Settings. (See "Password" on page 334.)</p> <p>Configure Password Authentication. (See "Authentication" on page 335.)</p> <p>Configure Audit Forwarding Destination. (See "Logs - Audit Forwarding" on page 313.)</p>
System Logs	View Audit Logs. (See "Logs - Audit Logs" on page 313.)
User/Groups	<p>Manage Users. (See "Users/Groups" on page 333.)</p> <p>Manage User Groups. (See "Users/Groups" on page 333.)</p> <p>Run User Entitlement Reports.</p>
Console Access	<p>Allow Console Access. (See "Using the Command Line Interface to Configure IP Address" on page 25.)</p> <p>Control support login access. (See "SSH Access to Logger" on page 311.)</p>
Application Options	<p>View Options. (See "Options" on page 65.)</p> <p>Edit, Save, and Remove Options. (See "Options" on page 65.)</p>

Table 7-2 Logger Rights Groups

Section	Privilege
Monitor	<p>Monitor Logger throughput. (See "Monitor" on page 66.)</p> <p>Monitor Logger throughput on remote peers. (See "Monitor" on page 66 and "Peer Loggers" on page 280.)</p>

Section	Privilege
Filters	<p>Use and view shared filters. (See “Filters” on page 270.)</p> <p>Edit, save, and remove shared filters. (See “Filters” on page 270.) Also, import and export filters.</p>
Peers	<p>View registered peers. (See “Peer Loggers” on page 280.)</p> <p>Edit, save, and remove registered peers. (See “Peer Loggers” on page 280.)</p>
Devices and Device Groups	<p>View devices. (See “Devices” on page 223.)</p> <p>Edit, save, and remove devices. (See “Devices” on page 223.)</p> <p>View device groups. (See “Device Groups” on page 225.)</p> <p>Edit, save, and remove device groups. (See “Device Groups” on page 225.)</p>
Receivers	<p>View receivers. (See “Receivers” on page 239.)</p> <p>Edit, save, and remove receivers. (See “Receivers” on page 239.)</p>
Forwarders and Alerts	<p>View forwarders and alerts. (See “Forwarders” on page 246 and “Alerts” on page 255.)</p> <p>Edit, save, and remove forwarders and alerts. (See “Forwarders” on page 246 and “Alerts” on page 255.) For alerts, this privilege enables you to import and export them.</p>
ESM Connectors	<p>View ESM connectors. (See “ESM Destinations” on page 251.)</p> <p>Edit, save, and remove ESM connectors. (See “ESM Destinations” on page 251.)</p>
Search Filters	<p>View search group filters (aka user group filters). (See “Search Group Filters” on page 272.)</p> <p>Edit, save, and remove search group filters. (See “Search Group Filters” on page 272.)</p>
Configuration Backup	<p>View backups. (See “Configuration Backup and Restore” on page 284.)</p> <p>Edit, save, and remove backups. (See “Configuration Backup and Restore” on page 284.)</p>
Retrieve Logs	Download system logs. (See “Retrieve Logs” on page 296.)
Scheduling	View scheduled tasks. (See “Scheduled Tasks” on page 269.)
Storage Groups	<p>View storage groups. (See “Storage Groups” on page 233.)</p> <p>Edit and add storage groups. (See “Storage Groups” on page 233.)</p>
Event Archive/Restore	<p>View event archives. (See “Archiving Events” on page 229.)</p> <p>Edit, save, and remove event archives. (See “Archiving Events” on page 229.)</p>

Section	Privilege
Saved Search	View Saved Search. (See “Saved Searches” on page 273.) Edit, save, and remove Saved Search. (See “Scheduled Saved Search” on page 275.)
Fieldsets	View fieldsets. (See “Field Set” on page 94.) Edit, save, and remove fieldsets. (See “Deleting Custom Field Sets” on page 280.)
Scheduled Searches and Alerts	View scheduled searches and alerts. (See “Creating and Managing Saved Search Alerts” on page 260.) Edit, save, and remove scheduled searches and alerts. (See “Creating and Managing Saved Search Alerts” on page 260.)

Table 7-3 Logger Search Groups

Section	Privilege
Search	Search for events. (See “The Need to Search Events” on page 71.) Search for events on remote peers. (See “Searching Peer Loggers (Distributed Search)” on page 111.)

Table 7-4 Logger Report Groups.

Section	Privilege
Report	Global access to all report objects and permission to change reporting configuration. (See Chapter 5, Reporting, on page 133.) 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 164.) Edit and save report style. This overrides the corresponding permission on individual report groups. (See “Applying Report Template Styles” on page 214.) View all published reports. This overrides the corresponding permission on individual report groups. (See Chapter 5, Reporting, on page 133.) View, run, and schedule all reports. This overrides the corresponding permission on individual report groups. (See “Running, Viewing, and Publishing Reports” on page 154 and “Scheduling Reports” on page 215.) Edit and save reports. This overrides the corresponding permission on individual report groups. (See Chapter 5, Reporting, on page 133.)

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.

Managing a User Group

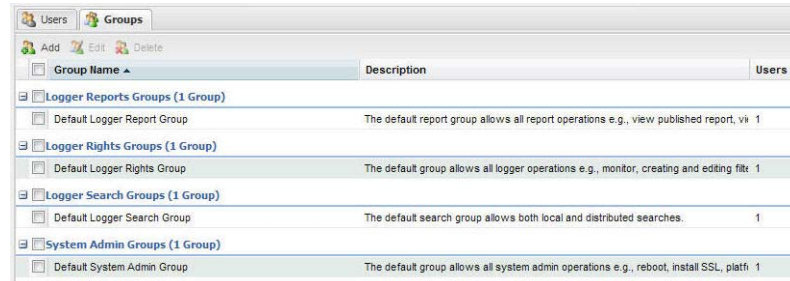


Figure 7-12 Groups page

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

To create a new user group:

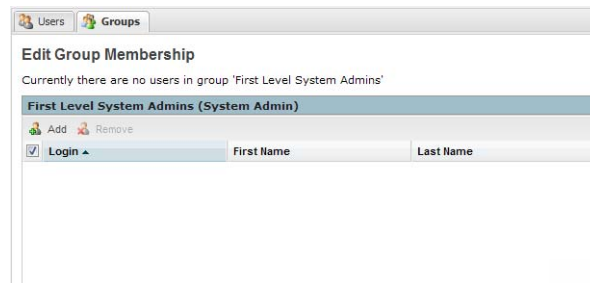
- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** under the Users/Groups section in the left panel to display the page shown in [Figure 7-12 on page 341](#).
- 3 Click the **Groups** tab.
- 4 Click **Add** from the top left side of the page.
- 5 Enter the definition of the new group.
 - a Enter a meaningful name for the group in the Group Name field.
 - b Enter a meaningful description for the group in the Description field.
 - c Select the group type—System Admin, Logger Rights, Logger Reports, Logger Search.
 - d Click the down arrow icon (▼) to define the group's rights and permissions.
- 6 Click **Save and Close** to save the settings of the group. OR click **Save and Edit Membership** to add users to this group.

To edit a user group:


- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** under the Users/Groups section in the left panel to display the page shown in [Figure 7-12 on page 341](#).
- 3 Click the **Groups** tab.
- 4 Select the Group that you want to edit.
- 5 Click **Edit** at the top left side of the page.
- 6 Update the user group information as necessary.

If you need to edit the group's membership:

- a Click **Save and Edit Membership** to display the Edit Group Membership page, as shown in the following figure.



- b Click **Add** from the top left of the Edit Group Membership page.
- c Select users you want to add. By default, you can only add users who do not belong to other groups of the type that you are editing. However, if you want to add such users, click **Show users that belong to other <group_type> groups**, as shown in the following figure. *When you add a user who belongs to another group of the same group type as the one you are updating, that user is automatically removed from the previous group.*



- d Click **OK**.

7 Click **Back to Group List**.

To delete a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** under the Users/Groups section in the left panel to display the page shown in [Figure 7-12 on page 341](#).
- 3 Click the **Groups** tab.
- 4 Select the Group (or Groups) that you want to delete.
- 5 Click **Delete** at the top left side of the page.

Users

Figure 7-13 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 **User Management** under the Users/Groups section in the left panel.
- 3** In the Users tab, click **Add** from the top left side of the page.
- 4** Enter the following parameters.

Parameter	Description
Login	A login name for the user
Password	A password for the user. This field is optional if you use RADIUS or CAC authentication.
Confirm Password	Reenter the password.

Parameter	Description
First Name	User's first name. If you enabled SSL client authentication (see "SSL Client Authentication (CAC Authentication)" on page 326), click Use Client DN to enter the Distinguished Name (Certificate Subject) information for the user. Distinguished Name should be in this format: <code>/ST=California/C=US/L=Cupertino/O=ArcSight, Inc./OU=Engg Team/CN=UserA/D/emailAddress=email@xyz.com</code> To determine the DN, use this URL to display the certificate: <code>https://<hostname or IP address of Logger>/app/cert</code> OR 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 > Advanced > Encryption > View Certificates > Your Certificates > Select the certificate > View .
Last Name	User's last name.
Email	An e-mail address for the user.
Phone Number	User's phone number.
Assign to Groups	Select the groups to which this user belongs. This setting controls the privileges a user has on this Logger.

5 Click **Save and Close**.

To edit a user:

- 1** Click **System Admin** from the top-level menu bar.
- 2** Click **User Management** under the Users/Groups section in the left panel.
- 3** In the Users tab, select the user (or users) you want to edit.
- 4** Click **Edit** from the top left side of the page.
- 5** Update the user information as necessary.
- 6** Click **Save User**.

To delete a user:



A system admin level user account that has been used to upgrade the Logger cannot be deleted. For example, if a system admin user Joe upgrades the Logger, Joe's user account can not be deleted from Logger once the upgrade is complete. To remove such a user, disable the user account. To disable a user account, edit the user account and disable the "Active" option.

- 1** Click **System Admin** from the top-level menu bar.
- 2** Click **User Management** under the Users/Groups section in the left panel.
- 3** In the Users tab, select the user (or users) you want to delete.
- 4** Click **Delete** from the top left side of the page.

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

Change Password for Default Admin

Old Password

New Password

New Password (confirm)

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 previous figure.
- 3 Enter the old password, the new password, and enter the new password a second time to confirm.
- 4 Click **Change Password**.



Passwords are subject to the password policy specified by the Admin user. See ["Password" on page 334](#).

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.

- ["System - System Locale" on page 345](#)
- ["System - SMTP Settings" on page 346](#)
- ["System - Process Status" on page 346](#)
- ["Logs - Audit Logs" on page 348](#)
- ["Logs - Audit Forwarding" on page 348](#)
- ["Users/Groups - User Management" on page 349](#)
- ["Users/Groups - Change Password" on page 355](#)
- ["Using a CA-signed Certificate on Software Version of Logger" on page 356](#)

System - System Locale

The Locale setting ensures that the user interface displays information such as date, time, numbers, and messages in the format and language appropriate for the selected country.

On a new Logger software, Locale is generally configured during the Logger configuration process using the Logger Configuration Wizard. However, if you skipped this step during

the configuration process, you can configure it later. Once configured, Locale cannot be changed.

To view or configure the System Locale:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **System Locale** from the System section.
- 3 To configure System Locale, select the new value.

You must reboot Logger for the changes to become effective.

System - SMTP Settings

Alerts use Simple Mail Transfer Protocol (SMTP) to send e-mail.

To change SMTP configuration:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SMTP** from the System section.
- 3 On the **SMTP** 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.
Backup SMTP Server	Enter the IP address or hostname of the SMTP server that will process outgoing e-mail in case the primary SMTP server is unavailable.
Outgoing Email Address	The e-mail address that will appear in the From: field of outbound e-mail.

- 4 Click **Save** to make the changes. You must reboot the Logger for changes to take effect.

System - Process Status

The Process Status page lists all Logger processes, including the Logger service, that are running on Logger. This page enables you to view the details of those processes and start, stop, or restart them.

To view the Process Status page:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Process Status** from the System section to display a page similar to the one shown in the following figure.

In the process list, processors refers to Forwarders.

Process Status

[Refresh Status](#)

System						
System	Status	Load	CPU Usage	Memory Usage	Data Collected	
n035-h027.qa.arcsight.co	running	[1.89] [1.65] [1.89]	16.8%us 1.9%sy 1.2%wa	68.0% [4076648 kB]	09/15/2010 15:01:23	

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

Processes

[Start](#) [Stop](#) [Restart](#)

Process	Status	Uptime	CPU Usage	Memory Usage
apache	running	4h 25m	0.0%	0.0% [3320 kB]
aps	running	4h 25m	0.3%	5.0% [303676 kB]
connector	running	4h 15m	0.0%	0.0% [608 kB]
mysqld	running	4h 23m	0.0%	1.0% [61668 kB]
postgresql	running	4h 25m	0.0%	0.1% [9260 kB]
processors	running	4h 15m	0.1%	4.3% [261348 kB]
receivers	running	4h 14m	0.0%	3.4% [204392 kB]
reportengine	running	4h 15m	0.0%	4.8% [289668 kB]
servers	running	4h 17m	0.0%	31.6% [1894552 kB]
web	running	4h 15m	3.2%	4.8% [292716 kB]

To view the details of a process, click the icon to the left of the process name, as shown in the following figure.

Process Status

[Refresh Status](#)

System						
System	Status	Load	CPU Usage	Memory Usage	Data Collected	
n035-h027.qa.arcsight.co	running	[1.50] [1.84] [1.90]	3.9%us 0.5%sy 2.1%wa	68.4% [4097992 kB]	09/15/2010 15:07:54	

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

Processes

[Start](#) [Stop](#) [Restart](#)

Process	Status	Uptime	CPU Usage	Memory Usage
apache	running	4h 32m	0.0%	0.0% [3320 kB]
Children 0 CPU Percent 0.0% CPU Percent Total 0.0% Data Collected 09/15/2010 15:08:09 Memory Kilobytes 3320 Memory Kilobytes Total 96128 Memory Percent 0.0% Memory Percent Total 1.6% Monitoring Status monitored Parent PID 1 PID 28151 Status running Uptime 4h 32m				
aps	running	4h 32m	0.2%	5.1% [311040 kB]
connector	running	4h 22m	0.0%	0.0% [608 kB]
mysqld	running	4h 30m	0.0%	1.0% [61668 kB]
postgresql	running	4h 32m	0.0%	0.1% [9260 kB]
processors	running	4h 21m	0.0%	4.3% [261496 kB]
receivers	running	4h 21m	0.0%	3.4% [204392 kB]
reportengine	running	4h 21m	0.0%	4.8% [289672 kB]
servers	running	4h 24m	0.0%	31.7% [1901108 kB]
web	running	4h 21m	3.1%	4.6% [281096 kB]

To start, stop, or restart a process, select the process and click **Start**, **Stop**, or **Restart** at the top of the process list.

Logs - Audit Logs

Logger audit logs are available for viewing.

Audit Logs

Search Audit Logs

Timestamp ---

Description

User

Search Results

User	Description	Timestamp ▼
admin	Session expired	09/16/2010 11:00:14
admin	Session expired	09/16/2010 10:38:13
admin	Session expired	09/16/2010 10:16:13
admin	Session expired	09/16/2010 10:12:12
admin	Session expired	09/16/2010 10:02:12

To view Audit logs:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Audit Logs** from the Logs section.
- 3 Select the date and time range for which you want to obtain the log.
- 4 To refine the audit log search, optionally specify a string in the Description field and user name in the User field. When a string is specified, only logs whose Description field contains the string are displayed. Similarly, when a user is specified, only logs whose User field contains the username are displayed.
- 5 Click **Search**.



Audit 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 348](#). For information about audit events that you can forward, see [Appendix E, Logger Audit Events, on page 483](#).

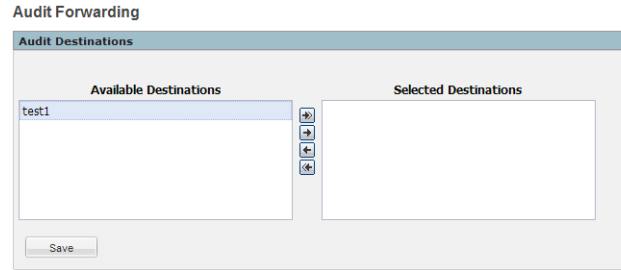
Logs - Audit Forwarding

For information about audit events that you can forward, see [Appendix E, Logger Audit Events, on page 483](#).

To forward audit events to specific destinations:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Audit Forwarding** from the Logs section.
- 3 Select destinations from the Available Destinations list and click the right arrow icon () to move the selected destination to the Selected Destinations list. You can select multiple destinations at the same time and move them over. Or you can move all available destinations by clicking the () icon.

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 - User Management

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 Search groups, or Logger Rights groups.

User Groups

Groups are organized by type, as shown in [Figure 7-14](#). Each user group is one of the following types: System Admin, Logger Rights, Logger Search, or Logger Reports.

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
Update	Enable Maintenance Mode (See "System Maintenance" on page 287.)
System Information	Process Status. (See "Process Status" on page 310.)
Platform Settings	Configure SMTP Settings. (See "Static Routes" on page 308.)
	Set the Application Locale. (See "System Locale" on page 302.)
Global Settings	Configure Audit Forwarding Destination. (See "Logs - Audit Forwarding" on page 313.)
System Logs	View Audit Logs. (See "Logs - Audit Logs" on page 313.)
User/Groups	Manage Users. (See "Users/Groups" on page 333.)
	Manage User Groups. (See "Users/Groups" on page 333.)
Application Options	View Options. (See "Options" on page 65.)
	Edit, Save, and Remove Options. (See "Options" on page 65.)

Table 7-6 Logger Rights Groups

Section	Privilege
Monitor	Monitor Logger throughput. (See "Monitor" on page 66.) Monitor Logger throughput on remote peers. (See "Monitor" on page 66 and "Peer Loggers" on page 280.)
Filters	Use and view shared filters. (See "Filters" on page 270.) Edit, save, and remove shared filters. (See "Filters" on page 270.) Also, import and export filters.
Peers	View registered peers. (See "Peer Loggers" on page 280.) Edit, save, and remove registered peers. (See "Peer Loggers" on page 280.)
Devices and Device Groups	View devices. (See "Devices" on page 223.) Edit, save, and remove devices. (See "Devices" on page 223.) View device groups. (See "Device Groups" on page 225.) Edit, save, and remove device groups. (See "Device Groups" on page 225.)
Receivers	View receivers. (See "Receivers" on page 239.) Edit, save, and remove receivers. (See "Receivers" on page 239.)
Forwarders and Alerts	View forwarders and alerts. (See "Forwarders" on page 246 and "Alerts" on page 255.) Edit, save, and remove forwarders and alerts. (See "Forwarders" on page 246 and "Alerts" on page 255.) For alerts, this privilege enables you to import and export them.
ESM Connectors	View ESM connectors. (See "ESM Destinations" on page 251.) Edit, save, and remove ESM connectors. (See "ESM Destinations" on page 251.)
Search Filters	View search group filters (aka user group filters). (See "Search Group Filters" on page 272.) Edit, save, and remove search group filters. (See "Search Group Filters" on page 272.)
Configuration Backup	View backups. (See "Configuration Backup and Restore" on page 284.) Edit, save, and remove backups. (See "Configuration Backup and Restore" on page 284.)
Retrieve Logs	Download system logs. (See "Retrieve Logs" on page 296.)
Scheduling	View scheduled tasks. (See "Scheduled Tasks" on page 269.)

Section	Privilege
Storage Groups	View storage groups. (See "Storage Groups" on page 233.) Edit and add storage groups. (See "Storage Groups" on page 233.)
Event Archive/Restore	View event archives. (See "Archiving Events" on page 229.) Edit, save, and remove event archives. (See "Archiving Events" on page 229.)
Saved Search	View Saved Search. (See "Saved Searches" on page 273.) Edit, save, and remove Saved Search. (See "Scheduled Saved Search" on page 275.)
Fieldsets	View fieldsets. (See "Field Set" on page 94.) Edit, save, and remove fieldsets. (See "Deleting Custom Field Sets" on page 280.)
Scheduled Searches and Alerts	View scheduled searches and alerts. (See "Creating and Managing Saved Search Alerts" on page 260.) Edit, save, and remove scheduled searches and alerts. (See "Creating and Managing Saved Search Alerts" on page 260.)

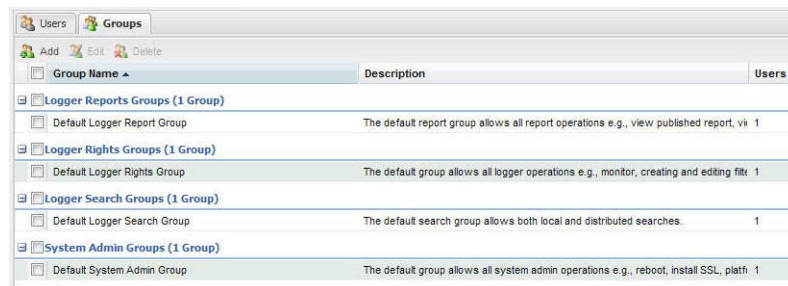
Table 7-7 Logger Search Groups

Section	Privilege
Search	Search for events. (See "The Need to Search Events" on page 71.) Search for events on remote peers. (See "Searching Peer Loggers (Distributed Search)" on page 111.)

Table 7-8 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 133.)</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 164.)</p> <p>Edit and save report style. This overrides the corresponding permission on individual report groups. (See "Applying Report Template Styles" on page 214.)</p> <p>View all published reports. This overrides the corresponding permission on individual report groups. (See Chapter 5, Reporting, on page 133.)</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 154 and "Scheduling Reports" on page 215.)</p> <p>Edit and save reports. This overrides the corresponding permission on individual report groups. (See Chapter 5, Reporting, on page 133.)</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.

**Figure 7-14** Groups page

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

Managing a User Group

To create a new user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** under the Users/Groups section in the left panel to display the page shown in [Figure 7-14](#).

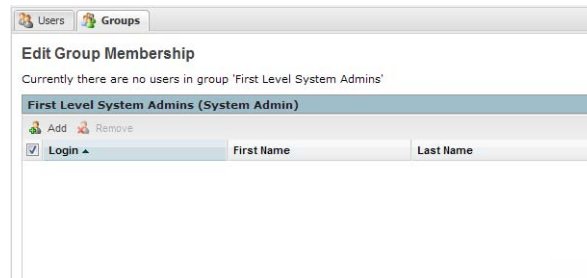
- 3 Click the **Groups** tab.
- 4 Click **Add** from the top left side of the page.
- 5 Enter the definition of the new group.
 - a Enter a meaningful name for the group in the Group Name field.
 - b Enter a meaningful description for the group in the Description field.
 - c Select the group type—System Admin, Logger Rights, Logger Reports, Logger Search.
 - d Click the down arrow icon (▼) to define the group's rights and permissions.
- 6 Click **Save and Close** to save the settings of the group. OR click **Save and Edit Membership** to add users to this group.

o edit a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** under the Users/Groups section in the left panel to display the page shown in Figure 7-14.
- 3 Click the **Groups** tab.
- 4 Select the Group that you want to edit.
- 5 Click **Edit** at the top left side of the page.
- 6 Update the user group information as necessary.


If you need to edit the group's membership:

- a Click **Save and Edit Membership** to display the Edit Group Membership page, as shown in the following figure.



- b Click **Add** from the top left of the Edit Group Membership page.
- c Select users you want to add. By default, you can only add users who do not belong to other groups of the type that you are editing. However, if you want to

add such users, click **Show users that belong to other <group_type> groups**, as shown in the following figure.



d Click **OK**.

7 Click **Back to Group List**.

To delete a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** under the Users/Groups section in the left panel to display the page shown in [Figure 7-14](#).
- 3 Click the **Groups** tab.
- 4 Select the Group (or Groups) that you want to delete.
- 5 Click **Delete** at the top left side of the page.

Users

Figure 7-15 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 **User Management** under the Users/Groups section in the left panel.

- 3 In the Users tab, click **Add** from the top left side of the page.
- 4 Enter the following parameters.

Parameter	Description
Login	A login name for the user
Password	A password for the user.
Confirm Password	Reenter the password.
First Name	User's first name.
Last Name	User's last name.
Email	An e-mail address for the user.
Phone Number	User's phone number.
Assign to Groups	Select the groups to which this user belongs. This setting controls the privileges a user has on this Logger.

- 5 Click **Save and Close**.

To edit a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** under the Users/Groups section in the left panel.
- 3 In the Users tab, select the user (or users) you want to edit.
- 4 Click **Edit** from the top left side of the page.
- 5 Update the user information as necessary.
- 6 Click **Save User**.

To delete a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** under the Users/Groups section in the left panel.
- 3 In the Users tab, select the user (or users) you want to delete.
- 4 Click **Delete** from the top left side of the page.

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

Change Password for Default Admin

Old Password

New Password

New Password (confirm)

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 previous figure.
- 3 Enter the old password, the new password, and enter the new password a second time to confirm.
- 4 Click **Change Password**.



Passwords are subject to the password policy specified by the Admin user. See ["Password" on page 334](#).

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 set the LD_LIBRARY_PATH:

```
export
LD_LIBRARY_PATH=<install_dir>/current/local/openssl/lib:$LD_LIBRARY_PATH
```

- 3 Run these commands to create a new private key and a server certificate request:

```
OPENSSL_FIPS=1 <install_dir>/current/local/openssl/bin/openssl
req -text -newkey rsa:1024 -keyout server.key -out newreq.csr
-config <install_dir>/current/local/openssl/ssl/openssl.cnf
```

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.

- 4 Send the generated CSR to a CA. Once you receive the signed certificate from the CA, go to the next step.
- 5 Run this command to archive any previously applied certificate files:

```
cd <install_dir>/current/local/apache/conf/ssl.crt
mv server.pem server.pem.good
mv server.crt server.crt.good
```

- 6 Run this command to install the signed certificate file you obtained from CA in Step 3:

```
cp newcertificate.crt
<install_dir>/current/local/apache/conf/ssl.crt/server.crt
```


- 7** Run this command to copy the server key to the Apache directory:

```
<install_dir>/current/local/openssl/bin/openssl rsa -in  
server.key -out <install_dir>/current/local/apache/  
conf/ssl.crt/server.pem
```

- 8** Stop the Logger service and related processes using this command:

```
<install_dir>/current/arcsight/logger/bin/loggerd stop
```

Wait until all processes have stopped, then run this command to start the service and processes:

```
<install_dir>/current/arcsight/logger/bin/loggerd start
```

Applying a License on the Software Version of Logger

See [“Applying a License on the Software Version of Logger”](#) on page 52.

Managing Connectors

The following topics are discussed here.

- ["Connector Overview" on page 360](#)
- ["Navigating the Manage Connectors Tab" on page 361](#)
- ["Locations" on page 363](#)
- ["Hosts" on page 367](#)
- ["Containers" on page 374](#)
- ["Connectors" on page 390](#)
- ["Configuration Suggestions for Connector Types" on page 424](#)

Connector Overview

You can manage the configuration of these kinds of connectors:

- **Local (on-board) connectors:** Pre-installed connectors on the Logger appliance running Connector Manager. Connector Manager (software edition) ships with no pre-installed connectors.
- **Remote Connector Appliance connectors:** Pre-installed connectors on a remotely-managed Connector Appliance.
- **Software-based connectors:** Software-based connectors installed manually on a remote host.

A connector configuration consists of properties such as name and type, and a set of *parameters* that customize how the connector works in a specific environment. Parameters vary based on the type of connector; for example, a connector for a firewall has different parameters than a connector that reads an intrusion detection system database.

You can manage connectors of many types, including syslog, Simple Network Management Protocol (SNMP), specific Intrusion Detection Systems (IDS), log files, vulnerability scanners, and operating system-specific security events. You can view the list of supported types in the drop-down menu when you configure a new connector.



The connectors you manage are configured automatically to run as *services* or *daemons*.

Individual software-based connectors are described in ArcSight documents specific to those connectors, including the connector-specific configuration guides available with each connector. You can also find general connector information in the *SmartConnector User's Guide*. All of these documents are available from the ArcSight Customer Support site.

Navigating the Manage Connectors Tab

The Manage Connectors tab enables you to configure and organize connectors. This section describes the user interface elements and explains how to use them effectively.

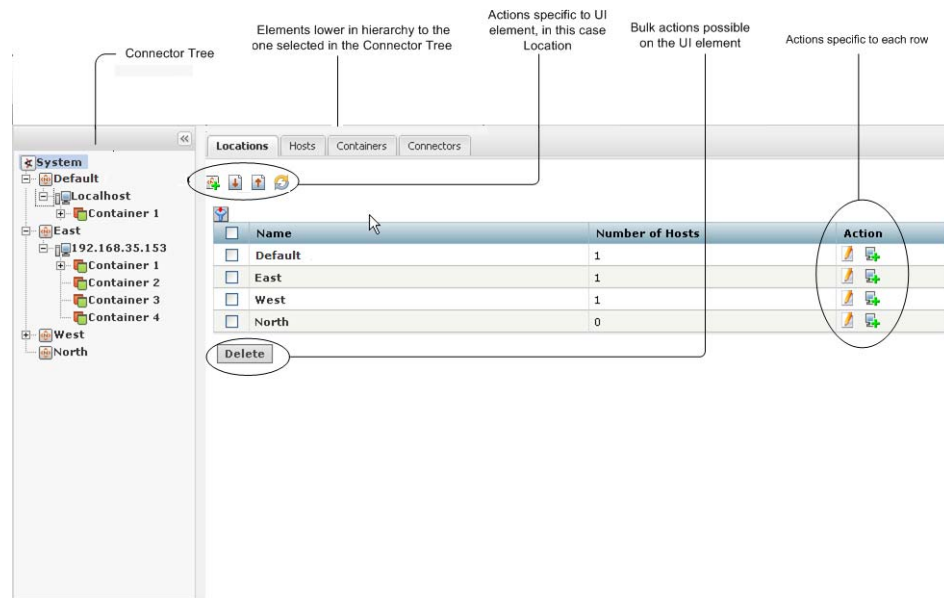



Figure 8-1 Managing Connectors

The Connector tree (the left panel of the window shown in [Figure 8-1](#)) organizes connectors into a hierarchy as follows:




Each connector you manage belongs to a container; each container belongs to a host; each host belongs to a location; and, all locations belong to root of the System.

When you click on an upper-level user interface element in the left panel, the interface displays elements lower in the hierarchy to it on the right panel. You can also perform management operations on the elements displayed on the right side.


For example, **System** provides the root (top-level) view. When you click System, all configured locations are listed in the left panel, as well as under the **Locations** tab in the right panel. You can perform various management tasks, such as editing, deleting, or adding a host, on those locations. In addition, all hosts, containers, and connectors on this system are displayed in specific tabs in the right panel. Click the **Hosts** tab to view all hosts on the system, and click **Containers** and **Connectors** to view the respective elements and perform management operations on them. Similarly, if you select a host (from the left

panel), all containers and connectors configured on that host are displayed on the right panel, as shown in the following figure.





When a container is down or a host is unreachable, the system waits for it to come online. There might be a delay of several minutes before the connector tree (in the left panel) and the Container tab (in the right panel) display.

On any user interface, you can perform three kinds of operations:

- A global operation—Listed on top of a user interface page; for example, you can upload a CSV file of locations.
- A localized operation—An operation on a single element displayed on the user interface page; for example, you can add a connector to a container by clicking the  icon in the Action column in the container's row.
- A bulk operation—A single operation performed on multiple elements on the user interface page; for example, you can upgrade multiple containers by selecting the containers (click the box to the left of the container to select it) and clicking Upgrade at the bottom of the page.



- The  icon refreshes a UI screen. This icon is available on the UI pages when relevant.
- Click the column filter icon () to display drop down lists of values on which to filter each table column. Click the 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. **Default** location is not provided on the Connector Manager (software edition).



ArcSight recommends that you do not delete the location **Default**.

You can view all locations on the system and view hosts, containers, and connectors in a location. You can add, edit, and delete a location. You can also add hosts to a location. All these procedures are described below.

Viewing All Locations

You can see all the locations that exist on the system.

To view all locations:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click **System** in the left panel.

All existing locations display on the Locations tab in the right panel.

Viewing Hosts, Containers, and Connectors in a Location

You can see all the hosts, containers, and connectors that exist in a location.

To view hosts, containers, and connectors in a location:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click the location (listed under System) from the left panel.

The hosts, containers, and connectors in the location display in the right panel, under specific tabs, as shown below.




Adding a Location

Before adding hosts, you need to add a location, which is a logical grouping of hosts.



You can also add locations in bulk using a comma-separated values (CSV) file. For more information see, [Adding Locations and Hosts from a File](#), below.

To add a location:

- 1 Click **Configuration** > **Manage Connectors**.
- 2 Click **System** in the left panel.
- 3 Click  (on top of the page) in the right panel.
- 4 Enter the name of the new location and click **Next**.
- 5 Click **Done**.

Exporting and Importing Remote Management Configuration

You can create a backup of the complete remote management configuration settings on the Connector Appliance (all remote software connectors and remote Connector Appliances that are managed by the appliance) and import the configuration on another appliance.


The remote management configuration is saved in AUP format in the Remote Management AUP repository so you can download it to your local computer.

You cannot manage the same connectors using two appliances at the same time. Before importing the remote management configuration to another Connector Appliance, you need to shut down the appliance from which you exported the configuration.



You can import the remote management configuration only on the same appliance model as the one from which the configuration is exported.

To export the remote management configuration:

- 1 Click **Configuration** > **Manage Connectors**.
- 2 Click **System** in the left panel.
- 3 Click  (on top of the page) in the right panel.
- 4 Follow the instructions in the wizard to export the configuration. The remote host configuration is saved in AUP format in the Remote Management AUP repository.

After you export the remote management configuration, you need to download it to your local computer from the Remote Management AUP repository.

After you have exported the remote management configuration and have downloaded it to your local computer, you can import the configuration to another appliance.



Importing the remote management configuration overwrites the current remote management configuration on the appliance.

To import the remote management configuration:

- 1 On the appliance where you want to copy the remote management configuration, click **Configuration > Manage Connectors**.
- 2 Click **System** in the left panel.
- 3 Click (on top of the page) in the right panel.
- 4 Follow the instructions in the wizard. When selecting the type of upload, choose **Full remote management (AUP format)**.



If there are no valid CA certificates for any connectors in the configuration, you see a question mark (?) next to the container that contains the connectors in the left panel. Refer to ["Resolving Invalid Certificate Errors" on page 385](#).

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.



A host is **not** added if:

- Any containers on the host are down.
- If you choose not to import the certificates that are retrieved.
- Authentication fails on any of the containers.

The CSV file needs to be in the format shown in the following example. Also, ensure that an end-of-line character is included in the last line of the CSV file if the file was created on a Windows system. However, an end-of-line character is not required if the file was created on a Linux system.

	A	B	C	D	E	F
1	Location	Hostname	Port	Type	User	Password
2	East	ernie.company.com	9006	8 Containers	admin	password
3	West	elmo.company.com	9008	Software	admin	password
4						

To add locations and hosts from a CSV file:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click **System** in the left panel.
- 3 Click (on top of the page) in the right panel to open the wizard.

- 4 Select **Remote hosts (CSV format)** and click **Next**. Follow the instructions in the wizard to upload the file.
- 5 Connector certificates are retrieved automatically so that the system can communicate with each connector in a container. The Upload CSV wizard lists the certificates. (To see certificate details, hover your mouse over the certificate.)
 - ◆ Select **Import the certificates to Connector Appliance from the containers**, then click **Next** to import the certificates and continue.
 - ◆ Select **Do not import the certificates to Connector Appliance from the containers** and click **Next** if you do not want to import the certificates. The Upload CSV wizard does not complete the upload CSV process.





The Upload CSV wizard does not complete the upload if certificate download failed for any of the connectors in a container or if any of the certificates failed to import into the trust store on the system.

Editing a Location

You can edit the name of a location from the System-level page or from a specific Location page.

To edit a location:

- 1 Click **Configuration > Manage Connectors**.
- 2 From the System-level page:
Click **System** (left panel) > **Locations** tab (right panel) >  in the Action column.
From a specific Location page:
Click **System** (left panel) > *Location* >  (on top of the page, in the right panel).
- 3 Enter the new name of the location and click **Next**.
- 4 Click **Done**.

Deleting a Location

When you delete a location, the hosts, containers, and connectors that it contains are also deleted.

To delete a location:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click **System** in the left panel.
- 3 Select the location you want to delete. You can select multiple locations.
- 4 Click **Delete** at the bottom of the page, in the right panel.

Adding Hosts to a Location

See [“Adding a Host” on page 368](#).

Hosts

A host is a computer on a network, associated with an IP address, on which connectors are installed. A host can be of three types:

- The Localhost (the local Connector Appliance or the Logger appliance running Connector Manager). By default, **Localhost** exists on a brand new Connector Appliance or Logger appliance running Connector Manager; it contains a default number of containers, which are empty.

Connector Manager (software edition) does not ship with Localhost.

- A remotely-managed Connector Appliance.
- A Software-type host (a Windows, Linux, or UNIX system running software-based connectors from ArcSight). A software-type host can contain up to 20 containers.

You can view all hosts on the system, and view containers and connectors in a host. You can add, scan, delete, and edit a host. You can move a host to a different location and upgrade a host remotely. You can also add a container to a host. All these procedures are described below.

Viewing All Hosts

You can see all the hosts you are managing.

To view all hosts:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click **System** in the left pane. All hosts display on the Hosts tab in the right panel.

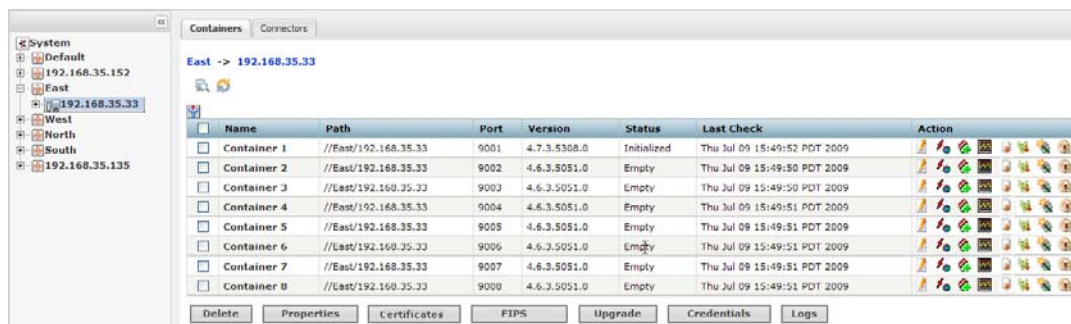
Viewing Containers and Connectors in a Host

You can see all the containers and connectors that exist on a host.

To view containers and connectors on a host:

- 1 Click **Configuration > Manage Connectors**.
- 2 In the left panel, click the location (under System) in which the host exists.
- 3 In the left panel, click the host to view the containers and connectors.

All containers display on the Containers tab and all connectors display on the Connectors tab in the right panel.



Adding a Host

By default, a local host **Localhost** exists on your Connector Appliance or Logger appliance running Connector Manager. However, Connector Appliance can manage connectors installed on other Connector Appliances and other systems such as Windows, UNIX, or Linux. To manage remote connectors, you need to add the hosts on which those connectors are running.



Connector Manager (software edition) does not ship with the default Localhost. You need to add the hosts that contain the connectors you want to manage.

When you add a host, the system also attempts to retrieve the CA certificates from the containers that reside on the host. Containers on the remote host can be managed only if the system can authenticate using the certificates and the credentials. When the certificates are retrieved, you are prompted to import them.



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.



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 365](#).

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 370](#).

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**.

Wizard

Host

Provide the remote host parameters

Location [Default]

Hostname/IP

Starting Port

Ending Port

User

Password

Comment

Hardware Type

Cancel Next

Wizard

Host

Provide the remote host parameters

Location [Default]

Hostname/IP

Starting Port

Ending Port

User

Password

Comment

Hardware Type

Cancel 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.
Ending Port	By default, Connector Appliance scans port 9001 to port 9020 when adding a host. If you select software in the Hardware Type field, you can specify the ending port number (for example, 9003) to speed up the add host process.
Password	The password for the user name you specify.
Comment	A meaningful description for the host you are adding.

Parameter	Description
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. <p>For the number of connectors 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 http://www.arcsight.com/supportportal.</p> <ul style="list-style-type: none"> If you want to remotely manage connectors running on a Windows, UNIX, or Linux system, select Software. <p>The system can detect the presence of software-based connectors on remote hosts using the Starting Port value you specified earlier. The system scans up to 20 configurable ports from the starting port to find the “listening” connectors.</p> <p>Any found connectors are added into the host. For more information, see “Scanning a Host” on page 370.</p>

- 4 Connector certificates are retrieved automatically so that the system can communicate with each connector in a container. The Add Host wizard lists the certificates. (To see certificate details, hover your mouse over the certificate.)

- ◆ Select **Import the certificates to Connector Appliance from the containers**, then click **Next** to import the certificates and add the host.
- ◆ Select **Do not import the certificates to Connector Appliance from the containers** and click **Next** if you do not want to import the certificates. Connector Appliance does not add the host.



Note

The Add Host wizard does not add the host if the certificate download failed for any of the connectors in a container or if any of the certificates failed to import into the trust store.

Scanning a Host

Scanning a host enables the system to detect new or removed containers from a remote **software-type** host. When a software-type host is added for the first time, it is scanned automatically for containers running at that time; however, to keep this information up-to-date, you need to scan the host manually whenever you add connectors to the remote host.

You can scan a host from the System-level page, a specific Location page, or a specific Host page.



Note

- You can scan only software-type hosts. See “[Hosts](#)” on page 367 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.




Note

A host cannot be scanned (the scan fails) if:

- Any containers on the host are down.
- If you choose *not* to import the certificates that are retrieved.
- Authentication fails on any of the containers.

To scan a host:

- 1 Click **Configuration** > **Manage Connectors**.
- 2 From the System-level page, click **System** (left panel) > **Locations** tab (right panel).
From a specific Location page, click **System** (left panel) > *Location* (under which the host exists).
From a specific Host page, click **System** (left panel) > *Location* (under which the host exists) > *Host*.
- 3 Click  in the Action column for the host that you want to scan.
- 4 Click **Next** in the Host Scan wizard.
- 5 Enter values for the parameters in the following table, then click **Next**.

Parameter	Description
Starting Port	The port number on the host on which Connector Appliance starts scanning for containers.
Ending Port	The port number on the host on which Connector Appliance ends scanning for containers.
User	The user name that the system uses to authenticate with the host.
Password	The password for the user name you provide.

- 6 Connector certificates are retrieved automatically so that the system can communicate with each connector in a container. The Host Scan wizard lists the certificates. (To see certificate details, hover your mouse over the certificate.)
 - ◆ Select **Import the certificates to Connector Appliance from the containers**, then click **Next** to import the certificates and continue.
 - ◆ Select **Do not import the certificates to Connector Appliance from the containers** and click **Next** if you do not want to import the certificates. The Host Scan wizard does not continue the scan.



Note

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 366](#)) or move an existing host (as described in [“Moving a Host to a Different Location” on page 372](#)).

Upgrading a Host Remotely



The following instructions only apply to upgrading a remotely-managed Connector Appliance.

You can upgrade a single remotely-managed Connector Appliance or several remotely-managed Connector Appliances at the same time (in bulk). Follow these guidelines:

- The containers of the appliance being upgraded need to be managed on the system from which you will initiate the upgrade.

Remotely upgrading a Connector Appliance is a two-step process.

To upgrade a Connector Appliance remotely:

- 1** Upload a Connector Appliance .aup upgrade file from the ArcSight Customer Support site to the AUP repository.

This step is only required if the version that you want to upgrade does not already exist in the repository.


- 2** Push the .aup upgrade file to the remote Connector Appliances, as follows:
 - a** Click **Configuration > Manage Connectors**.
 - b** From the System-level page, click **System** (left panel) > **Hosts** tab (right panel).
From a specific Location page, click **System** (left panel) > *Location* (under which the host exists).
 - c** Select the host you want to upgrade. You can select multiple hosts.
 - d** Click **Upgrade** at the bottom of the page.
 - e** Follow the instructions in the upgrade wizard.

Adding a Container to a Host

See [“Adding a Container” on page 375](#).

Containers

A container is a single Java Virtual Machine (JVM) that can run up to four connectors. The following illustration depicts Container 1 and the connectors it runs.




A default number of containers exist on each Connector Appliance. The number depends on the Connector Appliance hardware platform. Each container is identified with a label (Container Name) and an associated port number (9001 or higher).

Connector Manager on a Logger appliance contains one default container in the default host **Localhost**. You cannot delete this container. Connector Manager (software edition) does not contain any default containers.

You can perform many operations on containers. You can view all containers on the system and view the connectors in a container. You can add, delete, and edit a container. You can update container properties and change container credentials. You can manage certificates on a container, run a command on a container, and upgrade a container to a specific connector version. You can also view and delete container logs and run the Logfu utility. All these procedures are described below.

If you see a question mark (?) next to a container in the left panel, as shown below, the connectors in the container cannot be authenticated. The CA certificates for the connectors might be no longer valid. Refer to [“Resolving Invalid Certificate Errors” on page 385](#).



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 370](#). 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 390](#).

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).

- 3 Click  in the Action column of the container whose name you want to change.
If you are on the specific Container page, , is at the top of the page.
- 4 Enter the new name in the **Name** field and click **Next**.
- 5 Click **Done**.

Deleting a Container

You can delete containers from *software-type* hosts only. All other hosts (for example, a remotely-managed Connector Appliance) have a fixed number of containers.

When you delete a container, the connectors that it contains are also deleted.

To delete a container:

- 1 Click **Configuration** > **Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Select the container you want to delete. You can select multiple containers.
- 4 Click **Delete**.

Updating Container Properties

You can update existing container properties (located in the `agent.properties` file), delete them, or add new ones.

To update container properties:

- 1 Click **Configuration** > **Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).

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 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 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.

Caution

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.



Before enabling FIPS on a container that contains software connectors running as a service, review the caveats listed in document *Installing FIPS-Compliant SmartConnectors*, available from 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 379](#)).

To enable or disable FIPS mode on a container:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:


User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Select the container on which you want to enable or disable FIPS mode. You can select multiple containers.
- 4 Click **FIPS**, then click **Next** to run the FIPS Management wizard.



- 5 Click **Enable FIPS Mode** or **Disable FIPS Mode**, then click **Next**.

If FIPS mode is already enabled or disabled on the container, the FIPS Management wizard indicates this on the Summary page.



- 6 Check that the appropriate CA certificates are in the trust store so that the connectors in the container can validate their configured destinations successfully. If necessary, add the appropriate certificates to the container. Refer to [Managing Certificates on a Container](#).

Managing Certificates on a Container

Connectors require a Certificate Authority (CA) issued or self-signed SSL certificate to communicate securely with a destination. The Certificate Management wizard, available from the Containers tab, helps you add and remove certificates on a container. Using the wizard, you can:

- Enable or disable a demo certificate on a container.
You can enable a demo certificate on a container that is in non-FIPS mode only.
- Add a certificate on a container.
- Add a CA Certs file on a container.
You can add a CA Certs file on a container that is in non-FIPS mode only.
- Remove a certificate from a container.

From the Containers tab and the Connectors tab, you can view details about the certificates applied to a container. See [“Viewing Certificates on a Container”](#) on page 383.

For information about resolving invalid certificates, see [“Resolving Invalid Certificate Errors”](#) on page 385.

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.

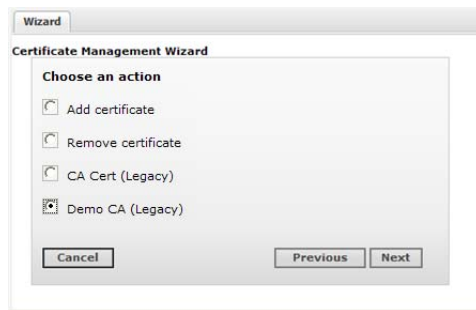
To enable or disable a demo certificate on a container:

- 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 **Certificates**, then click **Next** to run the Certificate Management wizard.
- 5 Click **Demo CA (Legacy)**, then click **Next**.



- 6 Follow the instructions in the Certificate Management wizard.

After you add the demo certificate on a container, the container restarts automatically.


Adding 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.

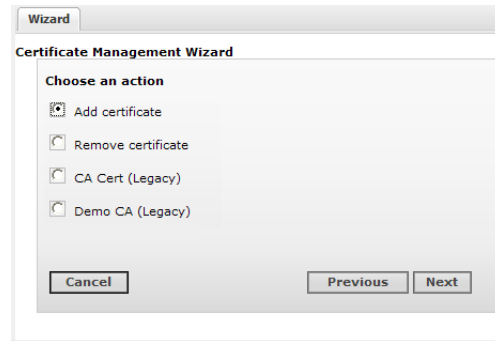
To apply a single CA certificate on a container:

- 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 certificate. You can select multiple containers.
- 4 Click **Certificates**, then click **Next** to run the Certificate Management wizard.
- 5 Click **Add Certificate** to add a certificate.



- 6 Follow the instructions in the wizard.

If a container is down or a connector is running an older build, the wizard reports errors in the progress bar and on the Summary page.



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.

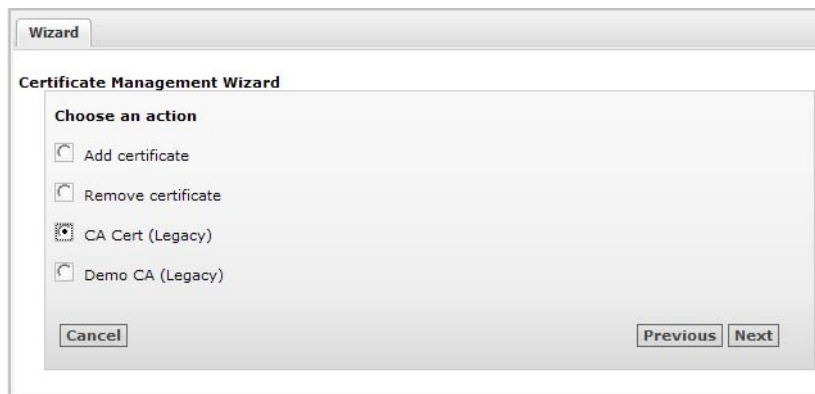
To add a CA Certs file on a container:

- 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 **Certificates**, then click **Next** to run the wizard.
- 5 Click **CA Cert (Legacy)**. You can add a CA Certs file to a container only if it is in non-FIPS mode.



- 6 Follow the instructions in the wizard.

After the CA Certs file has been added to a container, the container restarts automatically.

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:

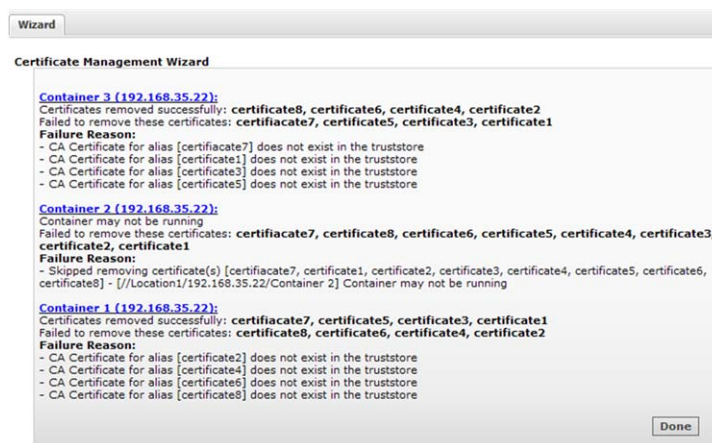
User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).

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 **Certificate**, then click **Next** to run the wizard.
- 5 Click **Remove certificate** and click **Next**.
- 6 Select one or more certificates from the certificate list and click **Next**.


The certificates are removed from the list of certificates and no longer used. When you remove a certificate from a container in FIPS mode, the container restarts automatically.

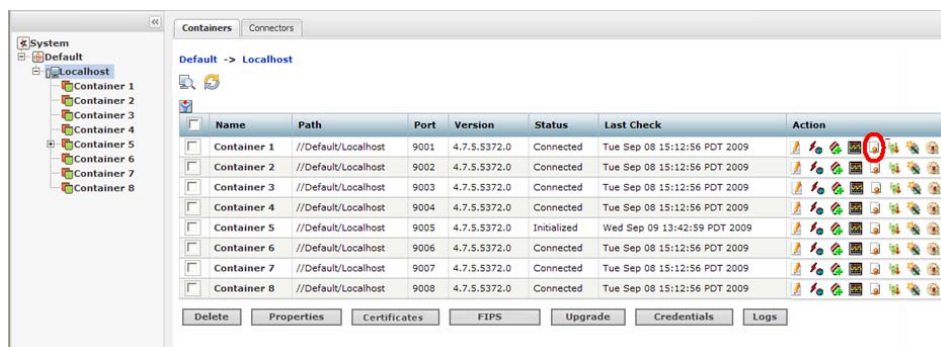
The Certificate Management wizard displays the certificates that are removed successfully in a comma-separated list. Certificates that cannot be removed are shown in a comma-separated list together with a reason why the certificate removal failed.




Viewing Certificates on a Container

From the Containers tab or the Connectors tab, you can display a list of the CA certificates applied to a container and view the details for a particular certificate in the list.

- On the **Containers** tab, click the  icon in the **Action** column for the container whose certificates you want to view.



- On the **Connectors** tab, select the  icon at the top of the page.



The Certificate List wizard displays the certificates applied to a container. To see details about a certificate, select the certificate and click **Next** at the bottom of the page.





Resolving Invalid Certificate Errors

If no valid CA certificates exist for the connectors in the container, you see a question mark (?) next to the container in the left panel, as shown below.



To resolve the invalid certificate error:

- 1 Click the container name in the left pane to view the certificate error on the Connectors tab.
- 2 Click the  icon to run the Certificate Download wizard.



- 3 Follow the instructions in the wizard to import the valid certificates.

Running a Command on a Container


You can run commands on a container to configure memory settings, pull an OPSEC certificate, or restart the container.

To run a command on a container:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).
From the Container page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Container .

- 3 Click  in the Action column of the container.

If you are on the specific Container page,  is at the top of the page.

- 4 Select the command you want to run and click **Next**.
- 5 Enter values for the parameters that the user interface displays and click **Done**.

Upgrading a Container to a Specific Connector Version

All connectors in a container are upgraded to the version you select.



You can't upgrade the same container more than once within a short period of time. After you upgrade a container, wait at least 15 minutes before upgrading it again.

To upgrade a container to a specific connector version:

- 1 Upload a connector build AUP from the 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.
- If you are upgrading an empty container, the system creates a temporary connector during the upgrade process. You can safely ignore this temporary connector; it is deleted shortly after being created.

Viewing Container Logs

You can retrieve and view the log files for a container. The log files are in `.zip` format.

To view container logs:

- 1 Load the logs to the Logs repository.

If the logs that you want to view are already in the Logs repository, skip this step.

- a Click **Configuration** > **Manage Connectors**.

- b Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- c Select the container whose logs you want to view. You can select multiple containers.

- d Click **Logs**.

The logs are loaded to the Logs repository. If you selected multiple containers, a log file for each container is loaded.

- 2 Retrieve and view the logs:

- a Click **Configuration** > **Repositories** from the top-level menu bar.

- b Click **Logs**.

- c Click  to retrieve the log files (in `.zip` format) you want to view.

Deleting Container Logs

To delete a container log file, click **Configuration** > **Repositories** > **Logs** > from the top-level menu bar. In the right panel, click  next to the log files you want to delete.

Running Logfu on a Container



The **Logfu** utility is a diagnostic tool that parses ArcSight logs to generate an interactive visual representation of the information contained within the logs.

When event flow problems occur (with a connector or the connected device), it is useful to have a visual representation of what happened over time. You can use Logfu to analyze this behavior.

To run Logfu on a container:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
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 connectors within the chosen container, plus many other types of data such as memory usage, and transport rates and logs.

Choose one of the Group box **data points**. Depending on which data point you chose, a list of fields appears in the Field box below.

- 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**.



Running Diagnostics on a Container

You can run certain diagnostics on a local or remote container. Currently, the **Edit a File** diagnostic action only is available:

To run diagnostics on a container:

- 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). The Connectors tab displays.

- 3 To open the Container Diagnostics wizard:
 - ◆ From the **Container** tab, click  in the **Action** column.
 - ◆ From the **Connectors** tab, click  at the top of the page.
- 4 Follow the steps in the wizard:
 - a Select the action you want to take on the selected container:
 - Select **Edit a configuration file** to edit a file in the *user/agent* folder on the container with the extension *.properties*, *.csv* or *.conf*.
 - Select **Edit a user file** to edit any file (except binary files, such as *.zip*, *.jar*, or *.exe*) in the *user/agent* folder on the container.
 - b From the list of available files, select the file you want to edit. The file displays in the Edit File panel. Make your edits, then click **Next** to save your edits and restart the container.



Tip

On Mozilla Firefox, if the text is underlined with red lines, right click on the text area and uncheck **Check Spelling**.



Note

When you click **Next**, Connector Appliance saves the updated file in the *user/agent* folder on the container; the original file is overwritten.

- c Click **Done** to close the Diagnostics wizard.

Connectors

A connector (also known as a Smartconnector) is an ArcSight software component that collects events and logs from various sources on your network. A connector can be configured on a Logger appliance running Connector Manager, on a Connector Appliance, or can be installed on a computer on your network and managed remotely. For a complete list of supported connectors, go to the 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, review the following important information.

- Make sure that the container, host, and location to which you want to add the connector exist on the system. If any of these elements do not exist, first create them using procedures described in [“Adding a Location” on page 364](#), [“Adding a Host” on page 368](#), and [“Adding a Container” on page 375](#).
- Follow the configuration best practices described in [“Configuration Suggestions for Connector Types” on page 424](#).

If you are configuring the Check Point OPSEC NG Connector, see [“Configuring the Check Point OPSEC NG Connector” on page 425](#).

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 428](#).

- If you are adding a software-based connector, make sure that the username and password for the connector match the username and password for the container to which you are adding the connector. Refer to [“Changing Container Credentials” on page 377](#).
- File-based connectors use the Network File System (NFS) or the Common Internet File System (CIFS).

For the file-based connectors on a Windows system, a CIFS share needs to be configured before you add those connectors. For information on creating a CIFS Mount, see [“CIFS Settings” on page 314](#).

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 316](#).

- 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 425](#).

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).
From the Container page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- 3 Click  in the Action column of the container to run the wizard to configure a connector.


If you are on the specific Container page,  is at the top of the page.

- 4 Select a connector type from the pull-down list of available types. Click **Next**.
- 5 Enter basic parameters for the connector. Parameters vary based on the connector type. You can hover the mouse pointer over a field for more information. When all fields have been entered, click **Next**.



When entering parameters that include a file path, enter the path in POSTIX format (for example, `/folder/filename`). If you enter the path in DOS/NTFS format (for example, `\folder\filename`), the backslash (`\`) is included as part of the file name and the path will be incorrect.

For file-based connectors on Windows systems, specify the name of the CIFS mount point you created for the connector, as shown in the following example. (You need to specify `/opt/mnt/CIFS_share_name`.)



Some connectors include table parameters. For example, the Microsoft Windows Event Log includes parameters for each host in the domain and one or more log types (security, application, system, directory service, DNS, file replication, and so on). You can import table parameters from a CSV file. See [“Adding Locations and Hosts from a File” on page 365](#) for the file format. You can import a CSV file that was exported from another connector as long as you export and import the CSV file from the same container. If the CSV file was exported from a different container, you need to change the secret parameters, such as the password, which appear in obfuscated format in the CSV file to plain text before you import the CSV file.



Note

For connectors that query Microsoft Active Directory to detect devices (for example, Microsoft Windows Event Log - Unified), if the “Network Security: LDAP Server Signing Requirements” policy is set to “Signing Required” on the Domain Controller, Connector Appliance will be unable to connect to the Active Directory or browse for devices. You see an error when selecting **Windows Host Browser** as the connector device browser type.



Note

For detailed information about individual connector parameters, refer to the specific *ArcSight SmartConnector Configuration Guide* for the type of connector chosen. The configuration guide also describes how to set up the source device for use with the connector.

- 6 Choose a primary destination for the connector and enter destination-specific parameters on the following page(s), then click **Next**. Destinations can be:

- ◆ ArcSight Logger SmartMessage (encrypted)
- ◆ ArcSight Manager (encrypted)
- ◆ CEF Syslog (cleartext, that is, unencrypted)



For containers running v5.1.2.5823 and later, Connector Appliance retrieves the certificate for the destination automatically and displays the certificate summary. To see certificate details, hover your mouse over the certificate.

- Select **Import the certificate to the connector from the destination**, then click **Next** to import the certificate and continue.
- Select **Do not import the certificate to the connector from the destination** and click **Next** if you do not want to import the certificate. The destination will not be added.

For containers running v5.1.2 and earlier, upload the certificate on the container and then add the destination.

Note: FIPS Suite B mode is not supported. Connector Appliance cannot download a manager certificate in Suite B mode.

- 7 Enter connector details:

Parameter	Description
Name	A descriptive name for this connector.
Location	The location of the connector (such as the hostname).
Device Location	The location of the device that sends events to the connector.
Comment	Additional comments.



Configuring a connector can take some time; the connector might initially display *Down* while it is restarting.

- 8 Click **Done**.

Editing Connector Parameters

ArcSight supports a large number of connector types to gather security events from a variety of sources, including syslog, log files, relational databases, and proprietary devices. Accordingly, configuration parameters vary widely depending on the type of connector being configured.

You can edit parameters (simple and table) for a specific connector or for multiple connectors at the same time.

Updating Simple Parameters for a Specific Connector

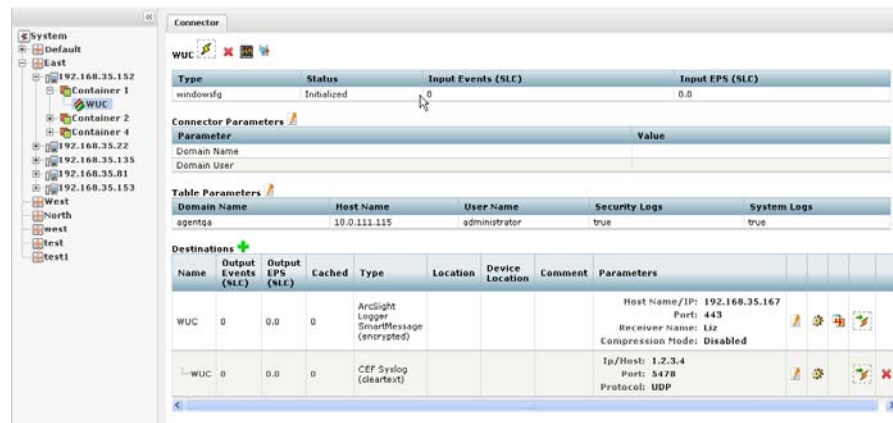
The following procedure describes how to update simple parameters for a specific connector. To update *table* parameters for a specific connector, see [“Updating Table Parameters for a Specific Connector” on page 396](#). To update both simple and table parameters for multiple connectors at the same time, see [“Updating Simple and Table Parameters for Multiple Connectors” on page 397](#).

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.



The screenshot shows the 'Connector' configuration page. On the left is a tree view of the system hierarchy. The main area displays the 'Connector Parameters' section, which is currently expanded. It contains three tables:

Type	Status	Input Events (SLC)	Input EPS (SLC)
windowsfg	Initialized	0	0.0

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: liz 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 connector being configured.
- When editing parameters that include a file path, enter the path in POSTIX format (for example, `/folder/filename`). If you enter the path in DOS/NTFS format (for example, `\folder\filename`), the backslash (\) is included as part of the file name and the path will be incorrect.

- 5 Click **Done** when complete.

The updated parameters display in the Connector Parameters section of the Connector page.


Updating Table Parameters for a Specific Connector

Certain connectors, such as the Microsoft Windows Event connector, have table parameters. You can update the table parameters for a specific connector when necessary.

To update table parameters for a specific connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click () next to the **Table Parameters** link.



Clicking the heading **Table Parameters** toggles between displaying and hiding the information in the Table Parameters section.

- 4 Modify parameters as necessary and then click **Next**.
 - ◆ To add more rows of parameter information, click the **Add Row** link.
 - ◆ You can use an Excel-compatible program to prepare a comma-separated values text file with the information and click the **Import File** button to load the entire table at once. The file needs to be in the same format as the rows shown on the Update Table Parameters page and needs to include a header row with parameter

labels in the order shown on that page. For fields that require checkbox values, enter True or False as the value. An example is shown below..

	A	B	C	D	E	F
1	Domain Name	Host Name	User Name	Password	Security Logs	System Logs
2	test	1.1.1.1	admin	password	TRUE	FALSE
3	test2	1.1.1.1.1	admin	password	TRUE	FALSE



Note

You can import a CSV file that was exported from another connector as long as you export and import the CSV file from the same container. If the CSV file was exported from a different container, you need to change the secret parameters, such as the password, which appear in obfuscated format in the CSV file to plain text before you import the CSV file.

- ◆ To export the table parameters to a CSV file for use as a backup or to import on another Connector Appliance, click the **Export File** button.

- 5 Click **Done** when complete.

The updated table parameters display in the Table Parameters section of the Connector page.

Updating Simple and Table Parameters for Multiple Connectors

If you have multiple connectors of the same type, you can change the simple and table parameters for all the connectors at the same time.

To edit parameters for multiple connectors at once:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel).
From the Connectors page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- 3 Select the connectors whose parameters you want to update.



Note

The connectors must be the same type; for example, you can change the parameters for several syslog connectors at the same time; however, you cannot change the parameters for several syslog and several SNMP connectors at the same time.

- 4 Click **Parameters**.
- 5 Follow the instructions in the wizard.

- ◆ You can choose to modify the simple parameters for all the selected connectors at once or modify the simple parameters per connector.
- ◆ If the connectors have table parameters, the table parameters are displayed so that you can modify them. If you have many table parameters to modify for multiple connectors, you can import the parameters from a CSV file (for information about adding rows and CSV file format, see [Step 3 on page 396](#)). You can also export the table parameters to a CSV file for use as a backup or to import on another Connector Appliance.

**Note**

When you update parameters for connectors that are of different versions, the newer connectors might have additional parameters. In this case, only the parameters that are the same for all connectors are displayed for updating.

Managing Destinations

Connectors can forward events to more than one destination, such as ArcSight 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
- Re-register destinations
- Manage alternate configurations for a destination
- Send a command to a destination

**Note**

- You cannot configure two connectors with the same ESM Manager destination if the destination (connector) name and location used for configuration is the same.
 - Logger receivers do not support encrypted data.
-

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).

User Interface Options**Path**

From the location in which the connector exists


Click **System** (left panel) > *Location* (left panel) > **Connectors** tab (right panel) > *Name of the Connector* (right panel).

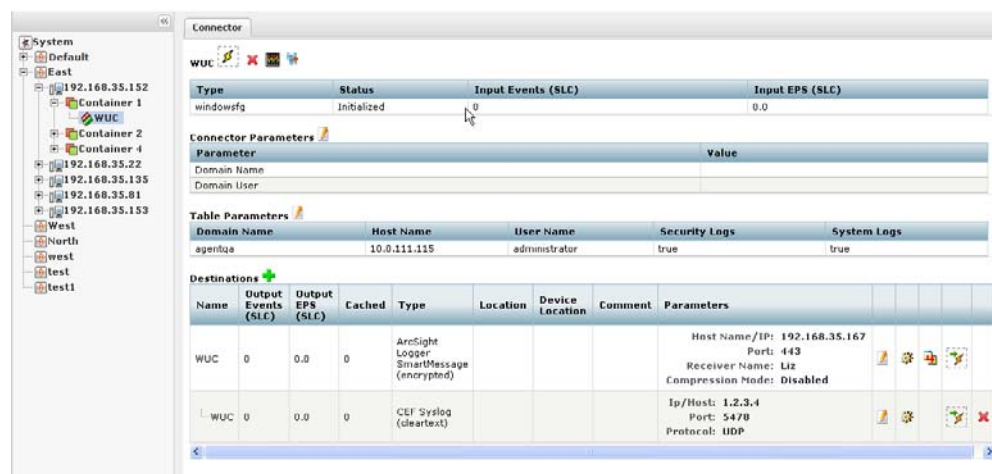
From the host on which the connector exists

Click **System** (left panel) > *Location* (left panel) > *Host* (left panel) > **Connectors** tab (right panel) > *Name of the Connector* (right panel).

From the Connector page

Click **System** (left panel) > *Location* (left panel) > *Host* (left panel) > *Container* > *Name of the Connector* (left panel).

3 Click () next to the **Destinations** link.



Note

Clicking the **Destinations** heading toggles between displaying and hiding the information in the Destinations section.

4 Follow the steps in the wizard.

You can either select an existing destination or add a new destination. If you are adding a new destination, select the destination type and enter parameters for the destination.



For containers running v5.1.2.5823 and later, Connector Appliance retrieves the certificate for the destination automatically and displays the certificate summary. To see certificate details, hover your mouse over the certificate.

- Select **Import the certificate to the connector from the destination**, then click **Next** to import the certificate and continue.
- Select **Do not import the certificate to the connector from the destination** and click **Next** if you do not want to import the certificate. The destination will not be added.

For containers running v5.1.2 and earlier, upload the certificate on the container and then add the destination.

Note: FIPS Suite B mode is not supported. Connector Appliance cannot download a manager certificate in Suite B mode.

- 5 Click **Done** when complete.

Adding a Failover Destination to a Specific Connector

Each destination can have a failover destination that is used if the connection with the primary destination fails.




UDP connections cannot detect transmission failure; use Raw TCP for CEF Syslog destinations.

To add a failover destination:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click () in the Destinations section to display the Add Connector Destination wizard.
- 4 Follow the steps in the wizard to select from available destinations and enter the destination details.



For containers running v5.1.2.5823 and later, Connector Appliance retrieves the certificate for the destination automatically and displays the certificate summary. To see certificate details, hover your mouse over the certificate.

- Select **Import the certificate to the connector from the destination**, then click **Next** to import the certificate and continue.
- Select **Do not import the certificate to the connector from the destination** and click **Next** if you do not want to import the certificate. The destination will not be added.

For containers running v5.1.2 and earlier, upload the certificate on the container and then add the destination.

Note: FIPS Suite B mode is not supported. Connector Appliance cannot download a manager certificate in Suite B mode.

Adding a Primary or Failover Destination to Multiple Connectors

You can add a primary or failover destination to several connectors at the same time.

To add a primary or failover destination to more than one connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel).
From the Connectors page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- 3 Select all connectors to which you want to assign a destination.
- 4 Click **Add Destinations** at the bottom of the page to open the wizard.
- 5 Select **Add a destination** and click **Next**.
- 6 Choose between a creating a new destination or selecting an existing destination, then click **Next**.

If you choose to **create a new destination**, select the destination type and then provide the destination parameters.

If you choose to **select an existing destination**, select a destination from the list.



Note

Connector Appliance retrieves the certificate for the destination automatically and displays the certificate summary. To see certificate details, hover your mouse over the certificate.

- Select **Import the certificate to the connector from destination**, then click **Next** to import the certificate and continue.
- Select **Do not import the certificate to the connector from the destination** and click **Next** if you do not want to import the certificate. The destination will not be added.

Note: FIPS Suite B mode is not supported. Connector Appliance cannot download a manager certificate in Suite B mode.

- 7 Define the destination function by choosing between a primary or failover destination.

If you choose **Primary destination**, click **Next** to update the configuration.

If you choose **Failover destination**:

 - 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.

8 Click **Done** when complete.

Removing Destinations

You can remove a destination from a connector at any time. The following procedures describe how to remove a single destination from a specific connector and how to remove multiple destinations from one or more connector.


To remove a single destination from a *specific* connector:

- 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.

To remove *multiple* destinations from one or more connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel).

User Interface Options	Path
From the Connectors page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- 3 Select the connectors whose destinations you want to remove.
- 4 Click the **Destinations** button to open the wizard.
- 5 Select **Remove destinations** and click **Next**.
- 6 Follow the instructions in the wizard and click **Done** when complete.

Re-Registering Destinations

At certain times, you might need to re-register the destinations for one or more connector; for example, after you upgrade ESM, or if a Logger appliance or ESM appliance becomes unresponsive.

To re-register destinations for one or more connector:

- 1 Click **Configuration** > **Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel).
From the Connectors page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- 3 Select the connectors whose destinations you want to re-register.
- 4 Click the **Destinations** button to open the wizard.
- 5 Select **Re-register destinations** and click **Next**.
- 6 Follow the instructions in the wizard and click **Done** when complete.

Editing Destination Parameters

The following procedures describe how to edit destination parameters for a specific connector and how to edit destination parameters for multiple connectors at the same time.



You cannot change the connector type. However, you can remove the unwanted connector configuration and create a new one.

To edit destination parameters for a specific connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 In the Destinations section, click () next to the destination you want to edit to display the Edit Destination Parameters page.

The screenshot shows the 'Connector' configuration page. On the left is a tree view of the system hierarchy. The main area is divided into sections: 'Connector Parameters', 'Table Parameters', and 'Destinations'. The 'Destinations' section is expanded, showing 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 selected, and its parameters are shown in a detailed view. A red circle highlights the 'Edit' button (pencil icon) next to the 'WUC' destination in the Destinations table.

- 4 Make your changes and click **Next**.
- 5 Click **Done** when complete.

To edit destination parameters for *multiple* connectors:

- 1** Click **Configuration** > **Manage Connectors**.
- 2** Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel).
From the Connectors page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- 3** Select the connectors whose destination parameters you want to edit.
- 4** Click **Destinations** to open the wizard.
- 5** Select **Edit a destination** and click **Next**.
- 6** Follow the instructions in the wizard and click **Done** when complete.

Editing Destination Runtime Parameters



The runtime parameters for a destination enable you to specify advanced processing options such as batching, time correction, and bandwidth control. The parameters you can configure are listed in [Appendix H, Destination Runtime Parameters, on page 511](#). 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 409](#).

- 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 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

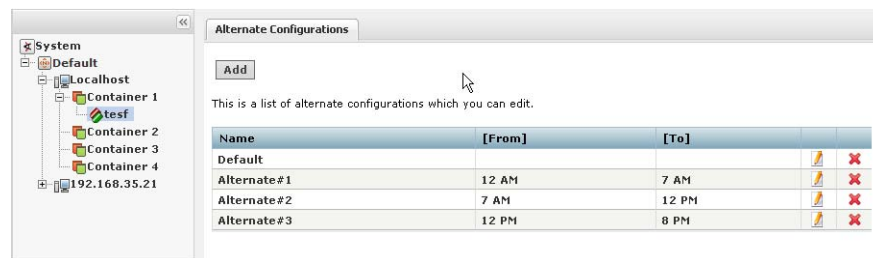
An alternate configuration is a set of runtime parameters that is used instead of the default configuration during a specified portion of every day. For example, you might want to specify different batching schemes (by severity or size) for different times of a day. You can define more than one alternate 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 410](#).

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 407](#).


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

**Note**

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.

**Note**

You can also delete a specific connector from its details page: Click **System** (left panel) > **Location** (left panel) > **Host** (left panel) > **Container** > **Connector** >  at the top of the page.

Sending a Command to a Connector


You can send a command to a connector.


To send a command to a connector:

1 Click **Configuration > Manage Connectors**.

2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

3 Click  in the Action column for the connector.

If you are on a specific Connector page,  is on top of the page.

4 From the **Command Type** drop-down list, select the command you want to send to the connector.

5 Click **Next**.

Running Logfu on a Connector

Run Logfu on a connector to parse ArcSight logs and generate an interactive visual representation of the information contained within the logs.

To run Logfu on a connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click () on top of the page. A separate window displays.

The system proceeds to retrieve and analyze system data logs. After this process is complete, a group of panels appears in the window.

- 4 From the **Group** box, choose which type of data you would like to view. The Group box lists all connectors within the chosen container, plus many other types of data such as memory usage, and transport rates and logs.

Choose one of the Group box **data points**. Depending on which data point you choose, a list of fields appears in the Field box below.

- 5 Choose a **field** to view. A graphic chart appears in the Chart box, providing rate and time information. The key at the bottom of the Chart box defines the data points mapped in the chart.
- 6 If you need to choose a different data point for analysis, click **Reset Data**.

Changing the Network Interface Address for Events

Connector Appliance has multiple network interfaces. By default, the connector determines which network interface address is used for events displayed in the ESM console or Logger, but typically uses `eth0`.

To use a specific network interface address for events, add the parameter `connector.network.interface.name` to the Connector's `agent.properties` file. For example, to use the IP address for `eth1`, specify the following parameter:

```
connector.network.interface.name=eth1
```

Developing FlexConnectors

FlexConnectors are custom SmartConnectors that can read and parse information from third-party devices and map that information to ArcSight's event schema.

Connector Appliance provides a FlexConnector Development wizard that lets you quickly and easily develop a FlexConnector by creating a parser file, and enables you to test and package your new FlexConnector before deploying it. The wizard generates regular expressions and provides event field mapping suggestions automatically so you do not need to be an expert in regular expression authoring, parser syntax, or ArcSight event schema.

Use the FlexConnector Development wizard to develop FlexConnectors for simple log files. For complex log files, use the FlexConnector SDK (available from the ArcSight Customer Support site).



Currently, the FlexConnector Development wizard supports Regex Files, Folder Follower, and Syslog (Daemon, File, Pipe) FlexConnectors only.

The FlexConnector Development wizard does not support the extra processors property or multiple sub messages. If you need these features, use the FlexConnector SDK to create your FlexConnector.




A FlexConnector that you develop with the FlexConnector Development wizard might perform more slowly than an ArcSight SmartConnector.

To develop a FlexConnector:

- 1 Click **Configuration** > **Manage Connectors**.
- 2 Use one of these navigation paths to go to the **Containers** tab:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Click  in the Action column of the container to which you want to add the FlexConnector. When the FlexConnector Development wizard opens, click **Next**.
- 4 Provide the vendor and product name of the device for which you are creating a FlexConnector, then click **Next**.



The device vendor and product name are required.

- 5 Select the data source type, then click **Next**:
 - ◆ Select **Syslog** to create a Syslog FlexConnector to read events from Syslog messages.
 - ◆ Select **File** to create a FlexConnector to parse variable-format log files using regular expressions (ArcSight FlexConnector Regex File) or to parse variable-format log files in batch mode (ArcSight FlexConnector Folder Follower).
- 6 Upload a sample log file for the data source type you selected in the previous step, then click **Next**.
- 7 The wizard finds the first unparsed line in the log file, generates a regular expression to match and extract tokens from that line, and displays the suggested field mappings for each extracted token in the Mappings table.

Wizard

FlexConnector Development Wizard

Enter regular expression corresponding to text
 Text 2005 Aug 24 13:57:54 EDT -04:00 %SPANTREE-6-PORTFWD: Port 3/16 state in VLAN 203 changed to forwarding
 Lines Skipped: 0% Lines Parsed: 0%

Regex (\\d+ \\S+ \\d+ \\d+:\\d+:\\d+ \\S+ \\S+) %SPANTREE-6-PORTFWD: Port (\\S+?) state in VLAN (\\d+) changed to forwarding
 Recalculate Reset

Mappings table			
Extracted Value	Type	Format	Event Field
1 2005 Aug 24 13:57:54	TimeStamp	yyyy MMM dd HH:mm:	deviceReceiptTime
2 3/16	String	String	deviceInboundInterface
3 203	Integer	String	deviceInboundInterface

Extra Mappings table	
Event Field	Value
name	__stringConstant(SPAN)

Add Row

Cancel Skip Line Skip To End Previous Next



The mappings are displayed in descending order of probability (based on ArcSight training data). You can change the mappings by selecting from the list.

The percentage of parsed lines in the file is shown in the top right of the panel. You can use this percentage to estimate where you are in the log file. The percentage of unparsed lines skipped in the file is also shown in the top right of the panel.

- ◆ To change the regular expression in the **Regex** box and recalculate the mappings, edit the expression and then click the **Recalculate** button. For information about regular expressions, see [Appendix B, Regular Expressions, on page 461](#). You can set the regular expression back to the suggested value by clicking the **Reset** button.
- ◆ Field mappings that do not correspond directly to the extracted tokens in the unparsed line of the log file are displayed in the Extra Mappings table. You can change the Event Field and provide a token operation. To add a new Event Field, click **Add Row**.

You can use extra mappings to:

- Remap an extracted token to a different Event Field in addition to the existing mapping. For example, you can add an Event Field with the value `$3` where `$3` is the third token in the list of suggested mappings.
- Map a modified token or combination of tokens to an Event Field. For example, you can add an Event Field with the value `__operation($1,$3)`.

- Map an Event Field to a constant string or integer. For example, you can add an Event Field with the value `__stringConstant(constant)`.

**Note**

The wizard always contains an extra mapping for the Event Field **name**, which maps all the words in the input log line. ArcSight strongly recommends that you do not simply delete the **name** Event Field but map it in either the Mappings or the Extra Mappings table.

For a list of the token operations used when tokens are mapped to ArcSight event fields, refer to the *FlexConnector Developer's Guide* (available from the ArcSight Customer Support site).

- 8 Click **Next** to save the mapping to the parser file and display the next unparsed line in the log file.

**Tip**

Click the **Skip Line** button to go to the next unparsed line in the log file without saving the mapping.

Click the **Skip to End** button to go to the end of the log file without processing any other lines and display the parser file for review.

Click the **Previous** button to go back to the previous line in the log file and make changes if necessary. If you configured any mappings for the previous line, the **Previous** button displays the configured mappings, not the default mappings.

After all unparsed lines in the log file have corresponding regular expressions and mappings, the wizard displays the parser file for review.

- 9 Review the parser file and make changes, if necessary, directly in the Review Parser File panel.

**Note**

In Mozilla Firefox, if certain text in the Review Parser File panel is underlined in red, you can disable Spell Check; Right-click in the panel and click **Check Spelling** to remove the check mark.

- 10 Click **Next** to save and package the parser file.

11 Choose how you want to deploy the FlexConnector:

- ◆ Select **Deploy parser to existing connector in container** and click **Next** to use the parser file with an existing connector. Click **Done** to close the FlexConnector wizard and redisplay the Container tab.



The **Deploy parser to existing connector in container** option displays only if the container already contains a connector of the same type.

- ◆ Select **Add new connector to container** and click **Next** to add the parser as a new connector. Follow the steps to add the connector to the container.




After deploying your FlexConnector, you can edit it any time from the **Connectors** tab. See [“Editing FlexConnectors” on page 418](#).


You can share FlexConnectors with other users. See [“Sharing Connectors \(ArcExchange\)” on page 419](#).

Editing FlexConnectors

After you have developed a FlexConnector with the FlexConnector wizard and have deployed it in a container, you can edit the FlexConnector to make changes to the parser file when needed.

The FlexConnector Edit wizard is available on the **Connectors** tab in the **Action** column.



Click  in the **Action** column for the FlexConnector to open the wizard. To edit the parser file, follow [Step 6](#) through [Step 11](#) in [“Developing FlexConnectors” on page 415](#).



Only edit a FlexConnector that is created with the FlexConnector wizard. Editing manually-created FlexConnectors might produce unpredictable results.



In addition to the FlexConnector Edit wizard, you can also use the Edit a File action in the Container Diagnostics wizard to edit your FlexConnector. Refer to [“Running Diagnostics on a Container” on page 389](#).

Sharing Connectors (ArcExchange)

You can share FlexConnectors and parser overrides with other users.

A FlexConnector is a custom connector that you define to gather security events from log files, databases, and other software and devices. You can share the following FlexConnector types:

- Syslog FlexConnectors (to read events from syslog messages)
- Log File FlexConnectors (to read fixed-format log files)
- Regular Expression Log File FlexConnectors (to read variable-format log files)
- Regular Expression Folder Follower FlexConnectors (to read variable-format log files recursively in a folder)
- Regular Expression Multiple Folder Follower FlexConnectors (to read events in real time or batch mode from multiple folders)
- XML FlexConnectors (to read events recursively from XML-based files in a folder)

A parser override is a file provided by ArcSight used to resolve an issue with the parser for a specific connector, or to support a newer version of a supported device where the log file format changed slightly or new event types were added. You can share parser overrides for all connector types that use a parser.

To share a FlexConnector or parser override, you need to package and upload it to ArcExchange on the ArcSight online community (Protect 724) or to your local machine. You can also download a FlexConnector or parser override that you need from ArcExchange or from your local machine and add it to a container.

Packaging and Uploading Connectors

Before uploading your FlexConnector or parser override to Protect 724 or to your local computer, you need to package it into a zip file, (called an AUP package) using the upload wizard.

A FlexConnector AUP package contains the connector properties file, categorization file, connector parameters, and a manifest file with all the metadata on the package required for successful deployment. Metadata includes information about the AUP package, such as the package type, connector type, connector description, and so on. You can create only one AUP package per connector per device type. You can package a FlexConnector in Basic or Advanced mode. In **Basic** mode:

- The wizard packages the FlexConnector properties file automatically. If the wizard finds more than one properties file, you are prompted to select the file you want to package.
- The wizard packages the categorization file automatically *only* if it can be determined based on the device vendor and product information found in the properties file.
- The wizard does not package connector parameters. You are prompted to configure the connector when it is downloaded and deployed.

In **Advanced** mode:

- The wizard packages the FlexConnector properties file automatically. If the wizard finds more than one properties file, you are prompted to select the file you want to package. (This is same as Basic mode.)
- The wizard packages the categorization file automatically if it can be determined based on the device vendor and product information found in the properties file. If the categorization file cannot be determined, you are prompted to select the categorization file you want to package from the list of files found in the container.
- The wizard displays connector parameters so you can configure the parameters you want to display and set the default values you want to provide during connector deployment (download). The parameters you do not configure for display are preconfigured with the current values and will not be displayed during connector deployment.

A parser override package contains the parser override properties file and the manifest file only.

Follow the steps below to package and upload a FlexConnector or parser override.





- To upload to ArcExchange, you must have a valid username and password for Protect 724.
- Make sure that you have configured Connector Appliance network settings under Setup > System Admin > Network and that the appliance can communicate with the Protect 724 server.

To package and upload a FlexConnector or parser override:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click  at the top of the Connector page to open the upload wizard. (From the Connectors page, select the connector in the right panel and click  in the **Action** column.)

4 Click **Next** and follow the steps in the wizard to:

- a** Select the type of AUP package you want to create for the selected connector.
Connector Appliance scans the container and displays the relevant files that can be packaged.
- b** For a FlexConnector, select **Basic** to create a default package or select **Advanced** to customize the package to meet your needs. For a description of Basic and Advanced mode, refer to [“Packaging and Uploading Connectors” on page 419](#).
- c** If the connector contains several properties files, you are prompted to select the properties file you want to package. Certain connectors, for example, syslog connectors, can have more than one parser override folder, in this case, you are prompted to select the folder you want to package.
- d** If you selected Advanced mode for a FlexConnector in [Step b](#) and the categorization file cannot be determined, you are prompted to select the categorization file you want to package from a list of files found in the container.



Categorization files are not packaged for parser overrides.

- e** If you selected Advanced mode for a FlexConnector in [Step b](#), select the configuration parameters you want to display when the connector is deployed and then provide default values for these parameters. Parameters you do not select are pre-configured with the current values.

If any advanced connector parameters were previously modified from their defaults, the wizard displays these parameters so that you can select which ones you want to be configured automatically during deployment.



Configuration parameters are not displayed for parser overrides.

If the connector has table parameters, they are not displayed during packaging. However, when the connector is downloaded to a container, you will be prompted to provide values for all the table parameters.

- f** Provide a description of the AUP package and instructions on how configure the device used by the connector.
- g** Provide the vendor, product, and version of the device used by the connector.
If the wizard can determine the vendor, product, and version of the device, the information is displayed in the fields provided. You can change the information to meet your needs.
- h** Upload the created AUP package to ArcExchange or to your local machine.



To upload the AUP package to ArcExchange, you must have a valid username and password for Protect 724.

Downloading Connectors

You can download a FlexConnector or parser override that is available from ArcExchange on Protect 724 or from your local computer. You download a FlexConnector or parser override directly to a container.

You can download only one FlexConnector per container using the download wizard. However, there is no limit to the number of parser overrides you can download to a container.



- When downloading a parser override to a container, the download wizard overwrites any existing parser override with the same name in the container without prompting for confirmation. To avoid overwriting an existing parser override, send a **Get Status** command to the existing parser override to check the parser information before you download a new parser override. For information on sending a Get Status command, refer to [“Sending a Command to a Connector” on page 413](#).
- ArcSight recommends that you back up the container to the Backup Files repository before downloading a connector or parser override so you can revert to the previous configuration if the download produces unexpected results.


Follow the steps below to download a FlexConnector or parser override to a container.

To download to ArcExchange, you must have a valid username and password for Protect 724. Also, make sure that you have configured Connector Appliance network settings under Setup > System Admin > Network and that the appliance can communicate with the Protect 724 server.

To download a FlexConnector or parser override:

- 1 Click **Configuration > Manage Connectors**.
- 2 Go to the **Containers** page. Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 In the right panel, select the container into which you want to download the connector, and then click  in the **Action** column to open the download wizard.

4 Click **Next** and follow the steps in the wizard to:

- a Select whether you want to download the connector from ArcExchange on Protect 724 or from your local computer.
- b Select the AUP package you want to download.

On Protect 724, you can search for a parser override or FlexConnector AUP package using a keyword or a combination of keywords.



You can only download a parser override package to a container that has a connector of the same type as the package.

You can download only one FlexConnector per container using the download wizard. If the container already contains a FlexConnector of the same type as the one you want to download, you can replace the existing FlexConnector with the one you are downloading, but you cannot create a new one.

- c For a FlexConnector, provide connector configuration parameters, if needed.

Preconfigured and advanced parameters are deployed automatically with the values that were packaged; you are not prompted to configure these parameters. The configurable parameters are displayed with suggested defaults, which you can modify if necessary. The table parameters are displayed with no configured values, you have to provide the values manually, as needed.

- d Add or select a destination for the connector.

If you are downloading the connector to a container that has an existing connector of the same type, you are *not* prompted for a destination.

The wizard copies the properties and categorization files to the appropriate locations and also installs the zip file for the AUP package in the `user/agent/deployedaups` folder on the Connector Appliance to keep track of the deployment history.

After a successful download, the container is restarted automatically.



To use memory efficiently, parser overrides for the Windows Unified connector only load when the first event is received.

Configuration Suggestions for Connector Types

The following table provides configuration suggestions for different types of connectors.

Connector Type	Effects of Limited Usage
Syslog connectors	<p>Due to the nature of UDP (the transport protocol typically used by Syslog), these connectors can potentially lose events if the configurable event rate is exceeded. This is because the connector delays processing to match the event rate configured, and while in this state, the UDP cache might fill and the operating system drop UDP messages.</p> <p>Note: ArcSight recommends that you do not use the Limit CPU Usage option with these connectors because of the possibility of event loss.</p>
SNMP connectors	<p>Similar to Syslog connectors, when the event rate is limited on SNMP connectors, they potentially lose events. SNMP is also typically UDP-based and has the same issues as Syslog.</p>
Database connectors	<p>Because connectors follow the database tables, limiting the event rate for database connectors can slow the operation of other connectors. The result can be an event backlog sufficient to delay the reporting of alerts by as much as minutes or hours. However, no events will be lost, unless the database tables are truncated. After the event burst is over, the connector might eventually catch up with the database if the event rate does not exceed the configured limit.</p>
File connectors	<p>Similar to database connectors, file-based connectors <i>follow</i> files and limiting their event rates causes an event backlog. This can eventually force the connector to fall behind by as much as minutes or hours, depending on the actual event rate. The connectors might catch up if the event rate does not exceed the configured rate.</p>
Asset Scanner connectors	<p>All connectors on Connector Appliance run as a service (not as an application). Therefore, asset scanner connectors running on Connector Appliance are <i>not</i> supported in Interactive mode.</p> <p>To run the asset scanner connector in Interactive mode, install the connector on a standalone system and manage it as a software-based connector.</p>
Proprietary API connectors	<p>The behavior of these connectors depends on the particular API, (for example, OPSEC behaves differently than PostOffice and RDEP). But in most cases, there will be no event loss unless the internal buffers and queues of the API implementation fill up. These connectors work much like database or file connectors.</p>

Deploying FlexConnectors

FlexConnectors are custom connectors that are user-defined. 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. You can share FlexConnectors with other Connector Appliance users. Refer to [“Sharing Connectors \(ArcExchange\)” on page 419](#).

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 connector can operate in clear channel or sslca mode.



Note

- This procedure is supported only for ArcSight connector release 4.6.2 or later.
- A hostname is called an Application Object Name on Check Point. A password is a Communication Activation Key on Check Point.

To configure a connector to operate in sslca mode

On the Check Point SmartDashboard:

- 1 Create an OPSEC Application Object using the Check Point SmartDashboard. You need to provide these parameters when creating the application object.


Parameter	Description
Name	A meaningful name for the application object you are creating; for example, ArcSightLea-1. This name is used to pull the OPSEC certificate in the system.
Host	The hostname of the system managing the connector.
Client Entities	Select LEA.

Parameter	Description
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 385](#). 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.pl2` in the above example).



If the certificate is not pulled successfully, check to ensure that the Application object name you specified is correct (including the case) and the container on which you are running the command is up and running.

- 3 Install Policy on the LEA client for the Check Point Gateway using the SmartDashboard.

On the Connector Appliance:

- 4 Add a Check Point connector by following instructions described in [“Adding a Connector” on page 390](#). 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 390](#).

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 444](#) 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 446](#).

After the driver file has been uploaded to a container, follow the instructions in [“Adding a Connector” on page 390](#) to add a connector that requires a JDBC driver.

Managing Repositories

The following topics are discussed here.

- [“Overview” on page 430](#)
- [“Logs Repository” on page 432](#)
- [“CA Certs Repository” on page 433](#)
- [“Upgrade AUP Repository” on page 436](#)
- [“Content AUP Repository” on page 437](#)
- [“Remote Management AUP Repository” on page 439](#)
- [“Emergency Restore” on page 441](#)
- [“User-Defined Repositories” on page 442](#)
- [“Pre-Defined Repositories” on page 447](#)

Overview

Certain management operations require a specific upgrade or content update (.enc) file, or a certificate. Other operations such as viewing the logs require you to load the logs to a Log repository. You can also maintain centralized repositories for files needed for connector configuration and management.

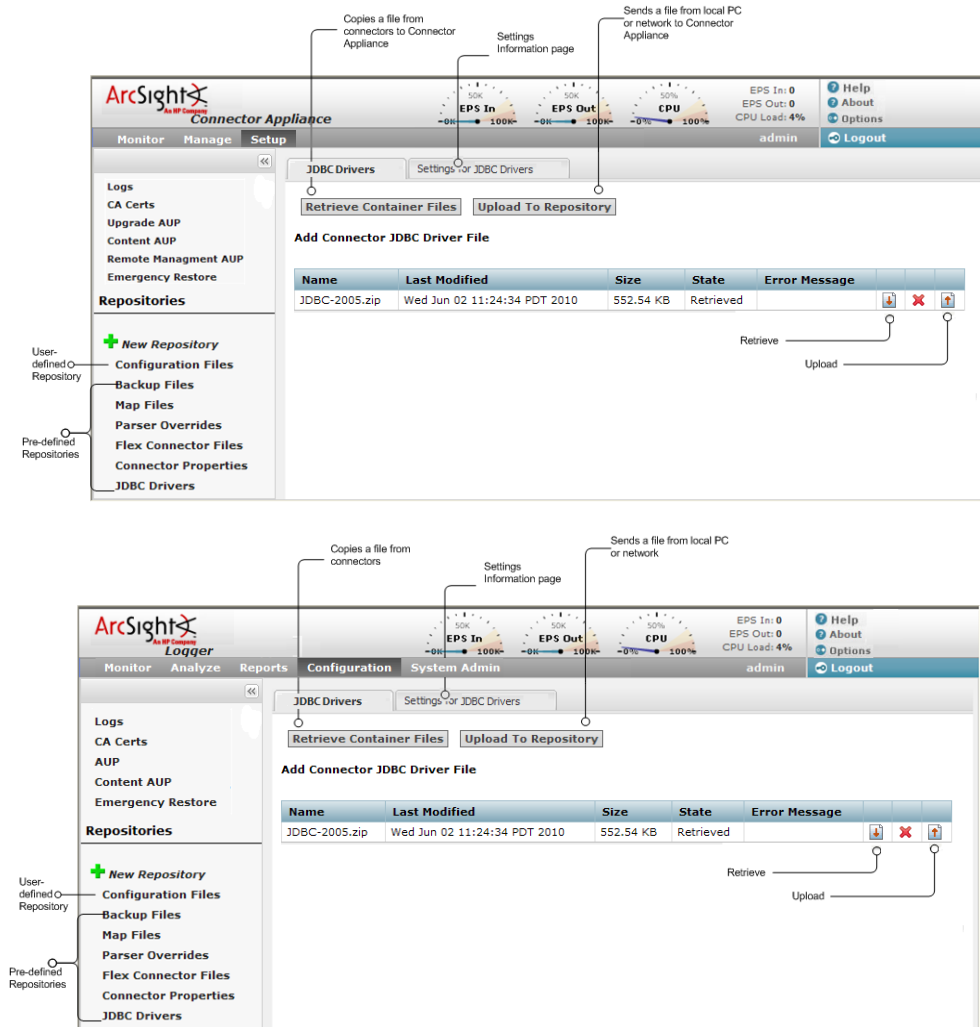





Figure 9-1 Repository Functions

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 connectors to the repository.
- **Upload to Repository** sends a file from your local computer (the computer running the browser) or a network host accessible from your local computer to the repository.
- **Retrieve**  downloads a file from the repository to your local computer network.

- **Upload**  copies a file from the repository to one or more connectors.



You can perform these operations using repositories:

- Manage logs in the Logs repository
- Manage CA certificates in the CA Certs repository
- Upgrade a connector using an upgrade file available in the Upgrade repository
- Apply a Content ArcSight Update Pack (AUP) on one or more connector
- Manage remote management configuration AUP files in the Remote Management AUP repository
- Restore a container when it is damaged and irrecoverable
- Maintain centralized repositories of files for connector configuration and management

Logs Repository

When you want to view connector logs, you need to first **Load** the logs of the container that contains the connector to the Logs repository, then **Retrieve** the logs to view them.



If a container contains more than one connector, logs for all connectors are retrieved.

For information on loading, retrieving, and deleting the logs, see [“Viewing Container Logs” on page 387](#).

Uploading a File to the Logs Repository

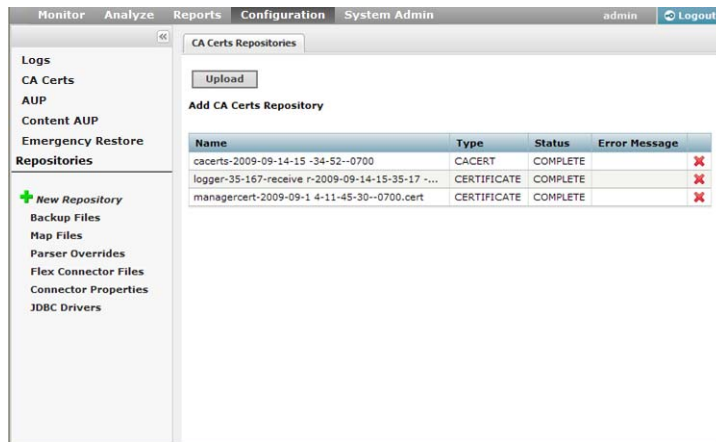
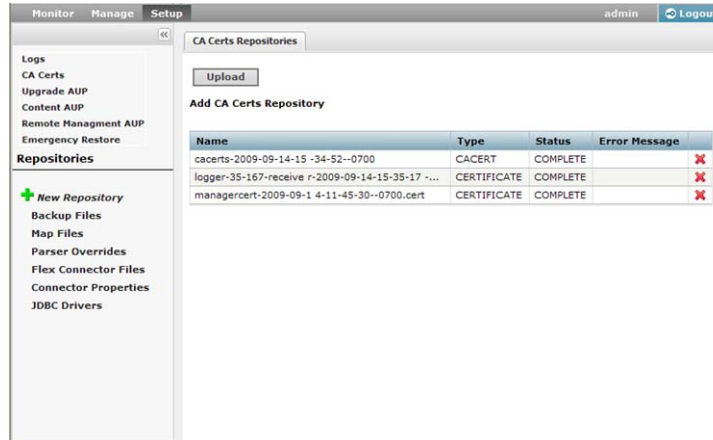
Uploading a file into the Log repository is useful for sharing annotated log or other files with other users. The file needs to be in `.zip` format.

To upload a file:

- 1 Click **SetupConfiguration** > **Repositories**.
- 2 Click **Logs** from the left panel.
- 3 Click **Upload** from the right panel.
- 4 Enter the local file path or click **Browse** to select the file.
- 5 Click **Submit** to add the specified file to the repository or **Cancel** to quit.

CA Certs Repository

Connectors require a Certificate Authority (CA) issued or self-signed SSL certificate to communicate securely with a destination. The CA Certs repository (shown below) enables you to store CA Certs files (that contain one or multiple certificates) and single CA certificates. When certificates are stored in the CA Certs repository, you can add the certificates to a container so that the connectors in the container can validate their configured destinations successfully.



To associate a CA certificate to a connector, you need to:

- Upload the CA certificate or CA Certs file to the CA Certs repository, as described below.
- Add a CA certificate from the CA Certs repository to the container that contains the connector, as described in [“Managing Certificates on a Container” on page 379](#).



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

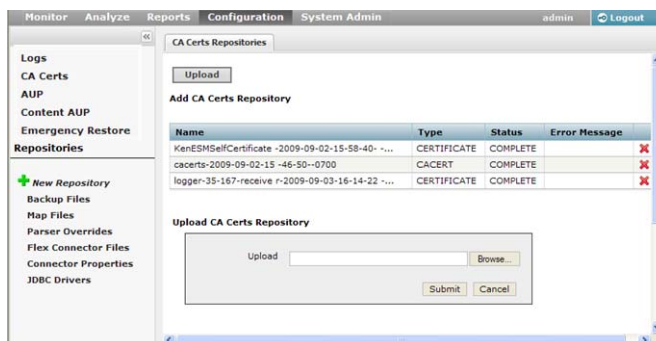
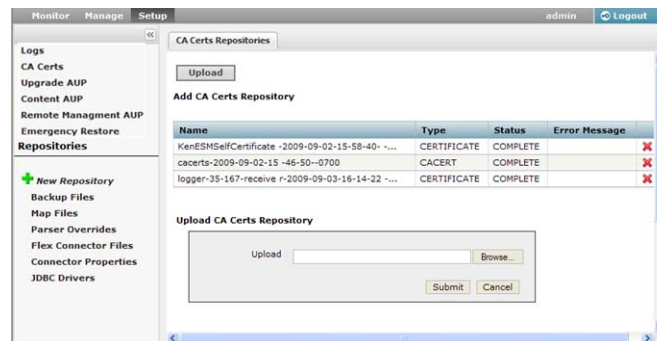
You can upload a CA Certs file or a single certificate to the CA Certs repository.



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 **SetupConfiguration** > **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 379](#).

To remove a certificate from the repository:

- 1 Click **SetupConfiguration** > **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 (✖).

Upgrade AUP Repository

The Upgrade AUP repository enables you to maintain a number of connector AUP (upgrade) files. You can apply any of these AUP upgrade files to containers when you need to upgrade to a specific version. As a result, all connectors in a container are upgraded to the version you apply to the container.

This repository can also maintain upgrade files for upgrading remotely-managed Connector Appliances. **The central appliance needs to be upgraded using the .enc file before you use it to upgrade other appliances remotely.**

About the AUP Upgrade Process



The process discussed in this section only applies to upgrading connectors and to upgrading a remotely-managed Connector Appliance. **If you are upgrading the local Connector Appliance (localhost), use a .enc file. Refer to the Release Notes for more information.**

To upgrade a connector or to upgrade a remotely-managed Connector Appliance, you need to:

- Upload the appropriate .aup upgrade file to the Upgrade AUP repository, as described below.
- Apply the .aup upgrade file from the Upgrade AUP repository to the container (see “Upgrading a Container to a Specific Connector Version” on page 386) or to a remote Connector Appliance (see “Upgrading a Host Remotely” on page 372).

Uploading an AUP Upgrade File to the Repository


To upload AUP upgrade files to the repository:

- 1 Download the upgrade AUP file for the connector or the remote Connector Appliance from the ArcSight Customer Support site at <http://www.arcsight.com/supportportal> 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 **SetupConfiguration** > **Repositories** from the top-level menu bar.
- 4 Click **Upgrade 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 386.
 - ◆ For a remotely-managed Connector Appliance upgrade, see “Upgrading a Host Remotely” on page 372.

Removing a Connector Upgrade from the Repository

You can remove a connector upgrade file from the repository when you no longer need it. When you remove a connector upgrade file from the repository, it is deleted from the system.

To remove a Connector upgrade from the repository:

- 1 Click **SetupConfiguration** > **Repositories** from the top-level menu bar.
- 2 Click **Upgrade 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 connector event categorization mappings, often called *content*. This content is packaged in ArcSight Update Packs (AUP) files. All existing content is included with major product releases, but it is possible to stay completely current by receiving up-to-date, regular content updates through ArcSight announcements and the Customer Support site. The AUP files are located under Content Subscription Downloads.

The ArcSight Content AUP feature enables you to apply an AUP file to applicable connector destinations that you are managing. Only the event categorization information can be applied to the connectors using this feature.

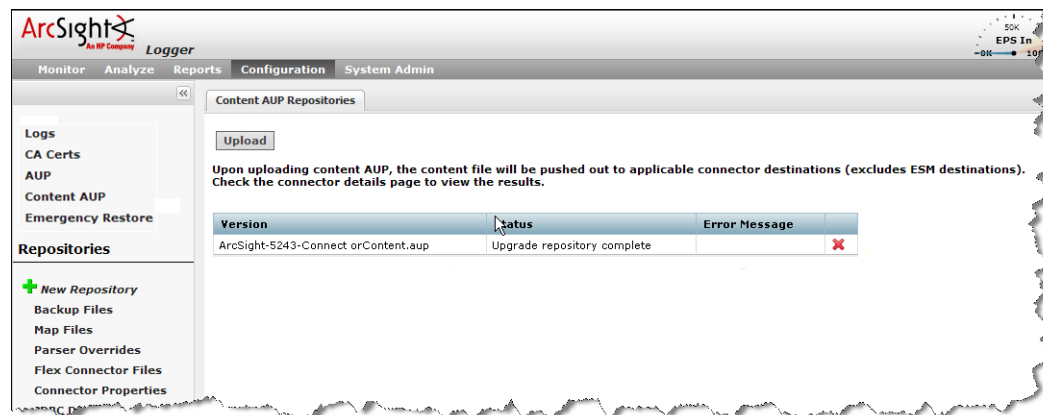
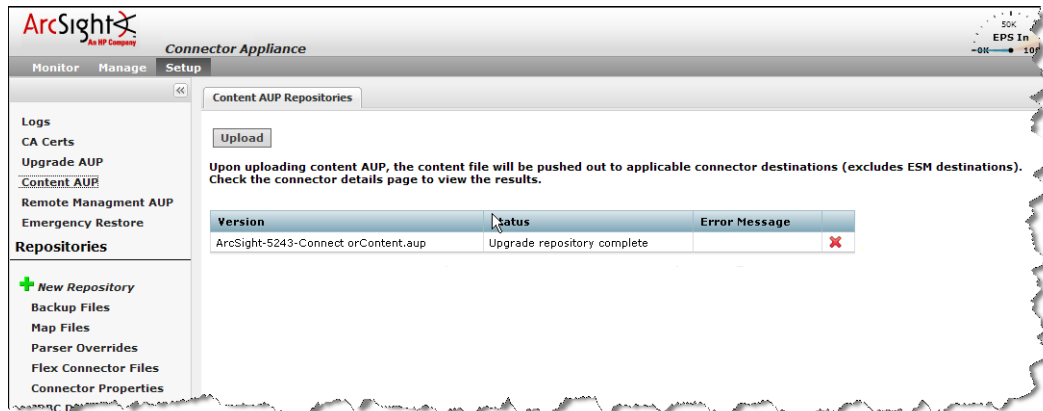
You can maintain a number of Content AUP files in the Content AUP repository. When an AUP file with a version number higher than the ones already in the repository is loaded, it is automatically pushed out to the connector destinations being managed. However, these connectors or connector destinations are skipped:

- Connectors that are unavailable at the time of the AUP file push
- Connectors whose current version does not fall in the range of versions that the Content AUP supports
- The ESM destination on a connector
- All destinations of a connector that have an ESM destination with the AUP Master flag set to Yes

Also, when a new connector is added, the highest number Content AUP is pushed automatically to its destinations.

Applying a New Content AUP

You can add a new content AUP file to the repository and push it automatically to all applicable connectors



To apply a new Content AUP:

- 1 Download the new Content AUP version from ArcSight Customer Support site at <http://www.arcsight.com/supportportal> 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 **SetupConfiguration** > **Repositories** from the top-level menu bar.
- 4 Click **Content AUP** from the left panel.
- 5 Click **Upload** from the right panel.
- 6 Click **Browse** and select the file you downloaded earlier.
- 7 Click **Submit** to add the specified file to the repository and push it automatically to all applicable connectors, or **Cancel** to quit.

You can verify the current Content AUP version on a connector by performing either of these steps:

- Run the `GetStatus` command on the connector destination and check that the value for `aup[acp].version` is the same as the AUP version you applied. For information


about running a command on a connector destination, see [“Sending a Command to a Destination” on page 411](#).

- Hover your mouse over a connector name to see the AUP version applied to all destinations of that connector.

Applying an Older Content AUP

If you need to apply an older Content AUP from the Content AUP repository, delete all versions newer than the one you want to apply in the repository. The latest version (of the remaining AUP files) is pushed automatically to all applicable connectors.

To delete a Content AUP from the Content AUP repository:

- 1 Click **SetupConfiguration** > **Repositories** from the top-level menu bar.
- 2 Click **Content AUP** from the left panel.
- 3 Locate the AUP file that you want to delete and click the associated  icon. Repeat for multiple files.







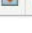

Remote Management AUP Repository

The Remote Management AUP repository stores AUP files that contain the remote management configuration of an appliance (a snapshot of all the remote software connectors and remote Connector Appliances that the appliance manages).

From the Remote Management AUP repository, you can:

- Download a Remote Management AUP file to your local computer (or network host accessible from the local computer) so that you can import the remote management configuration on another appliance.
- Upload Remote Management AUP files from your local computer (or network host accessible from the local computer) to the repository for storage.
- Delete Remote Management AUP files you no longer need.

The following example shows the Remote Management AUP repository.


Remote Management Repositories						
<input type="button" value="Upload"/>						
Add Remote Managment Repository						
Name	Last Modified	Size	State	Error Message		
RemoteManagementConfi g-5.6-C5606-127204...	Fri Apr 23 11:25:17 PDT 2010	1.46 KB	Retrieved			
RemoteManagementConfi g-5.6-C5606-127206...	Fri Apr 23 17:05:48 PDT 2010	1.3 KB	Retrieved			
RemoteManagementConfi g-5.6-C5606-127232...	Mon Apr 26 15:52:17 PDT 2010	1.89 KB	Retrieved			
RemoteManagementConfi g-5.6-C5606-127232...	Mon Apr 26 16:19:37 PDT 2010	1.11 KB	Retrieved			

Downloading Remote Management AUP Files

After you export the remote management configuration of a Connector Appliance, you can download the AUP file that contains the configuration to your local computer (or network host accessible from the local computer) so that it can be imported on another appliance.

For information on exporting and importing the remote management configuration of an appliance, refer to [“Exporting and Importing Remote Management Configuration” on page 364](#).

To download a Remote Management AUP file to your local computer:

- 1 Click **SetupConfiguration** > **Repositories** from the top-level menu bar.
- 2 Click **Remote Management AUP** from the left panel.
- 3 Locate the AUP file in the table and click  next to the file to download it to your local computer.

Uploading Remote Management AUP Files

You can upload remote management AUP files to the Remote Management AUP repository for storage.


To upload a Remote Management AUP file to the repository:

- 1 Click **SetupConfiguration** > **Repositories** from the top-level menu bar.
- 2 Click **Remote Management AUP** from the left panel.
- 3 Click the **Upload** button at the top of the page.
- 4 Click **Browse** and select the file you want to upload from the local computer (or network host accessible from the local computer).
- 5 Click **Submit** to add the specified file to the repository.

Deleting Remote Management AUP Files

When a remote management AUP file is no longer up-to-date or needed, you can remove it from the repository.

To delete a Remote Management AUP file:

- 1 Click **SetupConfiguration** > **Repositories** from the top-level menu bar.
- 2 Click **Remote Management AUP** from the left panel.
- 3 Locate the AUP file that you want to delete and click the associated  icon. Repeat for multiple files.


Emergency Restore

The Container Restore wizard guides you through the process of restoring a modified container. This feature is supported only for connectors and containers on the local host.



ArcSight recommends that you use this process only when a container is severely damaged and is no longer available. The Emergency Restore process deletes all information about that container and renders it empty. The connector is restored to the AUP version that you select.

To restore a container:

- 1 Click **SetupConfiguration** > **Repositories** from the top-level menu bar.
- 2 Click **Emergency Restore** from the left panel.
- 3 Follow the instructions in the Container Restore wizard.
- 4 Re-import the SSL certificate for the container. On the **Manage** tab, click the container name in the left panel. On the **Connectors** tab in the right panel, click the  icon to run the Certificate Download wizard and import the valid certificate.

User-Defined Repositories

A *user-defined repository* is a user-named collection of settings that control upload and download of particular files from connectors to the repository. Each repository uses a specified path, relative to `$ARCSIGHT_HOME/user/agent`, for files to be uploaded or for locations to download files. ArcSight connectors use a standard directory structure, so map files, for example, are always found in `$ARCSIGHT_HOME/user/agent`, (that is, the root directory, `$ARCSIGHT_HOME`, of the connector installation) in a folder called `map/`.

After they are created, user-defined repositories are listed on the left-side menu, under the **New Repository** heading, and appear with the user-specified display name.

User-defined repositories are expected to be grouped by file type and purpose, such as log files, certificate files, or map files. Each user-defined repository has a name, a display name, and an item display name, which are defined under the **Settings** tab that appears for user- or pre-defined repositories (for details about pre-defined repositories, see ["Pre-Defined Repositories" on page 447](#)).

Files viewed in the user-defined repository can be bulk processed with specified connectors and can be exchanged with the user's browser host.

Creating a User-Defined Repository

You can create a new repository at any time.



The repository requires correct directory paths. Your file will be applied to the wrong directory if the entered path contains errors, such as extra spaces or incorrect spellings. You can verify your directory paths by accessing the `Directory.txt` file, which lists the directory structure for every entered path. View the `Directory.txt` file by accessing your container logs and finding the `Directory.txt` file.

To create a new user-defined repository:

- 1 Click **SetupConfiguration** > **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.

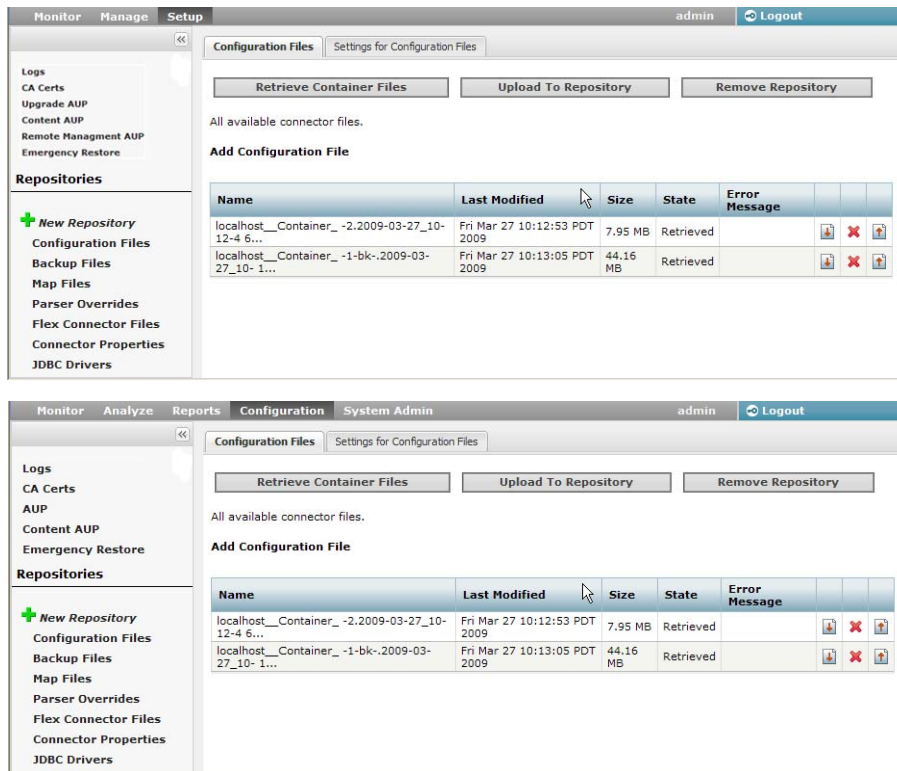
Parameter	Description
Filename Prefix	An identifying word that is included in the names of retrieved files. For example, map files are identified by Map in the file name: <code>localhost_Container_-1.Map-2009-04-06_12-22-25-607.zip</code>
Relative path (Download)	The path for download, relative to <code>\$ARCSIGHT_HOME</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 connectors to a repository. The specific files that are retrieved depend on the settings of a repository.



To retrieve a container file:

- 1 Click **SetupConfiguration** > **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 **SetupConfiguration** > **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 **SetupConfiguration** > **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.

The screenshot shows the 'Settings for Map Files' configuration window. The settings are as follows:

- 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 **SetupConfiguration** > **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.




Connectors require correct properties and proper files. Applying incorrect files, including empty files or files with binary content, can prevent a connector from functioning correctly.



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 **SetupConfiguration** > **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 **SetupConfiguration** > **Repositories** from the top-level menu bar.
- 2 In the left panel, click the name of the repository in which the file exists.
- 3 In the right panel, click  next for the file that you want to upload.
- 4 Follow the Upload Container Files wizard instructions to upload the file to the containers of your choice.
- 5 Verify that the file was uploaded correctly:
 - ◆ If you have SSH access to the connectors, connect to them and check the file structure.
 - ◆ Obtain the connector logs and check the contents of the `Directory.txt` file for each connector.

Removing a File from the Repository

To remove a file from the repository:

- 1 Click **SetupConfiguration** > **Repositories** from the top-level menu bar.
- 2 In the left panel, click the name of the repository in which the file exists.
- 3 In the right panel, click  for the file that you want to delete.

Pre-Defined Repositories

You can define repositories for any connector-related files. As a convenience, the following repositories are pre-defined.

- **Backup Files:** connector cloning (see [“Cloning Container Configuration” on page 453](#)).
- **Map Files:** enrich event data
- **Parser Overrides:** customize the parser (see [“Adding Parser Overrides” on page 454](#))
- **Flex Connector Files:** user-designed connector deployment
- **Connector Properties:** [agent.properties](#); subset of cloning
- **JDBC Drivers:** database connectors

To view the settings for a pre-defined repository, click the name of the repository and then click the **Settings** tab in the right panel.



The settings for pre-defined repositories are read-only; to modify the settings, click **New Repository** in the left panel to create a user-defined repository and provide the settings you want to use.

The following tables lists the settings for each pre-defined repository.

Settings for Backup Files

Name	Default Setting
Name	backup
Display Name	Backup Files
Item Display Name	Backup File
Recursive	Checked (Yes)
Sort Priority	0
Restart Connector Process	Checked (Yes)
Filename Prefix	ConnectorBackup
Download Relative Path	
Download Include regular expression	
Download Exclude regular expression	(agentdata/ cwsapi_fileset_).*\$
Delete before upload	Checked (Yes)
Delete groups	Checked (Yes)
Upload Relative Path	
Delete Relative Path	
Delete Include regular expression	
Delete Exclude regular expression	(agentdata/ cwsapi_fileset_).*\$

Table 9-1 Pre-Defined Settings for Backup Files

Settings for Map Files

Name	Default Setting
Name	map
Display Name	Map Files
Item Display Name	Map File
Recursive	Un-checked (No)
Sort Priority	5
Restart Connector Process	Un-checked (No)
Filename Prefix	Map
Download Relative Path	map
Download Include regular expression	map\[0-9]+\.\properties\$
Download Exclude regular expression	
Delete before upload	Checked (Yes)
Delete groups	Un-checked (No)
Upload Relative Path	
Delete Relative Path	map
Delete Include regular expression	map\[0-9]+\.\properties\$
Delete Exclude regular expression	

Table 9-2 Pre-Defined Settings for Map Files

Settings for Parser Overrides

Name	Default Setting
Name	parseroverrides
Display Name	Parser Overrides
Item Display Name	Parser Override
Recursive	Checked (Yes)
Sort Priority	10
Restart Connector Process	Checked (Yes)
Filename Prefix	Parsers
Download Relative Path	fcg
Download Include regular expression	. *
Download Exclude regular expression	
Delete before upload	Checked (Yes)
Delete groups	Checked (Yes)
Upload Relative Path	
Delete Relative Path	fcg
Delete Include regular expression	. *
Delete Exclude regular expression	

Table 9-3 Pre-Defined Settings for Parser Overrides

Settings for FlexConnector Files

Name	Default Setting
Name	flexconnectors
Display Name	Flex Connector Files
Item Display Name	Flex Connector File
Recursive	Checked (Yes)
Sort Priority	15
Restart Connector Process	Checked (Yes)
Filename Prefix	FlexConnector
Download Relative Path	flexagent
Download Include regular expression	. *
Download Exclude regular expression	
Delete before upload	Checked (Yes)
Delete groups	Checked (Yes)
Upload Relative Path	
Delete Relative Path	flexagent
Delete Include regular expression	. *
Delete Exclude regular expression	

Table 9-4 Pre-Defined Settings for FlexConnector Files

Settings for Connector Properties

Name	Default Setting
Name	connectorproperties
Display Name	Connector Properties
Item Display Name	Connector Property File
Recursive	Un-checked (No)
Sort Priority	20
Restart Connector Process	Checked (Yes)
Filename Prefix	ConnectorProperties
Download Relative Path	
Download Include regular expression	agent\..*
Download Exclude regular expression	
Delete before upload	Un-checked (No)
Delete groups	Un-checked (No)
Upload Relative Path	
Delete Relative Path	
Delete Include regular expression	agent\..*
Delete Exclude regular expression	

Table 9-5 Pre-Defined Settings for Connector Properties

Settings for JDBC Drivers

Name	Default Setting
Name	jdbcdrivers
Display Name	JDBC Drivers
Item Display Name	Connector JDBC Driver File
Recursive	Un-checked (No)
Sort Priority	25
Restart Connector Process	Checked (Yes)
Filename Prefix	
Download Relative Path	lib
Download Include regular expression	
Download Exclude regular expression	
Delete before upload	Un-checked (No)
Delete groups	Un-checked (No)
Upload Relative Path	
Delete Relative Path	lib
Delete Include regular expression	
Delete Exclude regular expression	

Table 9-6 Pre-Defined Settings for JDBC Drivers

Cloning Container Configuration

Using the **Backup Files** repository, you can quickly copy a container to other containers. As a result, all connectors in the source container are copied to the destination container. This process is called *cloning* a container configuration. You can clone a container to several containers at once. The contents of the source container are appended to the existing contents of the destination container.



Caution

Containers on Connector Appliance are pre-installed with the latest connector release. Do not clone older, software-based connectors (such as build 4.0.8.4964) to containers with newer connector builds (such as 4.0.8.4976 or later).

Cloning a connector using the Backup repository only works if the connector version numbers are the same.

To clone a container:

- 1 Click **Manage** from the top-level menu bar **Configuration > Manage Connectors** to list the containers and determine the source and destination for cloning.
- 2 Click **SetupConfiguration > Repositories** from the top-level menu bar.
- 3 Click **Backup Files** under the **Repositories** section in the right panel.
- 4 If the backup file that you need to use for cloning exists in the repository, go to the next step. Otherwise, follow the instructions in ["Retrieving a File from the Repository" on page 446](#) 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 446](#) to upload the backup file to one or more containers.

The destination containers are unavailable while the backup file is applied and the connectors are restarted.



Note

The backup file does not include the container certificates. You have to re-apply the certificates to the container after you upload the backup file. After applying the certificates, check the status of the destination container to make sure it is available.

Adding Parser Overrides

A parser override is a file provided by ArcSight used to resolve an issue with the parser for a specific connector, or to support a newer version of a supported device where the log file format changed slightly or new event types were added.

To use parser overrides, you need to:

- Upload a parser override file to the pre-defined **Parser Overrides** repository.
- Download the parser override file to the container that contains the connector that will use the parser override.

Follow the steps below.

To upload a parser override file:

- 1 Click **SetupConfiguration** > **Repositories** from the top-level menu bar.
- 2 Click **Parser Overrides** under the **Repositories** section in the right panel.
- 3 On the **Parser Overrides** tab, click the **Upload To Repository** button.
- 4 Follow the wizard to upload the file. When prompted by the wizard, make sure you:
 - ◆ Select the **Individual Files** option from the **Select the type of file that you want to upload** field.
 - ◆ Add a slash (/) after fcp before adding the folder name in the **Enter the sub folder where the files will be uploaded** field. For example, `fcp/multisqlserver_audit_db`.

When upload is complete, the parser override file is listed in the table on the Parser Overrides tab.

To download the parser override file to a container:

- 1 Click **SetupConfiguration** > **Repositories** from the top-level menu bar.
- 2 Click **Parser Overrides** under the **Repositories** section in the right panel.
- 3 In the table on the **Parser Overrides** tab, locate the parser override file you want to download and click the up arrow next to the file.
- 4 Follow the wizard to select the container to which you want to add the parser overrides.

When the wizard completes, the parser overrides will be deployed in the selected container.



You can download a parser override file from ArcExchange. For more information, refer to [“Sharing Connectors \(ArcExchange\)” on page 419](#).

To verify that the parser override has been applied successfully, issue a Get Status command to the connector. See [“Sending a Command to a Destination” on page 411](#). In the report that appears, check for the line starting with the text `ContentInputStreamOverrides`.

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, parseable 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". It's 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". It's 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.

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 108](#). 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/EPS 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/EPS 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 60280150", "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}\.\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>[^\s]*)" | 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>[^\s]*)"
```


Restoring Factory Settings

ArcSight Logger can be restored to its original factory settings using built-in Acronis True Image software.

**Caution**

Restoring Logger to factory settings will irrevocably delete all event data and configuration settings.

**Note**

The screens shown here are examples only. Your Logger partitions might vary, and the overall capacity might be different.

To restore factory settings

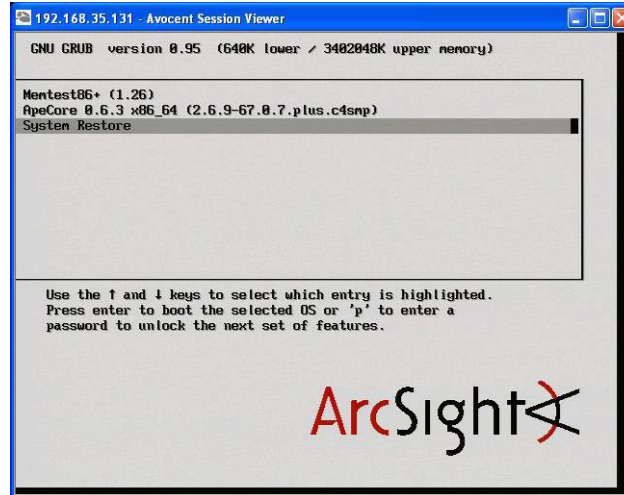
To restore Logger to its original factory settings, perform these steps:

- 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 Logger for Remote Access” on page 33](#) 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 message appears on the screen, press any key on your keyboard. This message 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.

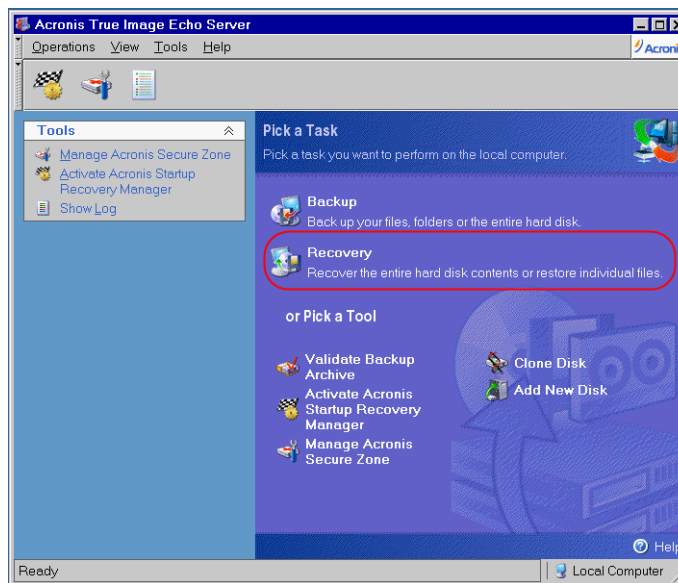
`Press any key to enter the menu`

`Booting Red Hat Enterprise Linux Server in 2 seconds...`

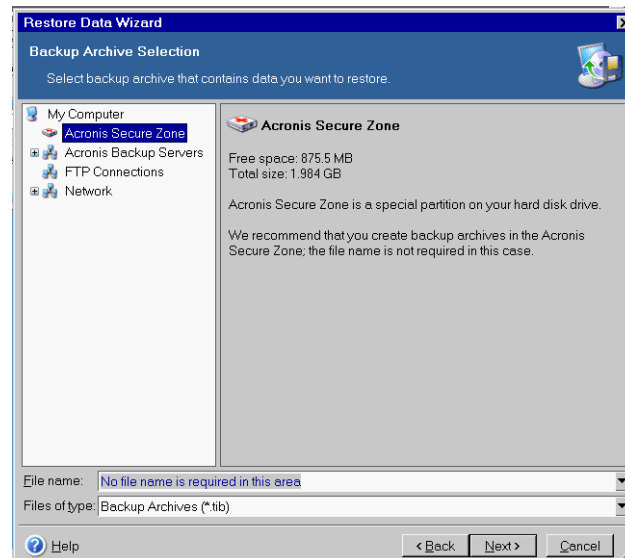
- 4 A screen similar to the following is displayed on the attached monitor. Use the mouse or arrow keys to select System Restore and press the **Enter** key on your keyboard.



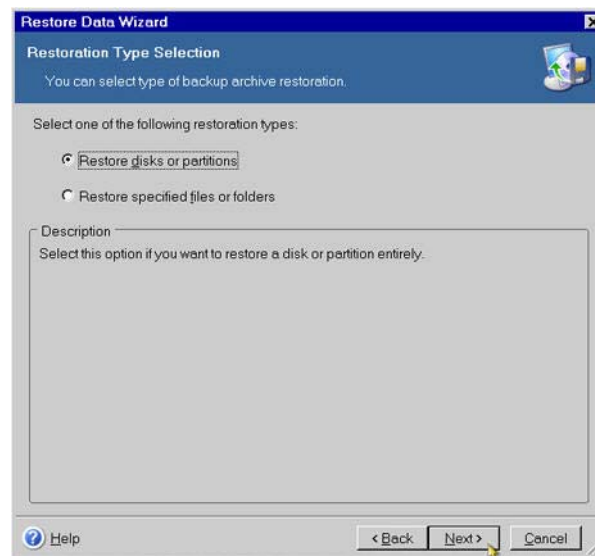
- 5 On the Pick a Task list, as shown in the following figure, choose **Recovery**. On the next page (Welcome to the Restore Data Wizard), click **Next** to continue.



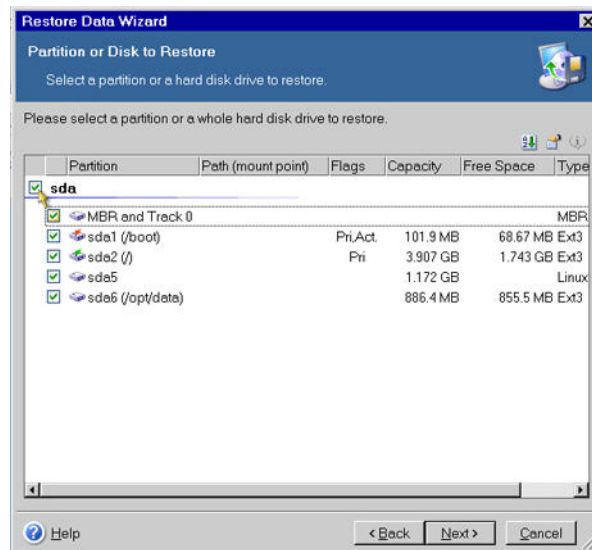
- 6 Select the Acronis Secure Zone, as shown in the following figure, and click **Next**. You will have a chance to review the choices you make on this page and the wizard pages that follow before initiating the restore process.



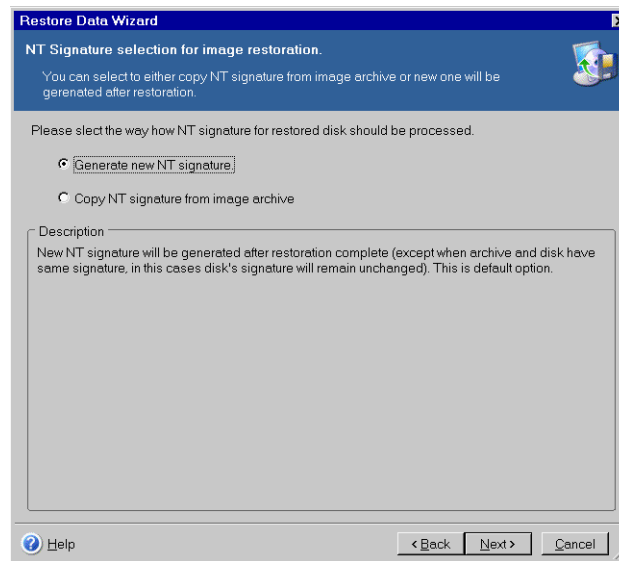
- 7 Select **Restore disks or partitions** and click **Next**. Only choose other options if specifically directed to do so by ArcSight Customer Support.



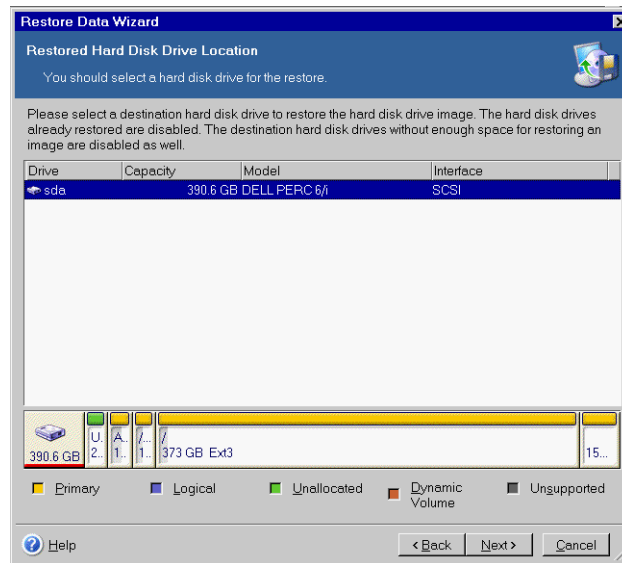
- 8 Select the entire drive, labeled 'sda' in the following figure. Click **Next** to continue.



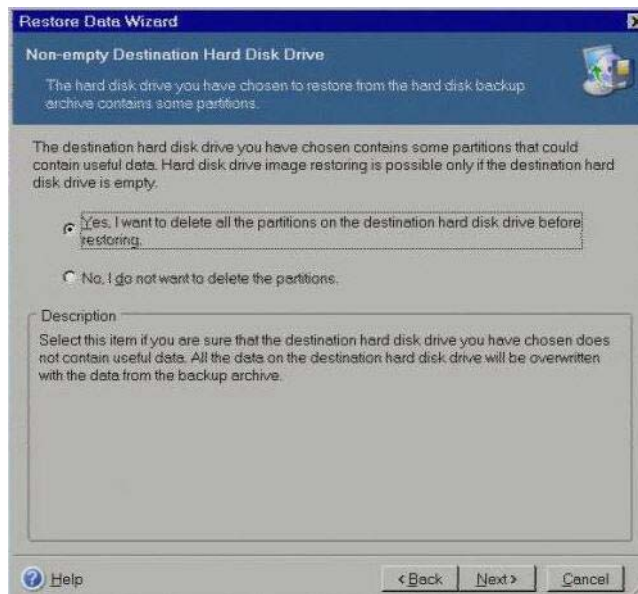
- 9 Select the way in which the NT signature for the restored disk should be processed and click **Next**.



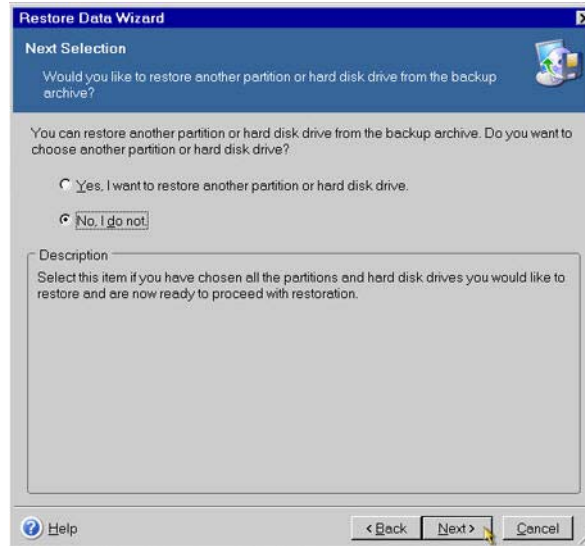
- 10** Choose the drive to restore ('sda') and click **Next**.



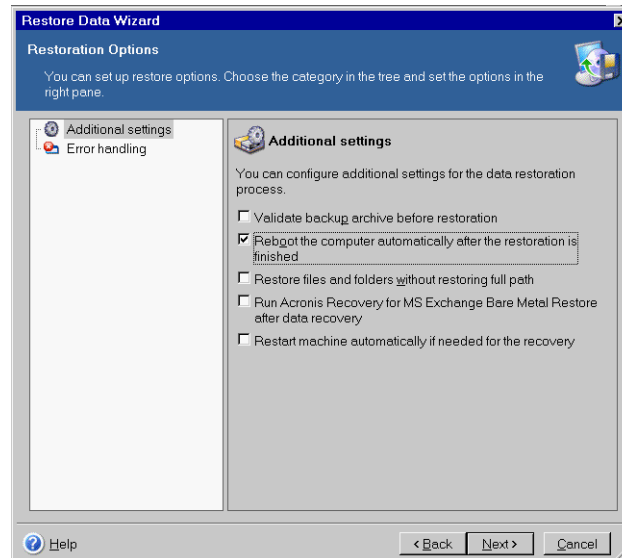
- 11** Select, "Yes, I want to delete all partitions on the destination hard disk drive before restoring", as shown in the following figure.



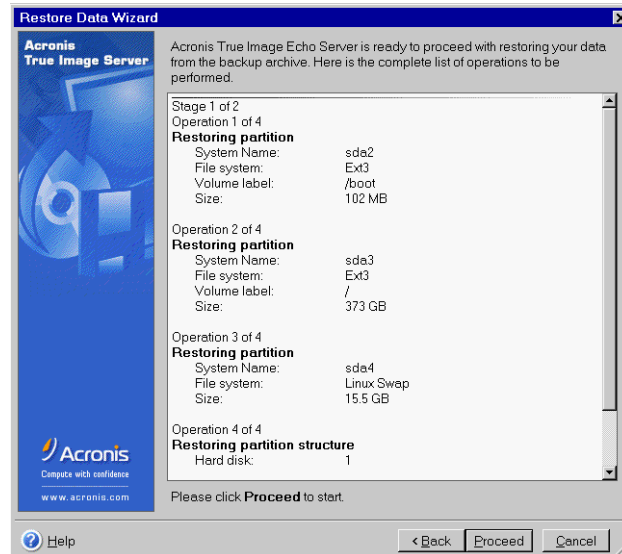
- 12 Because there are no other partitions or disks to restore, choose “No, I do not,” on the Next Selection page of the wizard. Click **Next**.



- 13 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**.

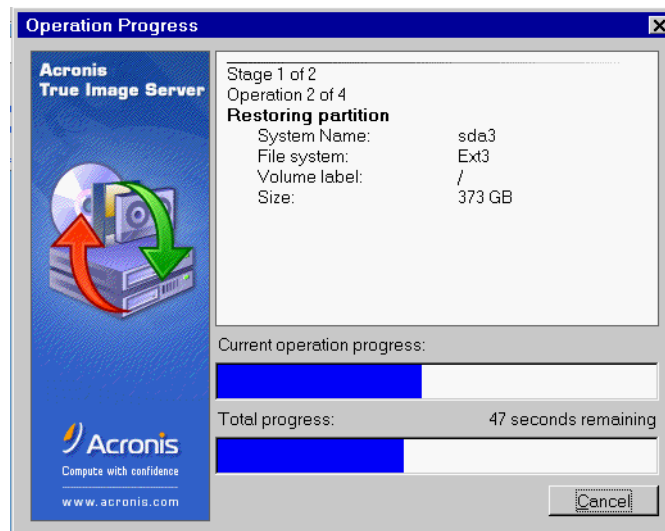


- 14 Review the checklist of operations to be performed, as shown in the following figure, and click **Proceed** to begin the restore process, or click **Back** to revisit previous wizard pages.



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.

- 15 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 313](#)) to forward the events. For a detailed understanding of the format of CEF events, see [Appendix A, Common Event Format, on page 455](#).

The following events are logged and available for Audit Forwarding to ArcSight ESM.

[“Platform Events” on page 484](#)
[“Logger Application Events” on page 490](#)

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

Both types of events are stored in the Logger Internal Storage group. As a result, these events can be searched using the Logger Search UI. For example, you can search for the this platform event:

“/Platform/Authentication/Failure/Password”

In addition to these events, a Logger appliance that has an ArcSight Connector Appliance installed on it generates Connector Appliance audit events. For a list of Connector Appliance audit events, see the *Administrator's Guide for Connector Appliance* for the version that applies to you.

Information in an Audit Event

A Logger audit event (in CEF format) contains information about the following prefix fields:

- Device Event Class ID
- Device Severity
- Message
- Device Event Category—(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: 203	7	Login attempt by inactive user	/Platform/Authentication/Inactive User/Failure	
platform: 205	5	Access enabled for support personnel	/Platform/Authorization/Support/Enable	
platform: 206	1	Access disabled for support personnel	/Platform/Authorization/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	

Device Event Class ID	Severity	Message	Device Event Category (cat)	Additional Fields (listed only if applicable)
platform: 300	5	Installed certificate	/Platform/Certificate/Install	
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: 301	5	Installed certificate revocation list	/Platform/Certificate/Revocation/Install	
platform: 302	5	Deleted trusted certificate	/Platform/Certificate/Delete	
platform: 303	5	Deleted certificate revocation list	/Platform/Certificate/Revocation/Delete	
platform: 304	7	Failed installing trusted certificate	/Platform/Certificate/Install/Failure	
platform: 305	7	Failed installing certificate revocation list	/Platform/Certificate/Revocation/Install/Failure	
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

Device Event Class ID	Severity	Message	Device Event Category (cat)	Additional Fields (listed only if applicable)
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
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

Device Event Class ID	Severity	Message	Device Event Category (cat)	Additional Fields (listed only if applicable)
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
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=fileserv cs2label=Remote directory cs3=trump cs3label=Mount point cs4=type cs4label=Mount type

Device Event Class ID	Severity	Message	Device Event Category (cat)	Additional Fields (listed only if applicable)
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
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: 284	5	Enabled SAN Multipathing	/Platform/Storage/Multi pathing/Enable	
platform: 285	5	Disabled SAN Multipathing	/Platform/Storage/Multi pathing/Disable	
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

Device Event Class ID	Severity	Message	Device Event Category (cat)	Additional Fields (listed only if applicable)
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	
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/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

Device Event Class ID	Severity	Message	Device Event Category (cat)	Additional Fields (listed only if applicable)
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	
platform: 306	5	Start process	/Platform/Process/Start	
platform: 307	5	Stop process	/Platform/Process/Stop	
platform: 308	5	Restart process	/Platform/Process/Restart	
platform: 310	5	Enabled FIPS mode	/Platform/Configuration/FIPS/Enable	
platform: 311	7	Disabled FIPS mode	/Platform/Configuration/FIPS/Disable	

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

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
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
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

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
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
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
logger:528	Event archive failed	/Logger/Resource/Archive/Failed	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
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
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

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
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
Loggers			
logger:550	Logger [name] has been added	/Logger/Resource/Logger/Configuration/Add	fname=Name duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType= Logger fileId=LoggerId
logger:551	Logger [name] has been deleted	/Logger/Resource/Logger/Configuration/Delete	fname=Name duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType= Logger fileId=LoggerId
logger:570	Logger authorization [name] has been added	/Logger/Resource/Logger/Authorizations/Configuration/Add	fname=Name duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType= Logger Authorization
logger:571	Logger authorization [name] has been deleted	/Logger/Resource/Logger/Authorizations/Configuration/Delete	fname=Name duser=username duid=userId cs4=sessionIdfile cs4Label=Session ID fileType= Logger Authorization fileId=LoggerId
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

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
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
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

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
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 e 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 e cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
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

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
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
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

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
logger: 641	ESM destination [name] has been deleted	/Logger/Component/EsmDestination/Configuration/Delete	fname=esmDestinationName duser=username duid=userId cs4=sessionId file cs4Label=Session ID fileType=ESM Destination fileId=esmDestinationId dvc=esmDestinationIp dvchost=esmDestinationHost cn1Label=ESM Destination Port cn1=esmDestinationPort cs1Label=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
logger: 644	SNMP destination [name] has been added	/Logger/Component/SnmpDestination/Configuration/Add	fname=snmpDestinationName duser=username duid=userId cs4=sessionId file 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

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
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
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

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
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
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

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
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
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=syslogOrSnmpOrE smHostName cn1Label=SyslogoOrSnmpO rEsmDestination Port cn1=syslogOrSnmpOrEsmP ort 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

Device Event Class ID	Message	Device Event Category (cat)	Additional Fields
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

Examples of System Health Events

The following table provides examples of system health events generated on Logger. These examples are intended to help you understand the format and various fields of the generated events.

Device Event Class ID	Example
hardware:101	CEF:0 ArcSight Logger 5.1.0.5784.0 hardware:101 Electrical (Current) OK 1 cat=/Monitor/Sensor/Current/Ok cs1=0.80 Amps cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=10.2 (Power Supply) cs5Label=Location cs6=Current 2 cs6Label=Sensor Name dst=192.168.36.5 dvc=192.168.36.5 end=1303937520837 rt=1303937520837
hardware:102	CEF:0 ArcSight Logger 5.1.0.5776.0 hardware:102 Electrical (Current) Degraded 5 cat=/Monitor/Sensor/Current/Degraded cs1=126 Watts cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Inc cs4Label=Raw Status cs5Label=Location cs6=Power Meter cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302817019262 rt=1302817019262
hardware:103	CEF:0 ArcSight Logger 5.1.0.5776.0 hardware:103 Electrical (Current) Failed 8 cat=/Monitor/Sensor/Current/Failed cs1=126 Watts cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=lcr cs4Label=Raw Status cs5Label=Location cs6=Power Meter cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302817019262 rt=1302817019262
hardware:111	CEF:0 ArcSight Logger 5.1.0.5780.0 hardware:111 Electrical (Voltage) OK 1 cat=/Monitor/Sensor/Voltage/Ok cs1=State Deasserted cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=3.1 (Processor) cs5Label=Location cs6=VCORE cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302819047959 rt=1302819047959
hardware:112	CEF:0 ArcSight Logger 5.1.0.5780.0 hardware:112 Electrical (Voltage) Degraded 5 cat=/Monitor/Sensor/Voltage/Degraded cs1=State Deasserted cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Inc cs4Label=Raw Status cs5=3.1 (Processor) cs5Label=Location cs6=VCORE cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302819047959 rt=1302819047959
hardware:113	CEF:0 ArcSight Logger 5.1.0.5780.0 hardware:113 Electrical (Voltage) Failed 8 cat=/Monitor/Sensor/Voltage/Failed cs1=State Deasserted cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=lcr cs4Label=Raw Status cs5=3.1 (Processor) cs5Label=Location cs6=VCORE cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302819047959 rt=1302819047959

Device Event Class ID	Example
hardware:121	CEF:0 ArcSight Logger 5.1.0.5803.0 hardware: 121 Battery OK 1 cat=/Monitor/Sensor/Battery/Ok cs1= cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=CMOS Battery cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303937972008 rt=1303937972008
hardware:122	CEF:0 ArcSight Logger 5.1.0.5803.0 hardware: 122 Battery Degraded 5 cat=/Monitor/Sensor/Battery/Degraded cs1= cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Inc cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=CMOS Battery cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303937972008 rt=1303937972008
hardware:123	CEF:0 ArcSight Logger 5.1.0.5803.0 hardware: 123 Battery Failed 8 cat=/Monitor/Sensor/Battery/Degraded cs1= cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=lcr cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=CMOS Battery cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303937972008 rt=1303937972008
hardware:131	CEF:0 ArcSight Logger 5.1.0.5780.0 hardware: 131 Fan OK 1 cat=/Monitor/Sensor/Fan/Ok cs1=29.01 unspecified cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=Inc cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=Fan Block 1 cs6Label=Sensor Name dst=192.168.37.21 dvc=192.168.37.21 end=1302823237825 rt=1302823237825
hardware:132	
hardware:133	CEF:0 ArcSight Logger 5.1.0.5780.0 hardware: 133 Fan Failure 8 cat=/Monitor/Sensor/Fan/Failed cs1=29.01 unspecified cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=lcr cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=Fan Block 1 cs6Label=Sensor Name dst=192.168.37.21 dvc=192.168.37.21 end=1302823237825 rt=1302823237825
hardware:141	CEF:0 ArcSight Logger 5.1.0.5803.0 hardware: 141 Power Supply OK 1 cat=/Monitor/Sensor/PowerSupply/Ok cs1= cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=10.1 (Power Supply) cs5Label=Location cs6=Status cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303938572149 rt=1303938572149
hardware:142	
hardware:143	CEF:0 ArcSight Logger 5.1.0.5776.0 hardware: 143 Power Supply Failed 8 cat=/Monitor/Sensor/PowerSupply/Failed cs1=0 Watts cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=lcr cs4Label=Raw Status cs5=10.2 (Power Supply) cs5Label=Location cs6=Power Supply 2 cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302817019263 rt=1302817019263
hardware:151	CEF:0 ArcSight Logger 5.1.0.5776.0 hardware: 151 Temperature OK 1 cat=/Monitor/Sensor/Temperature/Ok cs1=17 degrees C cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=64.1 (Unknown (0x40)) cs5Label=Location cs6=Temp 1 cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302823560051 rt=1302823560051
hardware:152	
hardware:153	

Device Event Class ID	Example
raid:101	CEF:0 ArcSight Logger 5.1.0.5780.0 raid:101 RAID Controller OK 1 cat=/Monitor/RAID/Controller/Ok cs1=Type: RAID-5 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=Optimal cs4Label=Raw Status cs5=RAIDController/Configuration cs5Label=Location cs6=RAIDController/Configuration cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302886250104 rt=1302886250104
raid:102	CEF:0 ArcSight Logger 5.1.0.5780.0 raid:102 RAID Controller Degraded 5 cat=/Monitor/RAID/ControllerDegraded cs1=Type: RAID-5 Critical Disks: 0 Failed Disks: 0 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Degraded cs4Label=Raw Status cs5=RAIDController/Configuration cs5Label=Location cs6=RAIDController/Configuration cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302826128482 rt=1302826128482
raid:103	
raid:111	CEF:0 ArcSight Logger 5.1.0.5776.0 raid:111 RAID BBU OK 1 cat=/Monitor/RAID/BBU/Ok cs1=Battery/Capacitor Count: 1 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=OK cs4Label=Raw Status cs5=RAIDController/Battery/bbu cs5Label=Location cs6=RAIDController/Battery/bbu cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302890169285 rt=1302890169285
raid:112	CEF:0 ArcSight Logger 5.1.0.5780.0 raid:112 RAID BBU Degraded 5 cat=/Monitor/RAID/BBU/Degraded cs1=Fully Charged: false Remaining Time Alarm: false Remaining Capacity Alarm: false Over Charged: false isSOHGood: true cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Degraded cs4Label=Raw Status cs5=RAIDController/Battery/bbu cs5Label=Location cs6=RAIDController/Battery/bbu cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302820608015 rt=1302820608015
raid:113	
raid:121	CEF:0 ArcSight Logger 5.1.0.5780.0 raid:121 RAID Disk OK 1 cat=/Monitor/RAID/DISK/Ok cs1=Port: 1I Box: 1 Bay: 1 Size: 500 GB Serial Number: 9SP24JD5 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=OK cs4Label=Raw Status cs5=Port: 1I Box: 1 Bay: 1 Serial Number: 9SP24JD5 cs5Label=Location cs6=RAIDController/Port/p0 cs6Label=Sensor Name dst=192.168.37.21 dvc=192.168.37.21 end=1302849041777 rt=1302849041777
raid:122	CEF:0 ArcSight Logger 5.1.0.5776.0 raid:122 RAID Disk Rebuilding 5 cat=/Monitor/RAID/DISK/Rebuilding cs1=Port: 2I Box: 1 Bay: 1 Size: 1 TB Serial Number: WMATV6348517 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Rebuilding cs3Label=Status cs4=Rebuilding cs4Label=Raw Status cs5=Port: 2I Box: 1 Bay: 1 Serial Number: WMATV6348517 cs5Label=Location cs6=RAIDController/Port/p0 cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302826980530 rt=1302826980530

Device Event Class ID	Example
raid:123	CEF:0 ArcSight Logger 5.1.0.5780.0 raid:123 RAID Disk Failed 8 cat=/Monitor/RAID/DISK/Failed cs1=Port: 1I Box: 1 Bay: 2 Size: 500 GB Serial Number: 9SP23M08 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=Failed cs4Label=Raw Status cs5=Port: 1I Box: 1 Bay: 2 Serial Number: 9SP23M08 cs5Label=Location cs6=RAIDController/Port/p1 cs6Label=Sensor Name dst=192.168.37.21 dvc=192.168.37.21 end=1302826358346 rt=1302826358346
network:100	CEF:0 ArcSight Logger 5.1.0.5780.0 network:100 Network Usage - Inbound 1 cat=/Monitor/Network/Usage/In cn1=41837428 cn1Label=Bytes Received cs2=SinceStartup cs2Label=timeframe cs6=eth0 cs6Label=Interface Name dst=192.168.35.115 dvc=192.168.35.115 end=1302733620026 rt=1302733620026
network:101	CEF:0 ArcSight Logger 5.1.0.5780.0 network:101 Network Usage - Outbound 1 cat=/Monitor/Network/Usage/Out cn1=158442791 cn1Label=Bytes Sent cs2=SinceStartup cs2Label=timeframe cs6=eth0 cs6Label=Interface Name dst=192.168.35.115 dvc=192.168.35.115 end=1302733620028 rt=1302733620028
memory:100	CEF:0 ArcSight Logger 5.1.0.5780.0 memory:100 Platform Memory Usage 1 cat=/Monitor/Memory/Usage/Platform cn1=2757 cn1Label=MB Used cs2=CurrentValue cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1302797940018 rt=1302797940018
cpu:100	CEF:0 ArcSight Logger 5.1.0.5780.0 cpu:100 CPU Usage 1 cat=/Monitor/CPU/Usage cn1=3 cn1Label=Percent Usage cs2=CurrentValue cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1302739080014 rt=1302739080014
disk:101	CEF:0 ArcSight Logger 5.1.0.5803.0 disk:101 Root Disk Space Remaining 1 cat=/Monitor/Disk/Space/Remaining/Root cn1=99 cn1Label=Percent Available cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=Ok cs4Label=Raw Status cs5=Root cs5Label=Location cs6=Disk/Space/Remaining/Root cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303927171790 rt=1303927171790
disk:102	CEF:0 ArcSight Logger 5.1.0.5780.0 disk:102 Disk bytes read 1 cat=/Monitor/Disk/Read cn1=373524 cn1Label=Kb Read cs2=SinceStartup cs2Label=timeframe cs6=sda cs6Label=Partition Name dst=192.168.35.115 dvc=192.168.35.115 end=1302743760036 rt=1302743760036
disk:103	CEF:0 ArcSight Logger 5.1.0.5780.0 disk:103 Disk bytes written 1 cat=/Monitor/Disk/Write cn1=24474998 cn1Label=Kb Written cs2=SinceStartup cs2Label=timeframe cs6=sda cs6Label=Partition Name dst=192.168.35.115 dvc=192.168.35.115 end=1302743760038 rt=1302743760038
eps:100	CEF:0 ArcSight Logger 5.1.0.5780.0 eps:100 Overall Receiver EPS 1 cat=/Monitor/Receiver/All/EPS cn1=0 cn1Label=EPS cs2=SinceLastMonitorEvent cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1302733680034 rt=1302733680034
eps:101	CEF:0 ArcSight Logger 5.1.0.5780.0 eps:101 Overall Forwarder EPS 1 cat=/Monitor/Forwarder/All/EPS cn1=0 cn1Label=EPS cs2=SinceLastMonitorEvent cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115

Device Event Class ID	Example
eps:102	CEF:0 ArcSight Logger 5.1.0.5803.0 eps:102 Individual Receiver EPS 1 cat=/Monitor/Receiver/EPS/Individual cn1=0 cn1Label=EPS cs1=N/A cs1Label=Receiver Type cs2=SinceLastMonitorEvent cs2Label=timeframe cs6=udp514 cs6Label=name dst=192.168.35.115 dvc=192.168.35.115 end=1303927500046 rt=1303927500046
eps:103	CEF:0 ArcSight Logger 5.1.0.5780.0 eps:103 Individual Forwarder EPS 1 cat=/Monitor/Forwarder/One/EPS cn1=0 cn1Label=EPS cs1=N/A cs1Label=Forwarder Type cs2=SinceLastMonitorEvent cs2Label=timeframe cs6=esm cs6Label=name dst=192.168.35.115 dvc=192.168.35.115 end=1302733080052 rt=1302733080052
search:100	CEF:0 ArcSight Logger 5.1.0.5780.0 search:100 Number of Searches Performed 1 cat=/Monitor/Search cn1=0 cn1Label=Number of Searches cs2=SinceStartup cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1302741300026 rt=1302741300026
storagegroup:100	CEF:0 ArcSight Logger 5.1.0.5803.0 storagegroup:100 Storage Group Space Used 1 cat=/Monitor/StorageGroup/Space/Used cn1=1 cn1Label=Percent Used cn2=7 cn2Label=retention period (days) cn3=1024 cn3Label=used (MB) cs2=CurrentValue cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1303928072008 fileType=storageGroup fname=Default Storage Group fsize=1534 rt=1303928072008

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 page 359](#) and [Chapter 9, Managing Repositories, on page 429](#) 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 Download Center at <https://arcsight.subscribenet.com>.

Destination Runtime Parameters

The following table describes the destination parameters you can configure. The parameters listed in the table are not available for all destinations. The user interface automatically displays the parameters valid for a destination. For step-by-step instructions on updating the runtime parameters of a destination, see [“Editing Destination Parameters” on page 405](#).

Name Fields	Value Fields
Batching	Connectors can batch events to increase performance and optimize network bandwidth. When activated, connectors create blocks of events and send them when they either (1) reach a certain size or (2) the time window expires, whichever occurs first. You can also prioritize batches by severity, forcing the connector to send the highest-severity event batches first and the lowest-severity event batches later.
Enable Batching (per event)	Create batches of events of this specified size (5, 10, 20, 50, 100 , 200, 300 events).
Enable Batching (in seconds)	The connector sends the events if this time window expires (1, 5 , 10, 15, 30, 60).
Batch By	This is Time Based if the connector should send batches as they arrive (the default) or Severity Based if the connector should send batches based on severity (batches of Highest Severity events sent first).
Time Correction	The values you set for these fields establish forward and backward time limits, that if exceeded, cause the connector to automatically correct the time reported by the device.
Use Connector Time as Device Time	Override the time the device reports and instead use the time at which the connector received the event. This option assumes that the connector will be more likely to report the correct time. (No Yes)
Enable Device Time Correction (in seconds)	The connector can adjust the time reported by the device 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 connector can also adjust the time reported by the connector itself, using this setting. This is for informational purposes only and allows you to modify the local time on the connector. This should be a temporary setting. The recommended way to synchronize clocks between Manager and connectors is the NTP protocol. The default is 0 .

Set Device Time Zone To Ordinarily, it is presumed that the original device is reporting its time zone along with its time. And if not, it is then presumed that the connector is doing so. If this is not true, or the device isn't reporting correctly, you can switch this option from Disabled to GMT or to a particular world time zone. That zone is then applied to the time reported. Default: **Disabled**.

Device Time Auto-correction

Future Threshold The connector sends the internal alert if the detect time is greater than the connector time by **Past Threshold** seconds.

Past Threshold The connector sends the internal alert if the detect time is earlier than the connector time by **Past Threshold** seconds.

Device List A comma-separated list of the devices to which the thresholds apply. The default, **(ALL)**, means all devices.

Time Checking

These are the time span and frequency factors for doing device-time auto-correction.

Future Threshold The number of seconds by which to extend the connector's forward threshold for time checking. The default is **5 minutes** (300 seconds).

Past Threshold The number of seconds by which to extend the connector's rear threshold for time checking. Default is **1 hour** (3,600 seconds).

Frequency The connector checks its future and past thresholds at intervals specified by this number of seconds. Default is **1 minute** (60 seconds).

Cache

Changing these settings will not affect the events cached, it will only affect new events sent to the cache.

Cache Size Connectors use a compressed disk cache to hold large volumes of events when the ArcSight ESM Manager is down or when the connector receives bursts of events. This parameter specifies the disk space to use. The default is **1 GB** which, depending on the connector, can hold about 15 million events, but it also can go down to **5 MB**. When this disk space is full, the connector drops the oldest events to free up disk cache space. (5 MB, 50 MB, 100 MB, 150 MB, 200 MB, 250 MB, 500 MB, 1 GB, 2.5 GB, 5 GB, 10 GB, 50 GB.)

Notification Threshold The size of the cache's contents at which to trigger a notification. Default is **10,000**.

Notification Frequency How often to send notifications after the Notification Threshold is reached. (1 minute, 5 minutes, **10 minutes**, 30 minutes, 60 minutes.)

Network

Heartbeat Frequency This setting controls how often the connector sends a heartbeat message to the destination. The default is **10 seconds**, but it can go from **5 seconds** to **10 minutes**. Note that the heartbeat is also used to communicate with the connector; therefore, if its frequency is set to **10 minutes**, then it could take as much as 10 minutes to send any configuration information or commands back to the connector.

Enable Name Resolution The connector tries to resolve IP addresses to hostnames, and hostnames to IP addresses, if required and if the event rate allows. This setting controls this functionality. The Source, Target and Device IP addresses and Hostnames might also be affected by this setting. By default, name resolution is enabled (**Yes**).

Name Resolution Host Name Only	Default: Yes .
Name Resolution Domain From E-mail	Default: Yes .
Clear Host Names Same as IP Addresses	Default: Yes .
Don't Resolve Host Names Matching	NA
Don't Reverse-Resolve IP Ranges	NA
Limit Bandwidth To	A list of bandwidth options you can use to constrain the connector's output over the network. (Disabled , 1 kbit/sec to 100 Mbits/sec.)
Transport Mode	You can configure the connector to cache to disk all the processed events it receives. This is equivalent to pausing the connector. However, you can use this setting to delay event-sending during particular time periods. For example, you could use this setting to cache events during the day and send them at night. You can also set the connector to cache all events, except for those marked with a very-high severity, during business hours, and send the rest at night. (Normal Cache Cache (but send Very High severity events)).
Address-based Zone Population Defaults Enabled	This field applies to 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 connector'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 connector'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 connector'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 connector'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 connector'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 connector'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.

Connector aggregation significantly reduces the amount of data received, and should be applied only when you use less than the total amount of information the event offers. For example, you could enable field-based aggregation to aggregate "accepts" and "rejects" in a firewall, but you should use it only if you are interested in the count of these events, instead of all the information provided by the firewall.

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	<p>Filter Aggregation is a way of capturing aggregated event data from events that would otherwise be discarded due to an agent filter. Only events that would be filtered out are considered for filter aggregation (unlike Field-based aggregation, which looks at all events).</p> <p>Connector aggregation significantly reduces the amount of data received, and should be applied only when you use less than the total amount of information the event offers.</p>
Time Interval	Choose a time interval, if applicable, to use as a basis for aggregating the events the connector collects. It is exclusive of Event Threshold. (Disabled , 1 sec, 5 sec, and so on, up to 1 hour.)
Event Threshold	Choose a number of events, if applicable, to use as a basis for aggregating the events the connector collects. This is the maximum count of events that can be aggregated; for example, if 150 events were found to be the same within the time interval selected (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 connectors by choosing one of two "turbo" (narrower data bandwidth) modes. The default transfer mode is called **Complete**, which passes all the data arriving from the device, including any additional data (custom, or vendor-specific).

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 connectors need to run in **Complete** mode, to capture the additional data.

Note: Connector 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 connector 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 connector's burden on the CPU by reducing its processing rate. This can also be a means of dealing with the effects of event bursts.</p> <p>The choices range from Disabled (no limitation on CPU demand) to 1 eps (pass just one event per second, making the smallest demand on the CPU).</p> <p>Note: The effect of this option varies with the category of connector in use, as described in the connector Processing Categories table below.</p>
Fields to Obfuscate	
Store Original Time in	Disabled or Flex Date 1.
Enable Port-Service Mapping	Default: No .
Enable User Name Splitting	Default: No .
Split File Name into Path and Name	Default: No .
Event Integrity Algorithm	Disabled , SHA-1, SHA-256, SHA-512, or MD5.
Generate Unparsed Events	Default: No .

Preserve System Health Events Yes, **No**, or Disabled.

Enable Device Status Monitoring (in minutes) **Disabled** or 1, 2, 3, 4, 5, 10, 30, 60, or 120 minutes.

Filters

Filter Out NA

"Very High Severity" Event Definition NA

"High Severity" Event Definition NA

"Medium Severity" Event Definition NA

"Low Severity" Event Definition NA

"Unknown Severity" Event Definition NA

Payload Sampling (When available.)

Max. Length Discard, 128 bytes, **256 bytes**, 512 bytes, 1 kbyte

Mask Non-Printable Characters Default: **False**.

Appendix I

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

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.

Starting with Logger v5.1, this functionality is also available on software Logger.

Understanding the Integrated Search Functionality

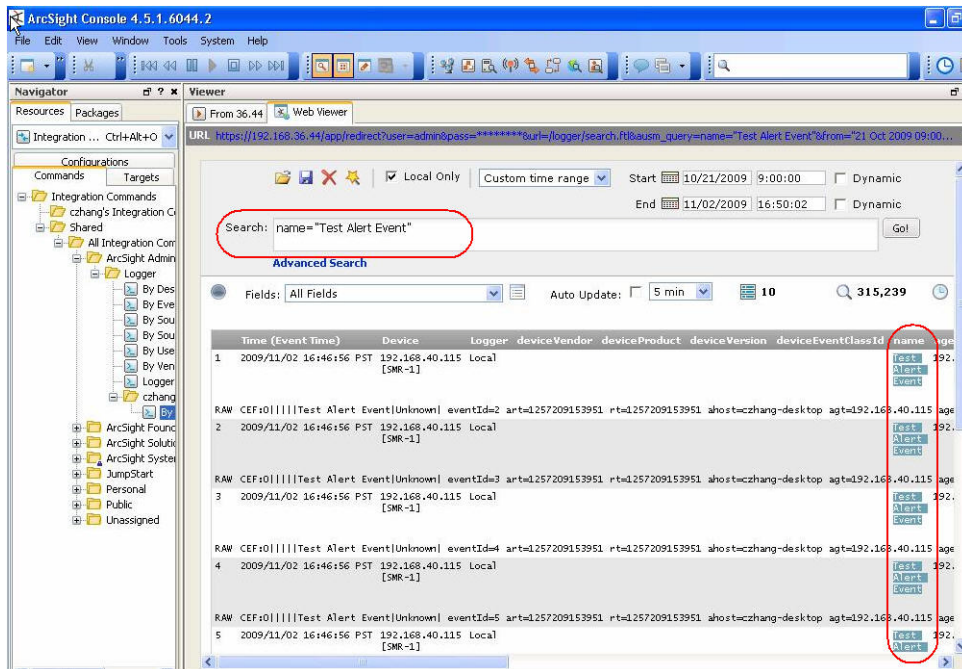
There are two ways to perform a search operation on Logger from an ESM Console:

- 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, right click on an event in an active channel of the ESM Console to display a menu to select a search method—Logger Search or Logger Quick Search.

If you select Logger Search, you need to select search options such as Event Name, Destination, Source, and so on, and the Logger appliances on which the search should be run (if there are multiple Logger appliances). If you select Logger Quick Search, you are not prompted to specify any search options because the search is run on the field you had clicked on.

The search results are displayed in the ESM Console, as shown in the following figure:



Before you can run a search operation on Logger from ESM Console, you need to set up parameters in the ESM Console that are used to authenticate the user who performs the search. Prior to Logger v5.1, only basic authentication (user name and password) was available; however, starting with Logger v5.1, a One Time Password (OTP) option is available. This option makes the user authentication between Logger and ESM Console highly secure. For OTP option to work, Logger must be running v5.1 or later, and the ESM Console must be running ESM v5.0 SP1 Patch 2 or later, as described in ["Prerequisites" on page 526](#).

By default, a Logger search from the ESM Console uses the OTP method to authenticate. However, if Logger or ESM Console is not running a release that supports the OTP option, an error message is displayed and basic authentication is used.

Prerequisites

The following table lists the minimum and recommended versions that Logger and ESM Console must be running.

Option	Requirement
Recommended	<p>Logger v5.1 (appliance and software)</p> <p>ESM v5.0 SP1 Patch 2</p> <p>By default, OTP authentication is used. These release versions must be installed for OTP authentication to work.</p>
Minimum	<p>Logger v4.0 or later (on appliance only)</p> <p>ESM v4.5 SP1 Patch 2 or later</p> <p>Basic authentication is used.</p>

Setup and Configuration

ESM

Follow these instructions to set up and configure ArcSight ESM Manager to run integrated search operations:

- 1 Ensure that the ESM Manager is running one of the recommended versions. (See ["Prerequisites" on page 526.](#))
- 2 Follow instructions in the ArcSight ESM *User's Guide* to set up ESM Console for integrated searches on Logger. When setting up a user for Logger access (as described in the "Set Up Users for Logger Access" section of the *User's Guide*), specify the following integration parameters, also shown in the figure below.

Parameter	Type	Value	Targets
OTPPassword	Password	••••••••	Logger Appliance 1
LoggerHost	Text	192.168.36.29	Logger Appliance 1
OTPUser	Text	logger_user	Logger Appliance 1
LoggerPort	Text	443	Logger Appliance 1
LoggerUser	Text	logger_user	Logger Appliance 1
LoggerPassword	Password	••••••••	Logger Appliance 1

Parameter	Description
LoggerUser	<p>The Logger user account to use when accessing a Logger target.</p> <p>For software Logger, this parameter is not applicable. Only OTP method is supported.</p>
LoggerPassword	<p>The password for that Logger account.</p> <p>For software Logger, this parameter is not applicable. Only OTP method is supported.</p>
LoggerHost	The IP address of the Logger host.
OTPUser	The Logger user account to use with the OTP authentication. This account must exist on the Logger.
OTPPassword	The password to use for the OTPUser specified above.
LoggerPort	<p>For OTP, you must specify the port number for the Logger.</p> <p>For a Logger appliance, the port number is 443.</p> <p>For software Logger, specify the port you configured on it during installation.</p>

The ArcSight ESM User's Guide is available from the ArcSight Download Center at <http://www.arcsight.com/supportportal>.

Logger

Make sure:

- 1 Your Logger appliance is running a version listed in ["Prerequisites" on page 526.](#)

- 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 527](#)) 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 528](#)) 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.

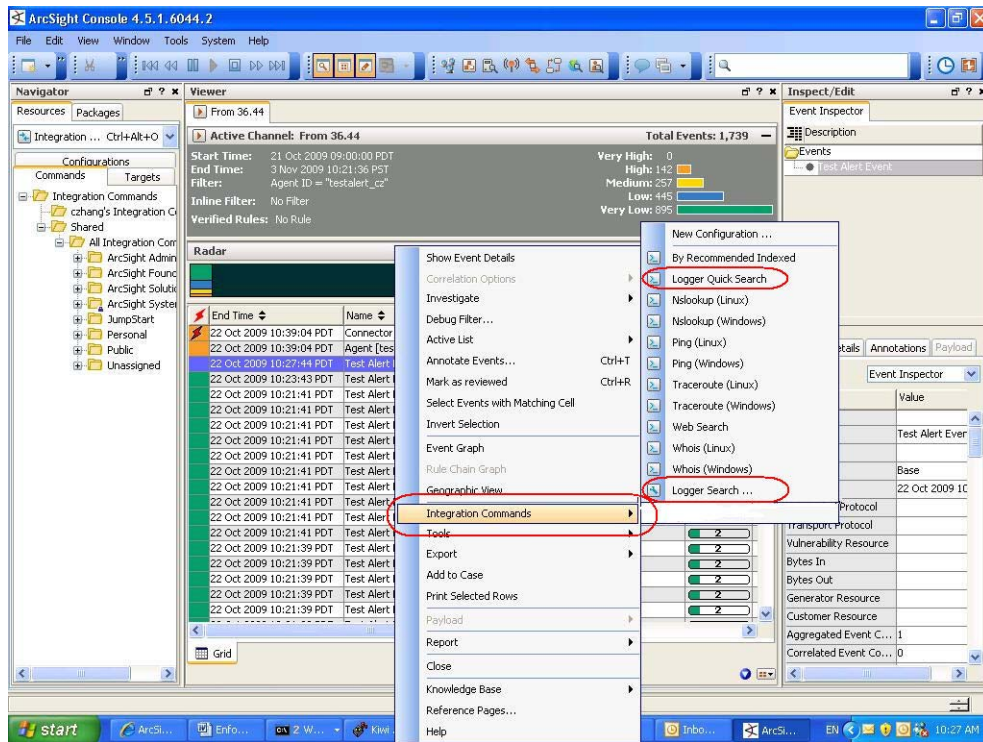
Additionally, even if a Logger peers with other Loggers, integrated search runs only on the local Logger that you select from the menu list on ESM Console.

- The One-Time Password (OTP) authentication is available for use only when Logger is running v5.1 or later and ESM Console is running v5.0 SP1 Patch 2 or later. If OTP cannot be used, the searches run from the ESM Console display a message that a single-use session token could not be negotiated thus regular authentication will be used. Click OK in that message window so the LoggerUser and LoggerPassword is used to authenticate.

Searching on Logger From ESM Console

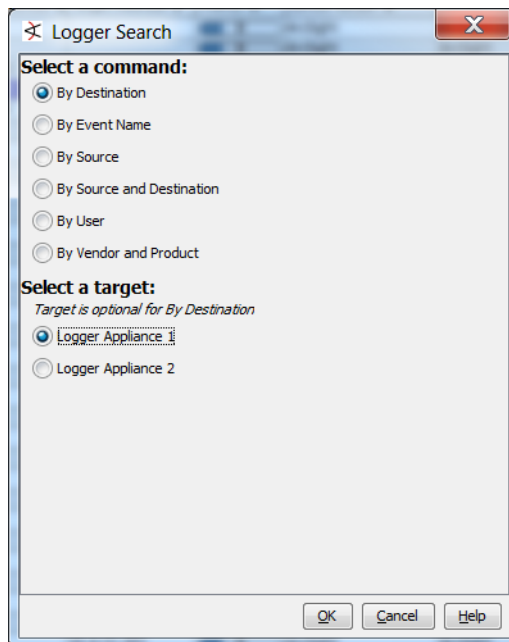
- To run a **Quick Search** on Logger (as described in [“Understanding the Integrated Search Functionality” on page 525](#)):
 - a Right click on the event field in an active channel of the ESM Console.

- 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.
- d** If Logger or ESM Console is not running a release that supports the OTP option, an error message is displayed indicating that a single-use session token was not negotiated and basic authentication will be used instead.

When such a message is displayed, click OK to proceed.

The search results are displayed in the ESM Console Web Viewer.

Symbols

.aup file for content update 437

A

accounts, user. See User.
Acronis True Image Server 475
advanced mode, packaging connectors 420
agents. See SmartConnectors
Alerts

- about 255
- adding 258
- disable 259, 266
- enable 259, 266
- remove 260

apache status 310, 346
ArcExchange 419
archive, scheduled 230
ArcSight ESM 16, 20, 21, 56, 59, 246, 254
ArcSight Logger Streaming SmartConnector 247
ArcSight Manager 22, 56, 59, 246, 252
ArcSight SmartConnectors 21, 57
Audit forwarding 483
Audit Log 313, 348, 483
AUP upgrade process 436
authentication, RADIUS 335
automatic timeout 333

B

backup, configuration 284
basic mode, packaging connectors 419
batching 511
browser requirements 62
bulk copy (see cloning) 453

C

CA certificate

- applying on container 380
- demo 379
- invalid errors 385
- managing 379
- removing from container 382
- viewing list 383

CAC support 326
CACERTS for ESM Destination 253
canonical equality check 65
cases 424
case-sensitive search 65
Categories tab 517
CEF 125, 455

CEF event filters 125
certificate revocation list 328
Certificate Signing Request 324
Certificate, installation 326
changing container credentials 377
character encoding 242
CIFS, configuring 390
CLI 25

- command table 26

cloning connectors 453
Comma Separated Values file, uploading 365
command line interface (CLI) 25
Common Event Format (CEF) 125, 455
Common Extension Dictionary 457
Configuration Backup 284
configuration monitoring, reports for 137
Configuration tab 223, 301
Connector Appliance

- remote upgrade 436

Connector Forwarder 247
connectors supported 390
connectors. See SmartConnectors
constraints, search 72
containers

- adding 375
- changing credentials 377
- definition 374
- deleting 376
- editing 375
- running commands 385
- updating properties 376
- upgrading 386
- viewing all 374
- viewing logs 387

content AUP 437
copying (see cloning) 453
CSR

- generating a certificate signing request 325

CSV file information 365
current time, changing 306
custom connector 419
customers 513

D

dashboard

- reports 140
- reports, preference for display 152

date/time format 246
default gateway 304
Default Storage Group 16, 32
demo certificate 379

- deploying
 - report package 219
- Device 223
 - delete 224
 - edit 224
 - pre-defining 224
- Device Group 225
 - creating 225
 - deleting 226
 - editing 226
- device group
 - maximum number 225
- devices
 - maximum number 224
- DNS Settings 303
 - changing 303, 304
- dynamic search 93

E

- e-mailing
 - reports 162
- encoding 242
- Error Log 313, 348
- ESM (ArcSight Enterprise Security Manager) 16, 20, 21, 56, 59, 246
- ESM Destination 251
 - creating 252
 - deleting 253
 - updating CACERTS 253
- ESM SmartConnector status 310, 346
- etc/hosts.txt 303
- eth0 414
- Event Archive 226, 229
 - adding 230
 - deleting 230
 - loading 232
 - settings 231
 - unloading 232
- event archive, scheduled 230
- Event Input/Output 238
- event storage, remote 236
- events
 - search 71
- export
 - search results 119
- exporting remote management configuration 364

F

- factory settings, restoring 475
- field query
 - indexing fields 120
- field set, search 72
- fields, indexing 120
- File events to ESM 254
- Filter 270
 - copying 271
 - creating 270
 - deleting 272
 - editing 271
- Filter, Report Category 222
- filter, search 72
- filtering information on UI page 362
- filters, system 124

- find, events 71
- FIPS 140-2
 - enabling on Connector Appliance 328
 - enabling on container 378
- Firefox (web browser) 62
- Forwarder 246
 - creating 247
 - deleting 250
 - editing 249
- Forwarder status 310, 346
- forwarding file events to ESM 254
- function tabs 63

G

- gateway, default 304
- gauge range 65
- gauges 63
- gid 317
- Global Settings 333

H

- health, system 127
- help 63
- hosts
 - adding 368
 - definition 367
 - deleting 372
 - editing 372
 - moving to different location 372
 - scanned 368
 - scanning 370
 - software-type 367
 - upgrading remotely 372
 - viewing all 367
- Hosts file 303

I

- i18n options 65
- importing remote management configuration 365
- indexing fields 120
- initialization 24
- insp status 310, 346
- interface homing 305
- Internal Storage Group 16, 233
- Internet Explorer (web browser) 62
- intrusion monitoring, reports for 136
- invalid certificate errors 385
- IP addresses
 - assigning 24
 - changing 304

L

- Localhost 367
- localhost 303
- locations
 - adding 364
 - definition 363
 - deleting 366
 - editing 366
 - viewing all 363
- Logfu utility 388

- Logger
 - rebooting 302
- login 61
- Login Settings 333
 - changing 333
- logout 63, 65
- logout, automatic 66
- Logs 313, 348
- logs, internal 296
 - retrieving 296

M

- maintenance mode 287
- Manager 22, 56, 59, 246, 252
- Monitor tab 66
- multi-homing 305

N

- navigation 63
- network interfaces 414
- Network Settings 304
 - changing 304
- network speed 304
- NFS, configuring 390
- NTP Server 307
- NTP setting 306, 308, 346

O

- online help 63
- options 63, 65

P

- packaging connectors
 - advanced mode 420
 - basic mode 419
- parameter value groups, in reports 212
- parameters
 - in report queries 205
 - quick run report 158
 - run reports 160
- parser override 419
- Password policy
 - changing 334
- Password, changing 345, 355, 356
- PCI Storage Group 16, 233
- Peer Logger 280
 - adding 282
 - deleting 284
- peer Logger, searching 111
- postgresql status 310, 346
- predefined filters 124, 125
- Process Status 310, 346
- Protect 724 419

Q

- queries
 - in reports 181
- query
 - events 71
- query controls 63

R

- RADIUS authentication 335
- RAID controller status 310, 323, 346
- range, gauge 65
- rebooting Logger 302
- Receiver 239
 - creating 241
 - deleting 242
 - editing 241
 - types 239, 242
- Receiver status 310, 346
- refreshing UI screen 362
- regular expressions (regex)
 - predefined 125
 - tutorial 462
- Remote Authentication Dial-In User Service (RADIUS) 335
- remote event storage 236
- remote file system mount
 - adding 317, 320
 - deleting 316, 318
 - editing 316, 318, 321
- remote management configuration 364
 - exporting 364
 - importing 365
- remote upgrade 436
- Report Category Filter 222
- reports 275
 - access rights 180
 - administration 220
 - categories 134
 - configuration monitoring 137
 - creating new 165
 - dashboard 140
 - delivery options 162
 - designing 164
 - editing 179
 - e-mailing 162
 - exporting 163
 - file formats 161
 - foundation 135
 - groups 134
 - intrusion monitoring 136
 - navigating to 133
 - parameter value groups 212
 - PCI solution add-on 137
 - publishing 161
 - query parameters 205
 - quick run parameters 158
 - remove scheduled 216
 - run parameters 160
 - running 154
 - SANS Top 5 135
 - saving 163
 - scheduling 215
 - solution add-ons 137
 - template styles 214
 - user-created 138
 - viewing published 164
 - viewing, editing schedules 216
- repositories, user-defined 442
- reset to factory settings 475
- restore to factory settings 475
- restoring a SAN 321

Retrieve Logs 296

S

safety precautions 23

SAN Storage 318

SAN, restore 321

SANS Top 5, reports for 135

saved

filter 123

search 123

Saved Search 273

adding 273

deleting 274

editing 274

Saved Search Files 278

Saved Search Job 275

adding 275

deleting 278

editing 277

scan a host 368, 370

scheduled event archive 230

Scheduled Task 269

currently running 270

finished 270

scheduling

export of search results 111

reports 215

SCP file receiver 242

search

constraints 72

events 71

exporting results 119

field set 72

filter 72, 123

peer Loggers 111

query, defining a 72

results, scheduling export of 111

saved 123

system filters 124

time range 72

Search Group Filter

associating with user group 273

report category filter 222

Search Group Filters 272

Search Results tab 112

servers status 310, 346

severity level 511, 513, 517

SFTP file receiver 242

SmartConnectors 21, 57, 511, 513

batching 511

defined 390

scanner 516

zones 513

SmartMessage 239

software-type host 367

solutions

reports 137

speed, network 304

SSL

Certificate Signing Request 324

SSL Settings 323

Static Route 308

adding 308

statistics 63

status

3Ware RAID Controller 310, 323, 346

process 310, 346

Storage 233

Storage Area Network 318

Storage Group 233

adding 233

Default 32

editing 234

Storage Group, default 16, 32

Storage Group, internal 16, 233

Storage Group, PCI 16, 233

Storage Rule 32, 235

adding 235

deleting 236

editing 236

Storage Settings 236

Storage Volume 236

settings 237

streaming SmartConnector 247

subnet mask 304

supported connectors 390

System Admin tab 301

system definition 361

system filters 124, 125

system health, monitoring 127

System Information 314

System Reboot 302

System Update 309

T

template styles

for reports 214

time configuration 306, 308, 346

time range, dynamic 93

time range, search 72

Time Settings 306

changing 306, 308, 346

time, changing 306

timeout, automatic 333

timezone 307

U

uid 317

Unicode options 65

update, content 437

updating container properties 376

upgrade

Connector Appliance 436

host 436

remote 436

US-ASCII encoding 242

User 337, 343, 349, 354

changing password 345, 356

creating 343, 354

deleting 344, 355

editing 344, 355

User Group 337, 349

associating with Search Group Filter 273

creating 341, 352

deleting 342, 354

editing 341, 353

user interface 63

- filtering information to display 362
- refresh 362
- Search Results tab 112
- User password, changing 345, 356
- user-defined repositories 442
- UTF-8 encoding 242

V

- version, component 309

W

- web browser requirements 62
- web status 310, 346
- What's New 22
- widgets
 - in report dashboards 150

