

Administrator's Guide

ArcSight Logger™ v4.0 GA

November 3, 2009



Administrator's Guide ArcSight Logger™ v4.0 GA

November 3, 2009

Copyright © 2006-2009 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
11/03/09	v4.0 GA	v4.0 GA release
07/22/09	v3.0 SP1	Service pack 1 for version 3.0.
01/07/09	v3.0 Patch 1	v3.0 Patch 1 version.
10/16/08	v3.0	v3.0 GA version.
06/30/08	v2.5.1	v2.5 SP1 release.
05/29/08	v2.5	v2.5 GA release.
11/19/07	v2.0 Patch 1	v2.0 guide; added Report Category Filter section.
10/19/07	v2.0	v2.0 guide.
10/17/07	v2.0	v2.0 guide.
5/24/07	v1.8	First draft with new v1.8 features.
5/4/07	v1.8 pre-release	International support.
4/24/07	v1.5 release	New v1.5 features and general improvements.
3/28/07	v1.1 release	New v1.1 features.

Document template version: 1.0.2.9

ArcSight Customer Support

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

Contents

About the Online Help	xi
Chapter 1: Overview	1
Introduction	1
Logger Features	3
Storage Configuration	3
Receiver Configuration	3
Analyzing Events	4
Grouping Events	4
Exporting	5
Forwarder Configuration	5
User Management	5
Other Setup and Maintenance	6
Deployment Scenarios	6
What's New in Version 4.0	8
Chapter 2: Installation	15
Hardware Device	15
Installation	15
ArcSight Logger Package Contents	15
Safety Precautions	16
Initialization	16
Connecting to the Command Line Interface	17
Using a Browser to Initialize Logger	17
Other CLI Commands	18
Deployment Planning	19
Storage Strategy	19
Retention Policy	19
Initialization Sequence	20
1 SAN	21
2 Storage Volume	21
3 Storage Groups	21
4 Time Settings	22
5 Index Fields	22

6 Reboot	23
7 Receivers	23
8 Devices	23
9 Device Groups	24
10 Storage Rules	24
Installing SmartConnectors to Send Events to Logger	24
SmartMessage	25
Forwarding Logger Events to an ESM Manager	26
Configuring SmartConnectors to Send Events to Both Logger and an ESM Manager	26
Sending Events from ArcSight ESM to Logger	27
Configuring SmartConnectors for Failover Destinations	28
Chapter 3: Using the User Interface	31
Logger User Interface	31
Browser Requirements	31
Navigating the User Interface	32
Help	32
Options	33
Logout	33
Monitor	33
Platform	35
Network	35
Logger	36
Receivers	36
Forwarders	36
Storage	36
Chapter 4: Searching and Analyzing Events	39
The Need to Search Events	39
The Process of Searching Events	40
Elements of a Search Query	40
Syntax Reference for Query Expression	49
Using the Search Builder Tool	53
Accessing Search Builder	53
Nested Conditions	56
Alternate Views for Query Building in Search Builder	56
Search Analyzer	57
Searching for Events on Logger	59
Searching Peer Loggers (Distributed Search)	60
Understanding the Search Results Display	61
Auto Updating Search Results	62
Exporting Search Results	62
Scheduling an Export Operation	63

Indexing	63
How indexing works	63
Full-text Indexing (Keyword Indexing)	64
Field-based Indexing	64
Guidelines for Field-based Indexing	66
Enabling Indexing	66
Saving Queries (Saved Filters and Searches)	67
Saving a Query	67
Using a Saved Filter or a Saved Search	68
System Filters/Predefined Filters	68
Using a System Filter	69
Monitoring System Health	69
System Health Events	70
Advanced Search Options	71
Alerts	71
Viewing Alerts	71
Receiving Alerts for Events	72
Base Event Fields	73
Go, Export, and Auto Update Options	73
Chapter 5: Reporting	75
Navigating to Reports	75
Report Groups	76
Foundation Reports	77
SANS Top 5 Reports	77
Network Monitoring Reports	78
Intrusion Monitoring Reports	78
Configuration Monitoring Reports	79
Solution Reports	79
Device Monitoring Reports	79
Anti-Virus Reports	80
Cross Device Reports	80
Database Reports	80
Firewall Reports	80
Identity Management Reports	80
IDS-IPS Reports	80
Network Reports	80
Operating System Reports	80
VPN Reports	80
User Reports	80
Reports Home Page	81
Using the Dashboard	82
Viewing the Dashboard	82

Designing Dashboards	83
What items can a dashboard include?	84
Quick Start - Creating a New Dashboard	84
Add an Empty Dashboard	85
Creating Widgets	87
Placing Dashboard Items on the Layout	87
Placing a Report on a Dashboard	87
Placing a Use Case on a Dashboard	90
Placing an External Link on a Dashboard	92
Swapping Items on Widgets	93
Setting Dashboard Preferences	94
Working with Available Dashboards	94
Selecting a Dashboard View	94
Modifying or Removing Existing Dashboards	95
Running, Viewing, and Publishing Reports	95
Best Practices	96
Finding Reports	96
Task Options on Available Reports	96
Running and Viewing Reports	98
About the Pagination of Reports	98
Quick Run / Run In Background Report Parameters	99
Run Report Parameters	101
Report File Formats	102
Publishing Reports	103
Report Delivery Options	104
Refreshing a Report	104
E-mailing a Report	104
Exporting and Saving a Report	105
Viewing the Output of a Published Report	106
Designing Reports	106
Opening the Report Designer	107
Creating New Reports	107
Quick Start: Base a New Report on an Existing One	107
Designing New Reports	110
Select Filter Criteria	112
Select Grouping	114
Select Totals	116
Sort Order	116
Highlighting	117
Create Matrix	117
Create Chart	118
Editing a Report	120
Adhoc Report Designer	121


Setting Access Rights on Reports	122
Setting up Queries	122
How Search and Report Queries Differ	123
Overview of Query Design Elements	123
Creating a Copy of an Existing Query	124
Designing a New SQL Query	124
Modifying a Query Object	137
Deleting a Query Object	137
Defining SQL in the Editor	137
Working with Parameters	145
Creating New Parameters	146
Modifying a Parameter	151
Deleting a Parameter	151
Configuring Parameter Value Groups	152
Applying Report Template Styles	154
Defining a New Template	155
Scheduling Reports	155
Viewing and Editing Scheduled Reports	155
Scheduling a Report	156
Deploying a Report Package	159
Report Server Administration	160
Using Report Category Filters	162
Backup and Restore of Report Content	162
Chapter 6: Configuration	163
Configuration	163
Devices	163
Devices	164
Device Groups	165
Event Archives	166
Event Archives	166
Scheduled Event Archive	167
Archive Storage Settings	168
Storage	168
Storage Groups	168
Storage Rules	170
Storage Volume	172
Event Input/Output	173
Receivers	173
Forwarders	179
ESM Destinations	183
Forwarding Log File Events to ESM	186
Alerts	187

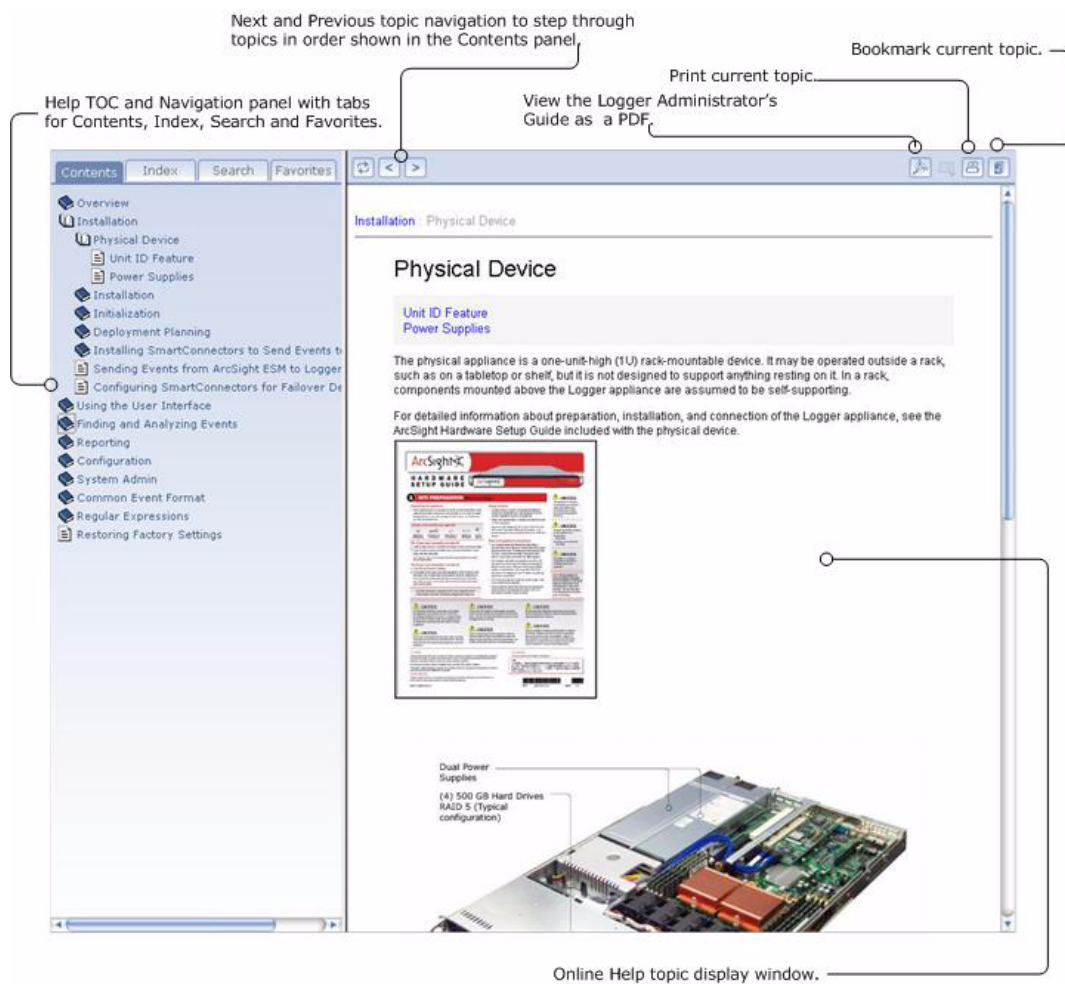
SNMP Destinations	191
Syslog Destinations	191
ESM Destinations	192
Export	193
Scheduled Tasks	193
Scheduled Tasks	193
Currently Running Tasks	194
Finished Tasks	194
Filters	194
Filters	194
Search Group Filters	195
Export	197
Saved Searches	197
Saved Search	197
Saved Search Jobs	198
Saved Search Files	200
Search Optimization	200
Add Search Indexes	200
Tuning Advanced Search Options	201
Deleting Custom Field Sets	202
Peer Loggers	202
Guidelines	203
Configuration Backup and Restore	206
Running a Configuration Backup (Ad-hoc or Scheduled)	207
Restoring from a Configuration Backup	208
Editing Configuration Backup Settings	208
System Maintenance	208
Database Defragmentation	210
Retrieve Logs	215
Exporting and Importing Content	215
Guidelines for Exporting and Importing	216
Exporting Content	217
Importing Content	218
Chapter 7: System Admin	219
System Admin	219
Reboot	219
DNS Settings	220
Hosts	221
Network	221
Time/NTP	223
Impact of Daylight Savings Time Change on Logger Operations	225
SMTP Settings	226

Static Routes	227
License & Update	227
Process Status	228
Support Login	229
Logs (Audit and Error)	230
Audit Forwarding	230
Storage	231
CIFS Settings	231
Network File System (NFS) Settings	233
SAN	235
Restoring a SAN	237
RAID Controller	237
Security	238
SSL Server Certificate	238
SSL Client Authentication (CAC Authentication)	240
FIPS 140-2	242
Users/Groups	247
Authentication Settings	247
Groups	251
Users	256
Changing Password	257
Appendix A: Common Event Format	259
Common Exchange Format	259
Common Extension Dictionary	261
Appendix B: Regular Expressions	265
Regex Overview	265
Simple Regular Expressions	265
Metacharacters	266
Forbidden Characters	271
Things To Remember	271
Appendix C: Restoring Factory Settings	273
Appendix D: Logger Audit Events	281
Types of Audit Events	281
Information in an Audit Event	281
Platform Events	282
Logger Application Events	284

Appendix E: Event Field Name Mappings	289
Appendix F: Connector Appliance Documentation	295
Appendix G: Logger Search From An ESM Console	297
Understanding the Integrated Search Functionality	297
Prerequisites	298
Setup and Configuration	298
ESM	298
Logger	298
Supported Search Options	299
Guidelines	299
Searching on Logger From ESM Console	299
Index	301






About the Online Help

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




The online Help includes the following features:

- Left panel Help navigation - Click a tab for Contents (TOC), Index, Search, or Favorites.

- Next, Previous topic navigation- Click the Previous button () to view the preceding topic in the history, or the Next button () to view the subsequent topic.
- Topic display window - Click a topic in the Contents, Index, Search hit list, or saved Favorites to view it in the display window.
- Access to the Administrator's Guide as an Adobe Acrobat PDF document. Click the PDF button () in the upper right of the Online Help toolbar to get PDF.
- Print capabilities - Click the Print () button to print a copy of the current topic.
- Bookmarks - Click the Bookmark () button and follow the instructions in the popup window to bookmark a topic.



On Logger platforms that have ArcSight's Connector Appliance integrated, context-sensitive Help is not available for the Connector Appliance UI screens. However, when you click the Help icon () in the upper right-hand corner, the Connector Appliance Administrator's Guide is displayed. This guide contains all the information you need to know to configure and use the integrated Connector Appliance.

Chapter 1

Overview

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

["Introduction" on page 1](#)

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

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

["What's New in Version 4.0" on page 8](#)

Introduction

ArcSight Logger is an appliance-based 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.

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

	Time (Event Time)	Device	Logger	deviceVendor	deviceProduct	deviceVersion	deviceEventClassId	name	baseEventCount	deviceCustomNumber1	deviceCustomNumber1L
1	2009/10/27 12:14:58 PDT	Logger	Local	ArcSight	Logger	4.0.0.4067.0	sensor:102	Logger Internal Event	1	24	value
	RAW CEF:0 ArcSight Logger 4.0.0.4067.0 sensor:102 Logger Internal Event 1 cat=Monitor/Sensor/System cs2=CurrentValue cnt=1 cs3=ok cs1=null type=0 cnd=24 cs1Label=unit rtd=1256670										
2	2009/10/27 12:14:58 PDT	Logger	Local	ArcSight	Logger	4.0.0.4067.0	sensor:111	Logger Internal Event	1	3600	value
	RAW CEF:0 ArcSight Logger 4.0.0.4067.0 sensor:111 Logger Internal Event 1 cat=Monitor/Sensor/FAN1 cs2=CurrentValue cnt=1 cs3=ok cs1=null type=0 cnd=3600 cs1Label=unit rtd=1256670										
3	2009/10/27 12:14:58 PDT	Logger	Local	ArcSight	Logger	4.0.0.4067.0	sensor:112	Logger Internal Event	1	3600	value
	RAW CEF:0 ArcSight Logger 4.0.0.4067.0 sensor:112 Logger Internal Event 1 cat=Monitor/Sensor/FAN2 cs2=CurrentValue cnt=1 cs3=ok cs1=null type=0 cnd=3600 cs1Label=unit rtd=1256670										
4	2009/10/27 12:14:58 PDT	Logger	Local	ArcSight	Logger	4.0.0.4067.0	sensor:113	Logger Internal Event	1	3600	value
	RAW CEF:0 ArcSight Logger 4.0.0.4067.0 sensor:113 Logger Internal Event 1 cat=Monitor/Sensor/FAN3 cs2=CurrentValue cnt=1 cs3=ok cs1=null type=0 cnd=3600 cs1Label=unit rtd=1256670										
5	2009/10/27 12:14:58 PDT	Logger	Local	ArcSight	Logger	4.0.0.4067.0	sensor:114	Logger Internal Event	1	3600	value
	RAW CEF:0 ArcSight Logger 4.0.0.4067.0 sensor:114 Logger Internal Event 1 cat=Monitor/Sensor/FAN4 cs2=CurrentValue cnt=1 cs3=ok cs1=null type=0 cnd=3600 cs1Label=unit rtd=1256670										
6	2009/10/27 12:14:58 PDT	Logger	Local	ArcSight	Logger	4.0.0.4067.0	sensor:115	Logger Internal Event	1	3600	value
	RAW CEF:0 ArcSight Logger 4.0.0.4067.0 sensor:115 Logger Internal Event 1 cat=Monitor/Sensor/FAN5 cs2=CurrentValue cnt=1 cs3=ok cs1=null type=0 cnd=3600 cs1Label=unit rtd=1256670										
7	2009/10/27 12:14:58 PDT	Logger	Local	ArcSight	Logger	4.0.0.4067.0	sensor:119	Logger Internal Event	1	0	value

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.

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

Syslog is a loose standard (characterized, not defined, in RFC 3164) for event messages. Although Logger is message-agnostic, it can do more with messages that adhere to the Common Event Format (CEF), an industry standard for the interoperability of event- or log-generating devices. (See [Appendix A, Common Event Format](#), on page 259 for more information.)

Any text can be treated as an event. 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.

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 202](#)
- [“Common Event Format” on page 259](#)

Logger Features

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

Storage Configuration

Logger includes onboard storage for events. Some Logger models include RAID 1 or RAID 5 storage systems. (See the *Appliance Specifications* document for more details. This document is available on the ArcSight Customer Support web site at <https://support.arcsight.com>.) On Logger models that support a Storage Area Network (SAN), you need to use the SAN for storage.

Events are stored compressed. You can not configure the compression level.

Use of a Network File System (NFS) as primary storage for Logger events is not recommended. However, an NFS or a CIFS system can be used for archiving Logger data.



Starting with v4.0 GA, Logger can mount Common Internet File System (CIFS) shares. However, a CIFS share can only be used to archive Logger data such as event archives, saved searches, exported filters and alerts, and configuration backup information. You can also configure the Logger to read event data or log files from a CIFS host.

Logger can interact with Network Attached Storage (NAS) or with a Storage Area Network (SAN) using a SAN gateway, as shown in [Figure 1-2](#).

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. Storage Groups and retention policies can be changed, but new Storage Groups cannot be added later, however, a Storage Group's size can be increased or decreased.

- [“Storage Strategy” on page 19](#)
- [“Initialization Sequence” on page 20](#)
- [“Storage” on page 168](#)
- [“Storage” on page 231](#)

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 are assumed to contain one event per line, including a timestamp for each event. 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 models, 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 173](#)
- [“Installing SmartConnectors to Send Events to Logger” on page 24](#)
- [“Sending Events from ArcSight ESM to Logger” on page 27](#)

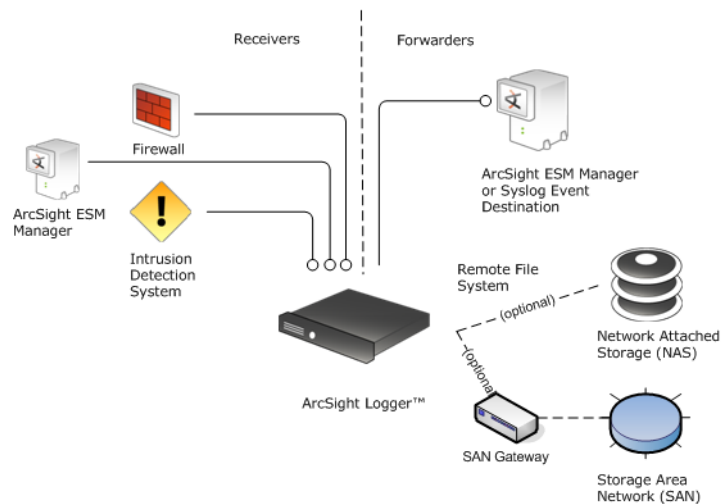


Figure 1-2 Logger has two network interface cards (NICs) so that it can receive events on one subnet and forward events on a different subnet.

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.

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

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

- [“Searching for Events on Logger” on page 59](#)
- [“Saving Queries \(Saved Filters and Searches\)” on page 67](#)
- [“Filters” on page 194](#)
- [“Saved Searches” on page 197](#)
- [“Peer Loggers” on page 202](#)

Grouping Events

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

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

Storage Rules associate a Device Group with a Storage Group. Storage Rules are ordered by priority, and the first matching rule determines to which Storage Group an incoming event will be sent.

Device Groups, Devices, Storage Groups, and Peer Loggers can each be used to filter events using Search Constraints, which can be specified interactively on the Analyze page as well as when creating Filters or Saved Searches.

- [“Devices” on page 163](#)
- [“Storage Rules” on page 170](#)
- [“Searching Peer Loggers \(Distributed Search\)” on page 60](#)

Exporting

Logger can export events that match the current query locally, to an NFS mount, a CIFS mount, a SAN (on specific Logger models), or to the browser as a file to be downloaded. Events are saved in Comma-Separated Values (CSV) format for easy processing by external applications. 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 and an NFS mount, a CIFS mount, or a SAN (on some Logger models) and file path.

- [“Exporting Search Results” on page 62](#)
- [“Impact of Daylight Savings Time Change on Logger Operations” on page 225](#)
- [“Saved Search Jobs” on page 198](#)

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 179](#)
- [“ESM Destinations” on page 183](#)

User Management

User accounts can be created by the Logger administrator to distinguish between different users of the system. User accounts inherit privileges from the User Group to which they

belong. User Groups can have an enforced event Filter applied to them, limiting the events that a specific user can see.

- [“Users” on page 256](#)
- [“Changing Password” on page 257](#)
- [“User Groups” on page 251](#)
- [“Search Group Filters” on page 195](#)

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 can be rebooted using controls in the browser user interface.

Various other system settings can be modified. Most require a System Reboot for the changes to take effect.

- [“Configuration Backup and Restore” on page 206](#)
- [“Retrieve Logs” on page 215](#)
- [“Storage” on page 231](#)
- [“Reboot” on page 219](#)
- [“License & Update” on page 227](#)
- [“Network” on page 221](#)

Deployment Scenarios

Typically, Logger is deployed inside the perimeter firewall with a high degree of physical security for the appliance itself to prevent tampering with the collected network security 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 security 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.

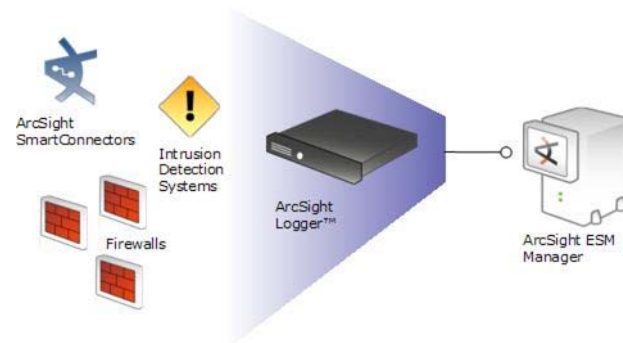


Figure 1-3 Logger can act as a funnel, forwarding selected events to ArcSight ESM.

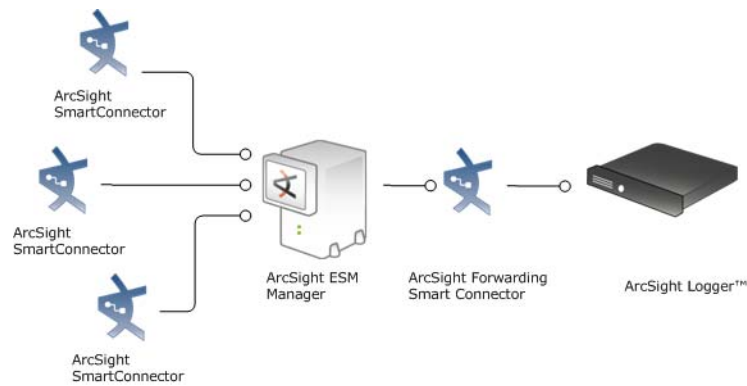


Figure 1-4 Logger can store events sent by ArcSight ESM.

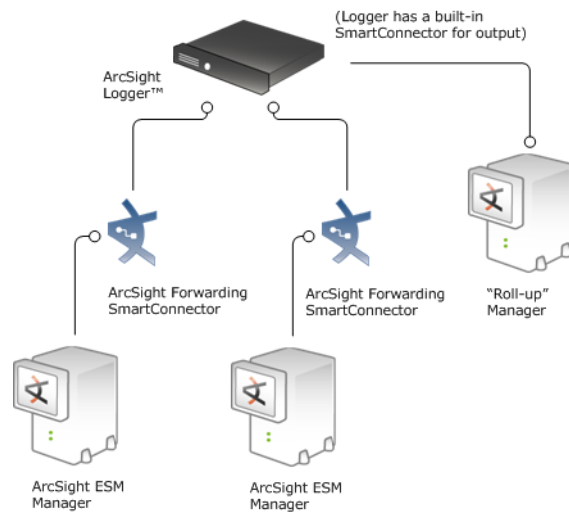


Figure 1-5 Logger can store and forward filtered events in a hierarchical ArcSight Manager deployment.

What's New in Version 4.0

This section lists the new features/enhancements introduced in the Logger v4.0 release. Also, refer to the Release Notes for this release for a list of fixed and open bugs, and late-breaking information.

■ Next Generation Hardware Platform

Logger v4.0 GA runs on the new Logger hardware platforms available from ArcSight. The new platforms (L3200, L7200, and L7200-SAN) are the next-generation Logger hardware systems for the existing platforms available from ArcSight.

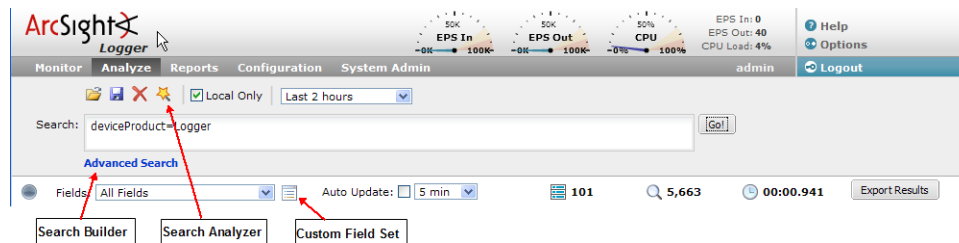
■ Enhanced Search function

Logger v4.0 introduces significant usability and functionality improvements for searching events. The Search user interface for field-based and regular expression queries has been unified to simplify user interaction. The queries for these search methods can be part of a single query expression, in which the field-based query searches for matching events and the regular expression query helps further refine search results.

Additionally, **a third search method called full-text search (also known as "keyword search") has been introduced**. When using this search method, you enter queries in plain English, as you would when using any of the popular Internet search engines.

The full-text search queries can be specified standalone or in conjunction with the field-based and regular expression queries. For example, a simple full-text search query searches for the word "failed" in the events stored on Logger, while a more complex version, which combines the three search methods, can be as follows:

```
failed AND name="*[4924TestAlert]*" AND ("192.168.*" OR
categoryBehavior CONTAINS Stop) | REGEX=":\d31"
```

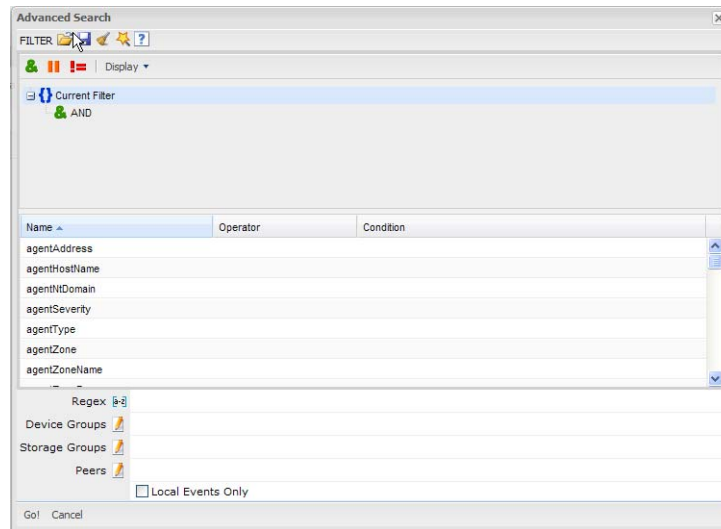


Other significant improvements related to search are:

◆ Search Builder tool (Advanced Search tool)

This tool is a boolean-logic conditions editor that enables you to quickly and accurately build keyword, field-based, and regular expression search queries, as shown in the following figure. The tool provides a visual representation of the conditions you are including in a query. For more information, see ["Using the Search Builder Tool" on page 53](#).

To access this tool, click **Advanced Search**, right below the Search text box.



◆ Search Analyzer tool

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.

◆ Auto-suggest facility

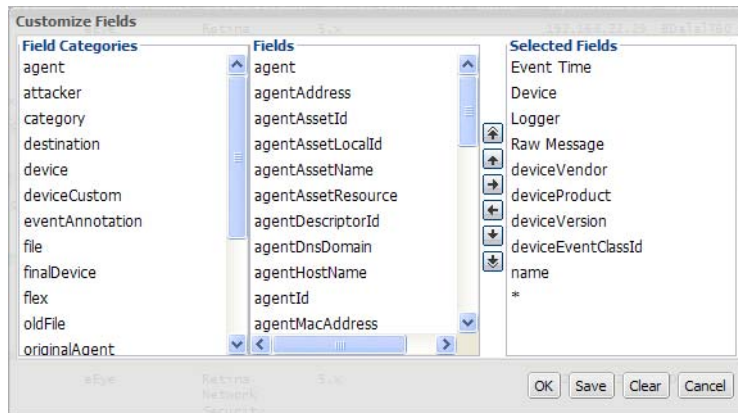
Although you can use the Search Builder tool to build your queries, if you choose to type them in the Search text box, Logger's auto-suggest facility enables you to quickly build query expression by automatically providing suggestions, possible matches, and applicable operators for the following:

- Fields in Logger schema
(See ["Indexing" on page 63](#) for a complete list of fields.)
- Metadata terms (`_storageGroup`, `_deviceGroup`, `_peerLogger`)
Type `"_s"` (for storage group), `"_d"` (for device group), or `"_p"` (for peerLogger) in the Search text box to obtain a drop-down list of constraint terms and operators.
- Regular expression term (`|REGEX=`)

◆ Custom field sets

A field set determines the fields that are displayed in the search results for each event that matches a search query. Starting with Logger v4.0, you can create your own field sets. By doing so, you can hide the event fields you do not need in the search results, and use the available screen space on your computer more efficiently to only view event fields you need.

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.



◆ Nested conditions

Starting with this release, you can include different boolean operators in a query. That is, you can include AND and OR in a single search query, or you can include AND, OR, and NOT in a single search query, and so on.

Consequently, you can create nested conditions by mixing any of the applicable operators.

For example, you can nest full-text search keywords using boolean operators. Similarly, you can nest field-based search queries using string, boolean, numeric, and so on operators. For example, `(name="John Doe" OR name="Jane Doe") AND message!="success"`.

When nesting conditions, metadata identifiers (`_storageGroup`, `_deviceGroup`, and `_peerLogger`) can only appear at the top level in a query expression.

◆ Drill-down search results

You can drill-down the search results to further refine them. Click a green-highlighted term in the search results 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.

■ Expanded LUN Storage Capacity

Logger platforms that support SAN now have the ability to support up to 5 TB volumes. To make use of this enhanced capacity, make sure you allocate a 5.4 TB LUN when initializing your Logger. The additional 0.4 TB need to be allocated to accommodate Logger system files.

Note: 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.

■ Storage Groups—Additional and Resizing

Logger can have a maximum of 6 storage groups now—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. (In Logger v3.x, you could only create two more storage groups in addition to the pre-existing ones, thus resulting in three for event storage.)

The storage groups need to be created during the Logger initialization phase and cannot be created or deleted once Logger has been initialized, as described in [“Initialization Sequence” on page 20](#). However, a Storage Group’s size can be increased and decreased any time (if there is sufficient disk space available on the Logger to perform the operation); therefore, create additional groups of minimal size even if you don’t need them at this point.

For more information, see [“Storage Groups” on page 168](#).

■ NFS and CIFS

Starting with this release, Logger can also mount a CIFS remote file system (Windows share) **for secondary storage**—for archiving data such as events, exported filters and alerts, and saved searches.

Starting with Logger v4.0, use of NFS as the primary storage device for storing Logger events is **not** recommended. However, you can continue to use NFS mounts to archive data such as events, exported filters and alerts, and saved searches.

For more information, see [“Storage” on page 231](#).

■ Configuration Audit Events

Logger now generates configuration audit events that provide greater visibility and change control for monitoring the activity of users and administrators who have access to Logger. For example, storage volume has been added, certificate added, and search indices added.

For more information, see [Appendix D, Logger Audit Events, on page 281](#).

■ Expanded Logger Schema

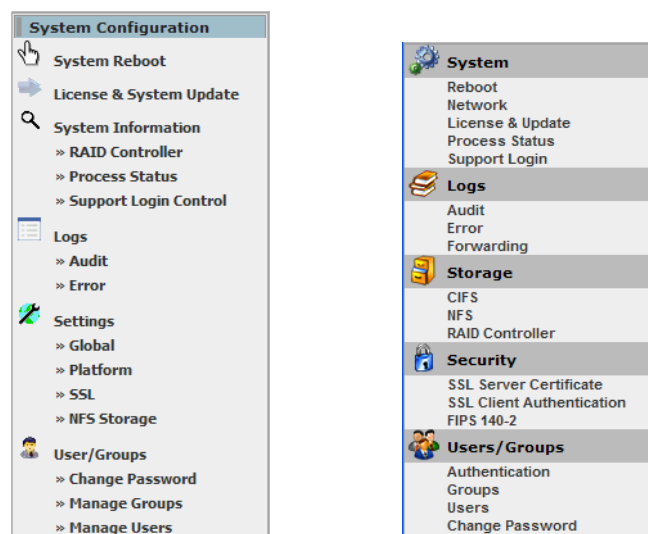
Fifteen additional fields such as `customerName` and `requestUrl` have been added to the Logger schema for enhanced searchability and reporting.

The `requestUrl` field cannot be indexed.

For a complete list of fields, see [“Indexing” on page 63](#).

■ System Admin Menu Changes

The System Admin menu has been reorganized to improve usability. The following figure illustrates the before (left figure) and after (right figure) layout. The system administration procedures in [Chapter 7, System Admin, on page 219](#) have been updated to reflect these changes.



■ Reporting Changes

The reporting section introduces following enhancements:

- ◆ User-definable templates: In addition to using pre-defined templates, now you can create your own templates.
- ◆ Search: Queries, Parameters, Parameter Value Groups, and templates include a search function to search for an existing element. For example, you can search for a query by specifying the letter its name starts with or by specifying a letter or a word that it contains in its name.
- ◆ Point Labels: When creating charts, you can include point labels that indicate the number of matches for each plotted element on the X-axis.
- ◆ Default Dashboard: If a dashboard is not defined, a Report Execution Status page indicates the status of recently run or accessed reports.
- ◆ Device Monitoring Reports: A new report group that address common device monitoring user cases for systems and devices on your network.
- ◆ Standard and Custom Reports: These reports are not listed on separate tabs. Instead, they are combined on a single screen.
- ◆ Run in Background: Reports that take longer to run or are not required immediately can be run in the background. This option enables you to perform other activities on Logger while a report is being generated.
- ◆ Report Title: Each report can be given a meaningful title. This title appears at the top of the report when it is generated.
- ◆ Adhoc Report Designer: Once you edit a report, you can run a report before saving it to ensure that the report output is as you expected.
- ◆ SQL Editor: The Design tab in the SQL Editor has been enhanced to improve usability.

For more information, see [Chapter 5, Reporting, on page 75](#).

■ Security Enhancements

The following security enhancements have been made in this release:

FIPS Support

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.

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.

FIPS mode is supported on SmartConnectors running version 4.7.5.5372 or later. Follow instructions in [“Installing or Updating a SmartConnector to be FIPS-compliant” on page 245](#) to ensure that your connector is FIPS compliant.

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.

CAC Support

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.

Support Login

Starting with this release, when ArcSight Customer Support needs access to your appliance for troubleshooting and diagnostics, they work with you to assign a single-use password to the appliance. Doing so enables Support Login access to the appliance. This password is valid only for one support session and is automatically disabled after the session ends.

For more information, see [“Support Login” on page 229](#).

■ ESM Manager as an Alert Destination

Alerts can be forwarded to an ESM Manager, in addition to the previously supported SNMP, Syslog, and E-mail destinations. For more information, see [“ESM Destinations” on page 183](#).

■ Performance Enhancements

This release includes many system- and application-level enhancements to improve performance and reliability of various Logger operations. For example,

- ◆ As a result of file system optimization in this release, pre-allocation now takes significantly lesser time than before.
- ◆ The enhanced Logger forwarding software enables you to pause and resume forwarding 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. For more information, see [“Forwarders” on page 179](#).
- ◆ Due to Search optimizations implemented in this release, search queries that contain multiple search components—field-based, regular expression—run faster than before.

■ Documentation—System Health Events

The system health events that Logger generates are stored in the Internal Storage Group, which is created by default. Until now, these events were not documented. Starting with this release, a list of these events is available in this document, Logger Administrator's Guide. See [“System Health Events” on page 70](#) for more information.

■ Documentation Access

This information only applies to you if your Logger platform is a Logger-Connector Appliance integrated solution.

Connector Appliance documentation is now available as follows:

- ◆ Through the Help icon (?) on any user interface page, when you are in the Connector Appliance context. When you click this icon, a PDF of the Connector Appliance Administrator's Guide is displayed. All information in this guide except system administration is applicable to your product.
- ◆ Through the ArcSight Customer Support site at <https://support.arcsight.com>.

Chapter 2

Installation

This chapter describes installing the Logger appliance and initializing the hardware for a specific network environment.

In this chapter:

[“Hardware Device” on page 15](#)

[“Installation” on page 15](#)

[“Initialization” on page 16](#)

[“Deployment Planning” on page 19](#)

[“Installing SmartConnectors to Send Events to Logger” on page 24.](#)

[“Sending Events from ArcSight ESM to Logger” on page 27](#)

[“Configuring SmartConnectors for Failover Destinations” on page 28](#)

Hardware Device

The Logger appliance is a rack-mountable device. It may be operated outside a rack, such as on a tabletop or shelf, but it is not designed to support anything resting on it. In a rack, components mounted above the Logger appliance are assumed to be self-supporting.

For detailed information about preparation, installation, and connection of the Logger appliance, see the rack installation instructions included with the appliance.

Installation

To install the Logger appliance, follow the instructions in the rack installation instructions, included in the package. Once the hardware is installed and connected, see [“Initialization” on page 16](#) to continue setup and prepare for deployment.

ArcSight Logger Package Contents

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

- Logger appliance chassis (main system)
- Face plate (bezel) if not attached
- One or two power cables (North America)
- Slide rail/rack mount parts kit

- Packaging Checklist, *ArcSight Getting Started Guide* document, rack installation instructions for your platform.

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

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

Safety Precautions

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

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



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

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

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

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

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

Initialization

Before logging in to Logger for the first time, the unit must be initialized with at least one valid IP address. There are three ways to accomplish this:

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

Connecting to the Command Line Interface

To use the Command Line Interface (CLI), attach a terminal to the serial port on Logger or attach a monitor and keyboard. The serial port expects a standard VT100-compatible terminal: 9600 bps, 8-bits, no parity, 1 stop bit (8N1), no flow control.

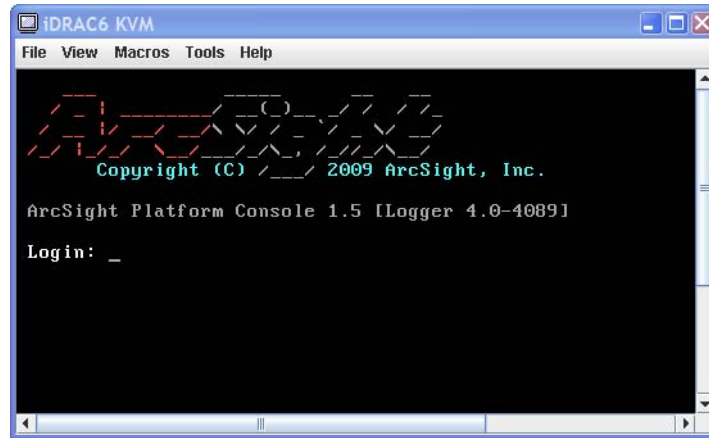


Figure 2-1 The ArcSight Platform Console sign-on. Note that the platform version will not match the current Logger application version.

To Initialize Logger using the CLI

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

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

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

Using a Browser to Initialize Logger

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

<https://192.168.35.35/>

- 2 At the login screen, enter **admin** and password **password**. Logger reminds you to set up a Storage Volume, but that step is described later, in [“Initialization Sequence” on page 20](#). It is very important that you do not specify a Storage Volume or make other critical deployment decisions at this time.
- 3 Click the **System Admin** tab.
- 4 On the sub-menu, click **Platform** (under Settings).
- 5 Click the **Network** tab and enter the desired host name, default gateway, IP address(es) and other information. Click **Update Settings**.



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

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

Other CLI Commands

The following commands are available at the CLI prompt:

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

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

Deployment Planning

Logger initialization requires planning because there are several initial settings which cannot be changed once they are set. The number of Logger appliances is one aspect that does not require pre-planning. Loggers can be added or removed dynamically, except of course that events stored on Loggers that are no longer peer are not available for distributed search.

Storage Strategy

Logger events can be stored either locally or remotely on a Storage Area Network (SAN). SAN storage is available with some Logger models only. If your Logger model supports a SAN, it should be available before you bring the Logger online.

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.

Starting with v4.0 GA, Logger can mount Common Internet File System (CIFS) shares. However, a CIFS share can only be used to archive Logger data such as event archives, saved searches, exported filters and alerts, and configuration backup information. 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 with a different retention policy. Retention policy is specified in terms of number of days that events are stored, or overall maximum size (in GB). Events from specific IP addresses can be routed to particular Storage Groups, making it possible to store all router events, for example, to a Storage Group with short retention, and business-critical host events to another Storage Group with a longer retention.

The Logger receipt time of an event is used to determine the starting time for its retention period.

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

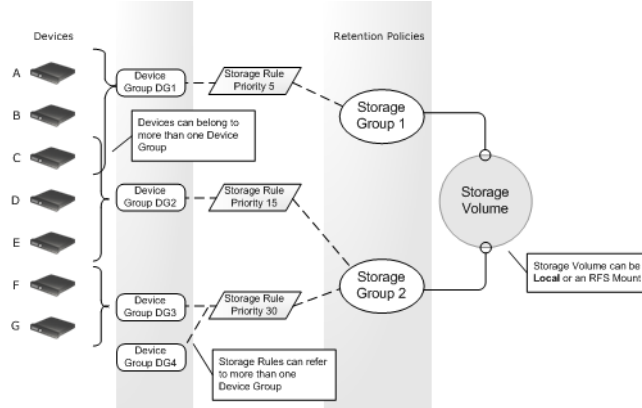


Figure 2-2 Devices participate in Retention Policy

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



Note

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

Initialization Sequence

It is very important that you initialize Logger in the sequence shown here. Logger can be reset to its initial condition, but other than that, several of the settings described here cannot be changed once set.



Caution

One-time initialization can only be changed by performing a factory reset (see [Appendix C, Restoring Factory Settings, on page 273](#)). Be sure you know how you want Logger storage set up before performing the first steps of the initialization sequence (up to rebooting).

This sequence ensures that resources are created and parameters are set in the proper order:

- 1 SAN (on selected Logger models only)
- 2 Storage Volume - establish where Logger stores event data
- 3 Storage Groups - apply retention policies to the Storage Volume
- 4 Time Settings

- 5 Index Fields
- 6 Reboot - commit the changes made in previous steps
- 7 Receivers
- 8 Devices
- 9 Device Groups
- 10 Storage Rules

1 SAN

If you will use Logger's built-in storage, skip this step.

If you are using a SAN as your primary storage, the SAN must be up before initializing the Logger. Logger can attach to only one LUN (on SAN) at a time for primary storage. You can add more LUNs for event archival, configuration backup, and export. See ["SAN" on page 235](#) for instructions. (Only certain Logger models support SANs.)

2 Storage Volume

Establish the Logger's Storage Volume. See ["Storage Volume" on page 172](#). Choose **Local** to use Logger's built-in storage or choose SAN if your Logger model supports SAN. If you will use SAN, enter a folder path to the SAN. This folder path must already exist on the remote storage.

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

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.



Storage Volume cannot be extended after initialization.

3 Storage Groups

Logger can have a maximum of 6 storage groups—two that pre-exist on your Logger (Internal Storage Group and Default Storage Group) and **four that you can create**. As a result, you have five storage groups available for event storage and one for Logger's internal events.

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

Each Storage Group can have a different retention policy.



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

Maximum number of Storage Groups on Logger (including preexisting groups): 6
Storage Groups created by default: 2 (Default Storage Group and Internal Storage Group)

Number of Storage Groups available for event storage: 5

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

Number of Storage Groups you can create: 4

See ["Storage Groups" on page 168](#) for the details of adding Storage Groups.

4 Time Settings

Configure the system time manually. Follow instructions in ["Time/NTP" on page 223](#).

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 223](#) for more information.

5 Index Fields

As shown in the [Figure 2-3](#), during the initialization process, Logger prompts you to add a recommended set of fields to the index. You are not required to index event fields at this point, but ArcSight strongly recommends that you do so. 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**. (For full-text indexing, each event is scanned and divided into keywords and stored on the Logger.) See ["Indexing" on page 63](#) for more information.

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

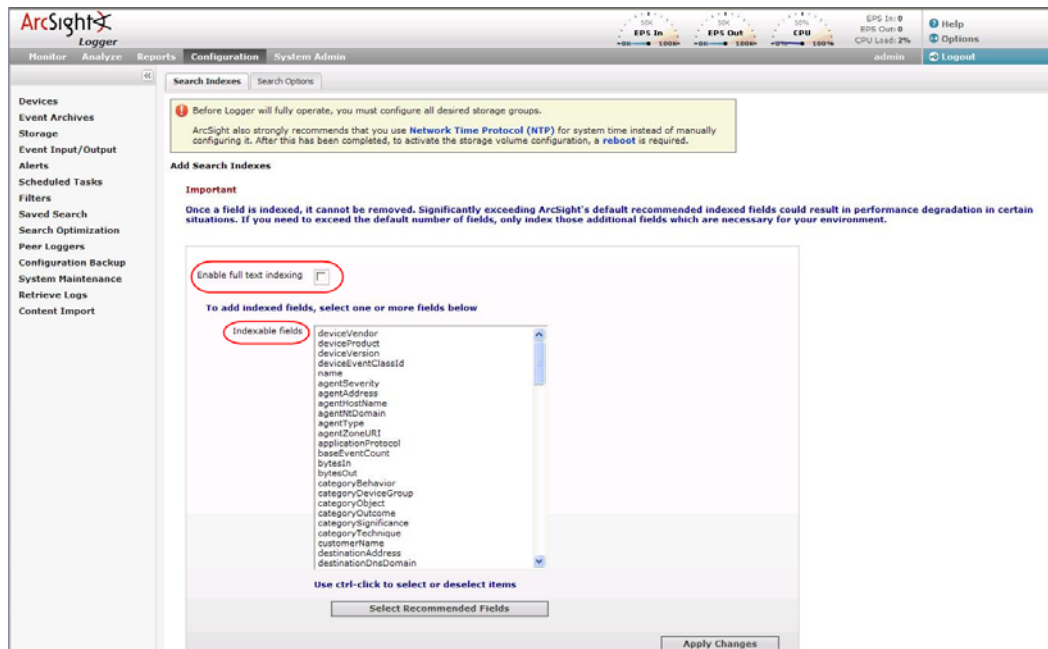


Figure 2-3 Initializing screen that prompts you to select recommended fields to index.

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

6 Reboot

After the Storage Volume and Storage Groups have been created, reboot the system to commit changes before other resources can be created and Logger can begin processing events. See [“Reboot” on page 219](#).



When Logger is rebooted, the Storage Volume and Storage Group settings become permanent. Only certain settings of non-default Storage Groups can be changed.

7 Receivers

Now that you have established a Storage and retention policy configuration for Logger, you can create Receivers to listen for events. Unlike the previous configuration choices in this Initialization Sequence, Receivers can be changed and deleted as needed in the future. Receivers can also be disabled and re-enabled later. For more information about setting up Receivers, see [“Receivers” on page 173](#).

8 Devices

When at least one Receiver is enabled, Logger begins storing events. Using a process called auto-discovery, Logger automatically creates resources called Devices to keep track of source IP addresses and uses DNS to map them to hostnames. Eventually, a Device is created for each device from which Logger received events.

You can also create Devices preemptively, by entering the IP addresses that you expect to be sending events to Logger. You might do this if you don't want to wait for autodiscovery, or if you want to control the initial naming of each Device. (Auto-discovered Devices are

named for their host, or if the DNS lookup fails, for their IP address, and their Receiver.)
For information about manually creating Devices, see [“Devices” on page 164](#).

9 Device Groups

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

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

You can change and delete Device Groups freely as your needs change. Setting up Device Groups initially is not critical; incoming events that are not assigned to a Device Group are automatically sent to the Default Storage Group. For the details of setting up Device Groups, see [“Device Groups” on page 165](#).

10 Storage Rules

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

If you only implemented extra Storage Groups because ArcSight recommended that you do so (back in step 3), then you do not need to create any Storage Rules and you can skip this step. Events from all Devices will be sent to the Default Storage Group and use its specified retention policy.

If you want to implement multiple retention policies, create Storage Rules that associate the appropriate Device Groups with the Storage Groups that implement the correct retention policy. See [“Storage Rules” on page 170](#) for more information.

Storage Rules are tested in order; the first matching rule determines to which Storage Group an event is sent. This approach means that a single Device can belong to several Device Groups without ambiguity about which Storage Group it will end up in.

Installing SmartConnectors to Send Events to Logger

ArcSight Logger is a hardware storage solution optimized for extremely high event throughput. Logger stores time-stamped text messages, called events, at high sustained input rates. Unlike ArcSight SmartConnectors, Logger does not “normalize” events. Events consist of an event time, a receipt time, a source (host name or IP address), and an

un-parsed message portion. Logger compresses raw data, but can also retrieve it in an unmodified form for forensics-quality litigation reporting.

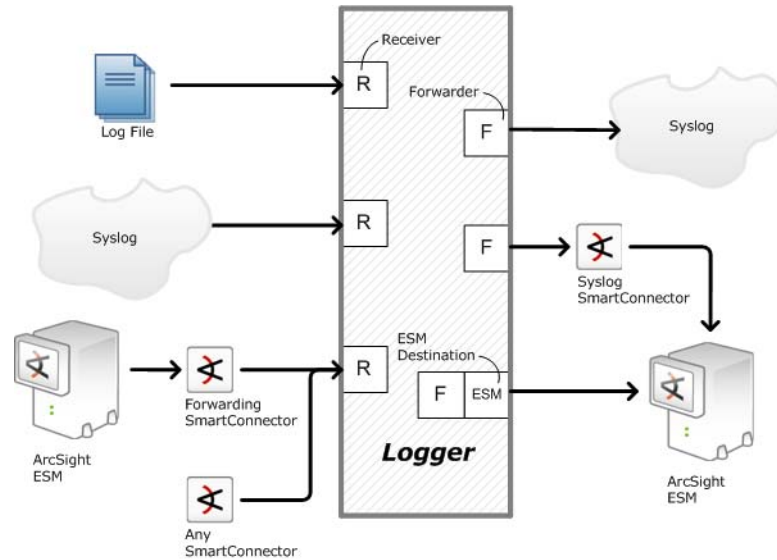


Figure 2-4 ArcSight SmartConnectors interact in a variety of ways with Logger Receivers (R) and Forwarders (F).

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

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

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

SmartMessage

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



Caution

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

Older SmartConnectors will work with Logger, but will 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.



The SmartMessage secure channel uses secure sockets layer (SSL) protocol to send encrypted events to Logger. This is similar to, but different from, the encrypted binary protocol used between SmartConnectors and ESM Manager. Use port 443 (instead of ArcSight's traditional port, 8443) because the secure channel uses SSL.

Set up the SmartMessage Receiver on Logger first (see [“Receivers” on page 173](#)) 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.



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

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

For more information about the Common Event Format (CEF), see [“Common Event Format” on page 259](#).

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.

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



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

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

Use build 4810 or later for compatibility with Logger.

- 1 Install the SmartConnector component normally, but cancel the installation when the SmartConnector Wizard asks whether the target Manager uses a demo certificate (see [Figure 2-5](#)). Confirm that you want to exit, then click **Done** to dismiss the Install Wizard. This will install the SmartConnector, but leave it un-configured.
- 2 Create a file called **agent.properties** in the directory `$ARCSIGHT_HOME/current/user/agent`, where `$ARCSIGHT_HOME` is the root directory where the SmartConnector component was installed. This file should contain a single line:

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



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

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

Parameter	Description
Ip/Host	IP or host name of the Logger
Port	514 or another port that matches the Receiver
Protocol	UDP or Raw TCP
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 26.

For more information about the Common Event Format (CEF), see [“Common Event Format”](#) on page 259.

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.

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.

Chapter 3

Using the User Interface

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

Navigation: see [“Logger User Interface” on page 31](#)

Performance monitoring: see [“Monitor” on page 33](#)

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

Logger User Interface

The Logger user interface is a web browser application using Secure Sockets Layer (SSL) encryption. Users must be authenticated with a name and password before they can use the interface.

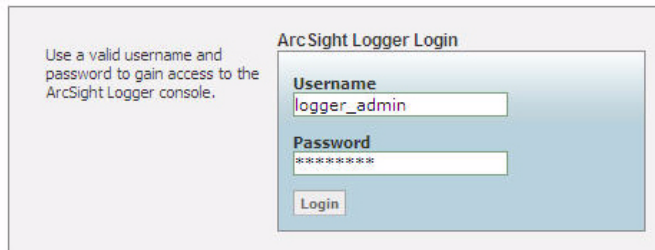


Figure 3-1 Users must login to authenticate themselves to Logger.

Browser Requirements

Logger works with most modern browsers, including Firefox 2.0 and Firefox 3.0.6, or Internet Explorer 6.0 and 7.0. Javascript and cookies must be enabled. An Adobe Flash Player plug-in is required for Internet Explorer 6.0 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/>.

Navigating the User Interface

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

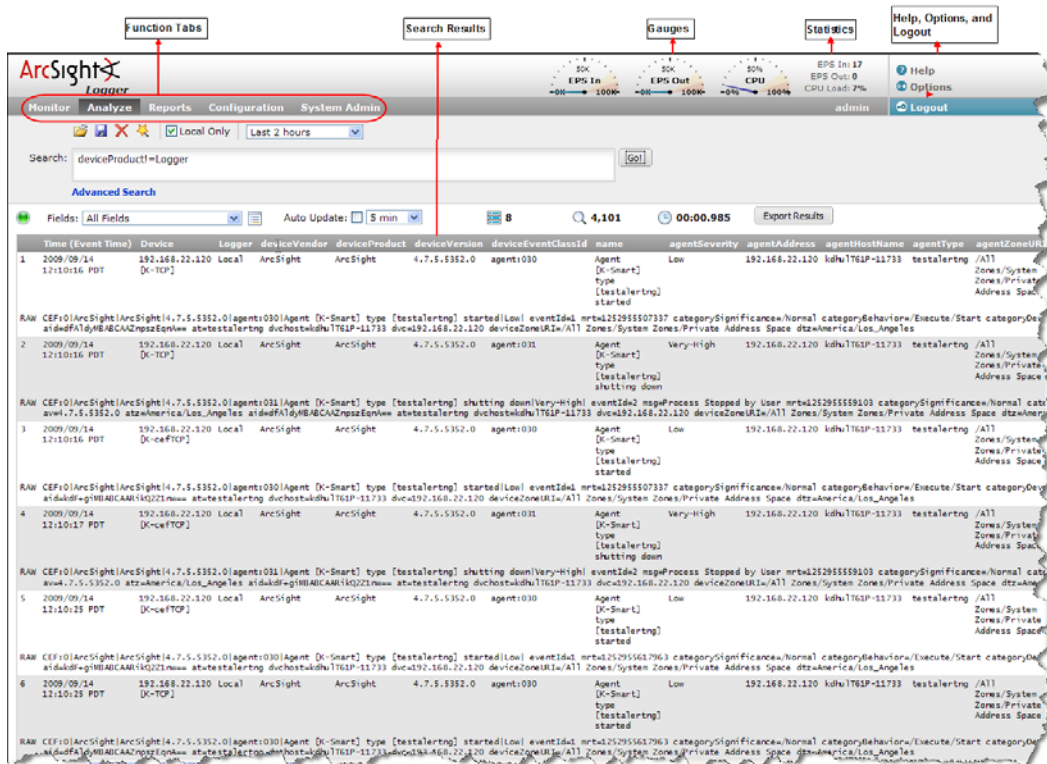


Figure 3-2 Overall layout of the Logger interface.

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



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

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

Help

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

Options

The Options page, shown in [Figure 3-4](#), allows you to set the range on the EPS In and EPS Out gauges. If the event rate exceeds the specified maximum, the range is automatically increased.

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

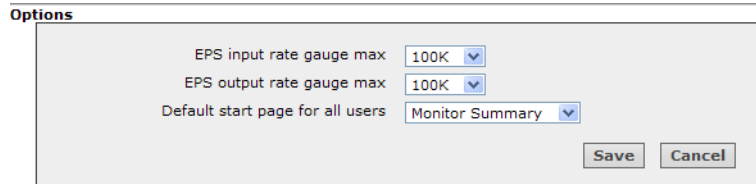
The screenshot shows a web interface titled "Options". It contains three configuration items, each with a label and a dropdown menu: "EPS input rate gauge max" set to "100K", "EPS output rate gauge max" set to "100K", and "Default start page for all users" set to "Monitor Summary". At the bottom right of the form are two buttons: "Save" and "Cancel".

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

Logout

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

Logger automatically logs you out after a user-configurable length of time (15 minutes by default). To change this length of time, see ["Users/Groups" on page 247](#).



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. Under the Monitor tab,

select monitor pages for Summary, Platform, Network, Logger, Receivers, Forwarders, and Storage.

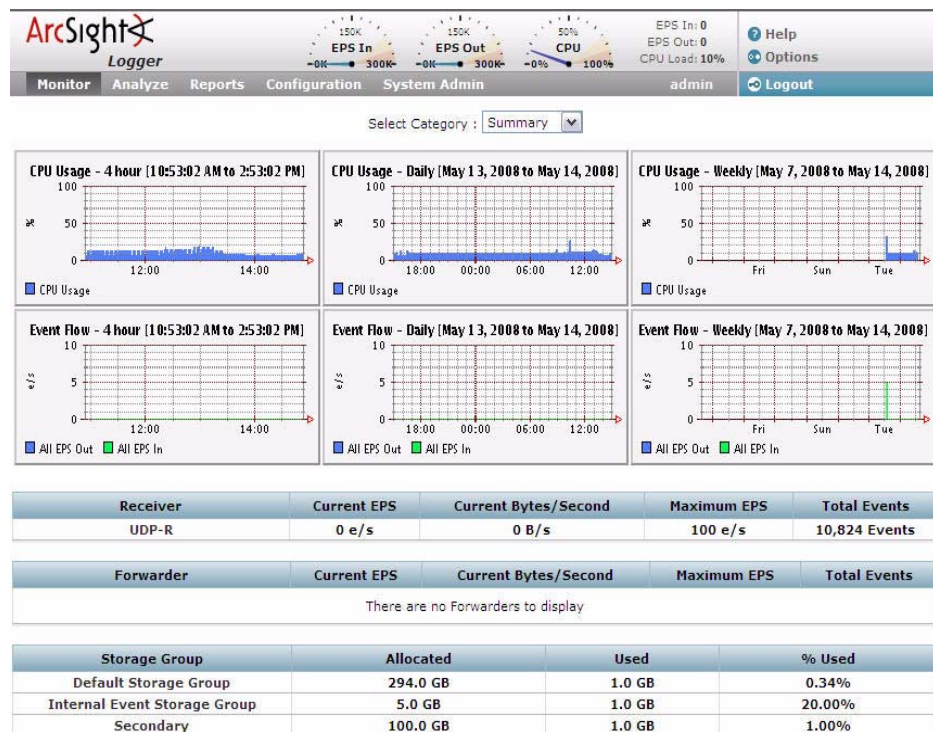


Figure 3-5 The Monitor tab displays summary information by default

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

- 4-hours
- Daily
- Weekly

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

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

Platform

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



Figure 3-6 Platform page of the Monitor tab

Network

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

Logger

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

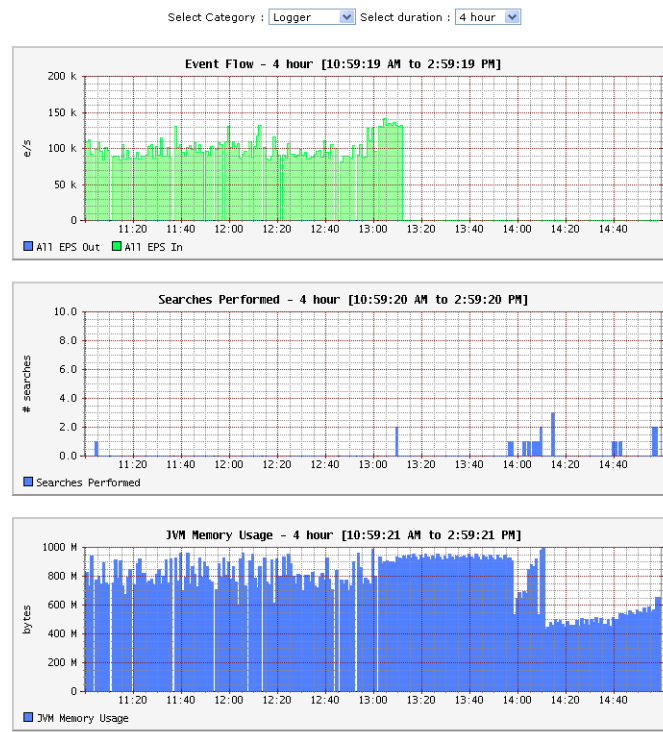


Figure 3-7 Memory usage displayed on the Logger page of the Monitor tab

Receivers

The Receivers monitor page shows total Events per Second (EPS) received and displays values for each configured Receiver.

The list of Receivers includes all Receivers known to the system, including those that are disabled.

To create a new Receiver, or to enable or disable one, see [“Receivers” on page 173](#).

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 179](#).

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 168](#).

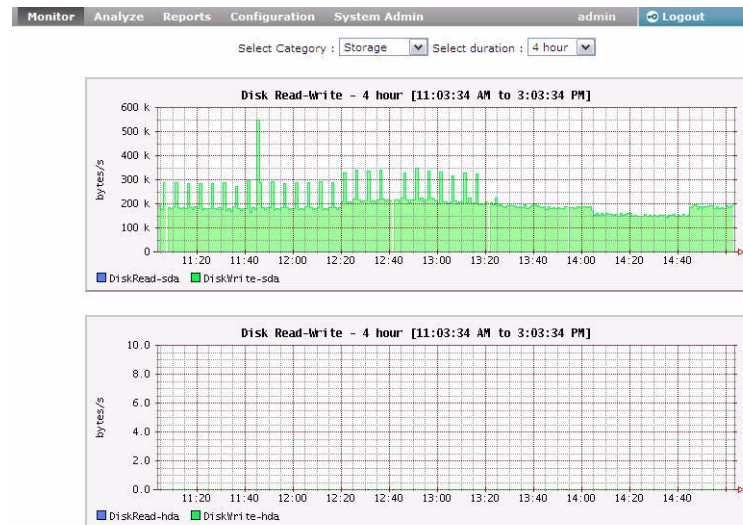


Figure 3-8 Monitor tab, Storage page

Searching and Analyzing Events

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

[“The Need to Search Events” on page 39](#)
[“The Process of Searching Events” on page 40](#)
[“Elements of a Search Query” on page 40](#)
[“Syntax Reference for Query Expression” on page 49](#)
[“Using the Search Builder Tool” on page 53](#)
[“Search Analyzer” on page 57](#)
[“Searching for Events on Logger” on page 59](#)
[“Understanding the Search Results Display” on page 61](#)
[“Exporting Search Results” on page 62](#)
[“Indexing” on page 63](#)
[“Saving Queries \(Saved Filters and Searches\)” on page 67](#)
[“System Filters/Predefined Filters” on page 68](#)
[“Advanced Search Options” on page 71](#)
[“Alerts” on page 71](#)

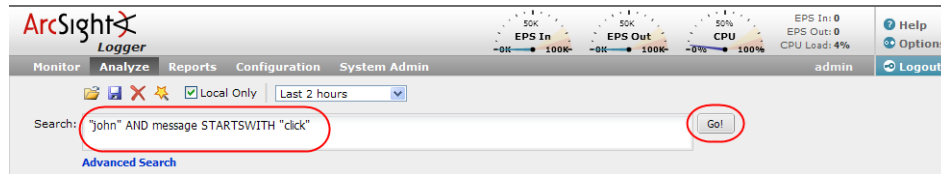
The Need to Search Events

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

Logger provides a user friendly and intuitive search interface to search through all events stored on it.



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

Logger offers several convenient ways to enter a search query—typing the query in the Search text box, using Logger's Search Builder tool to create a query, or using a previously saved query (referred to as filter or saved search). When you type a query, the auto-suggest facility in the user interface provides suggestions and possible matches for the fields you are entering and the applicable operators for those fields, thus enabling you to quickly build a query expression. The auto-suggest facility is available only for fields in the Logger schema, metadata terms (`_storageGroup`, `_deviceGroup`, `_peerLogger`), and the regular expression term (`|REGEX=`). (See ["Indexing" on page 63](#) for a complete list of fields.)

Although a search query on Logger is as simple as entering a keyword to match, you will utilize the full potential of Logger's search operation if you are familiar with all the elements of a query, as described in the next section, ["Elements of a Search Query" on page 40](#).

Elements of a Search Query

A simple Logger search query consists of these elements:

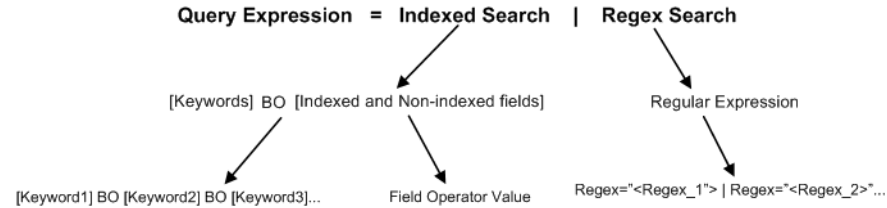
- Query Expression
- Time range
- Field Set

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

Query Expression

A query expression is a set of conditions that are used to select or reject an event when a search is performed. The expression can specify a very simple term to match such as "login" or an IP address; or it can be more complex, such as events that include several IP addresses, ports, and occurred between specific time ranges from devices that belong to a specific device group.

A query expression is what you specify in the Search text box and is specified in the following syntax.



Where

BO is a Boolean Operator—AND, OR, NOT.
If a boolean operator is not specified between keywords or keywords and indexed fields, AND is used.
Operator is one of the operators listed in a table, later in this section

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.

The *Regex search* enables you to define a **regular expression** to refine the data that matched the indexed search filter.

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 49](#).

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 63](#). For discussion on field-based query performance, see ["Performance Optimizations for Indexed Fields in Search Queries" on page 58](#).)

Field search is case sensitive. Make sure you understand and follow other requirements and guidelines listed in ["Syntax Reference for Query Expression" on page 49](#).

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")
```

A complete list of fields you can specify is available in ["Indexing" on page 63](#) 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
>=	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

Field Operator	Example
SQL Operator	
IS	sessionId IS NULL sessionId IS NOT NULL
Boolean Operators	
AND	name="Data List" AND message="Hello" AND 1.2.3.4
OR	(name="TestEvent" OR message="Hello") AND type=2 AND 1.2.4.3
NOT	NOT name="test 123"
List Operator	
IN	priority IN [2,5,4,3] destinationAddress IN ["10.0.20.40", "209.128.98.147"] _deviceGroup IN ["DM1"] _storageGroup NOT IN ["Internal Event Storage Group", "SG1"] _peerLogger IN ["192.0.2.10", "192.0.2.11"]

Regular Expressions

You can also include a regular expression in your search query that defines the pattern you want to match or not match.

A regular expression in a query expression can further refine the search results. For example, to match any IP address, specify

```
|REGEX="\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}"
```

You can specify multiple regular expressions in one query expression. To do so, separate each regular expression with the | character; for example,

```
|REGEX="<REGEX1>" |REGEX="<REGEX2>"
```

Regular expression search is case insensitive. Make sure you understand and follow the requirements and guidelines listed in [“Syntax Reference for Query Expression” on page 49](#).

The following is a list of the characters used in a regular expression syntax and their meaning. For a tutorial introduction to regular expressions, see Appendix B, [“Regular Expressions” on page 265](#).

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]

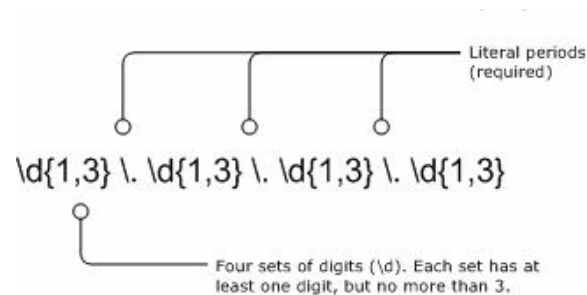
\W A non-word character: [^\w]

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

The following figure illustrates a regular expression that matches any IP address in the dotted-decimal format.



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:

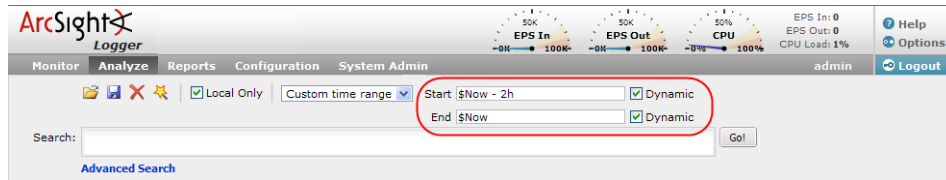
Start: 8/13/2008 13:36:30

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 45](#). The <units> portion, if given, consists of an integer and a single, case-sensitive letter, as shown in [Table 4-2 on page 45](#).

Table 4-1 Current Period


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

Table 4-2 Units

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

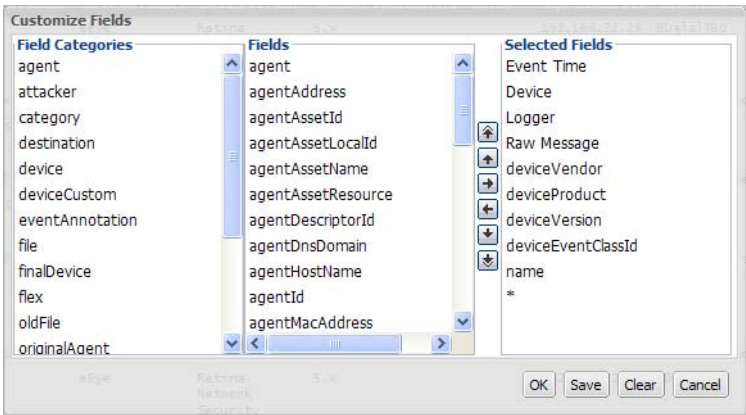
Field Set

A field set determines the fields that are displayed in the search results for each event that matched a search query. Logger provides a number of predefined field sets, as listed in the

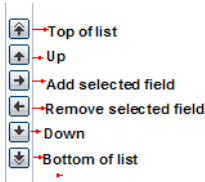
following table. To view the fields includes in each of the predefined field sets, click the  (Customize Fieldset) icon. When you run the first search operation in a new browser window, you might not be able to select the field sets as they are hidden. The field sets list is displayed after you have run the first search operation.

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

Starting with Logger v4.0, you can 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.

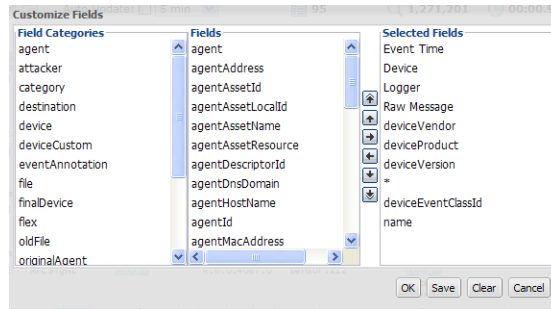


Use these buttons to create and edit a custom field set.



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

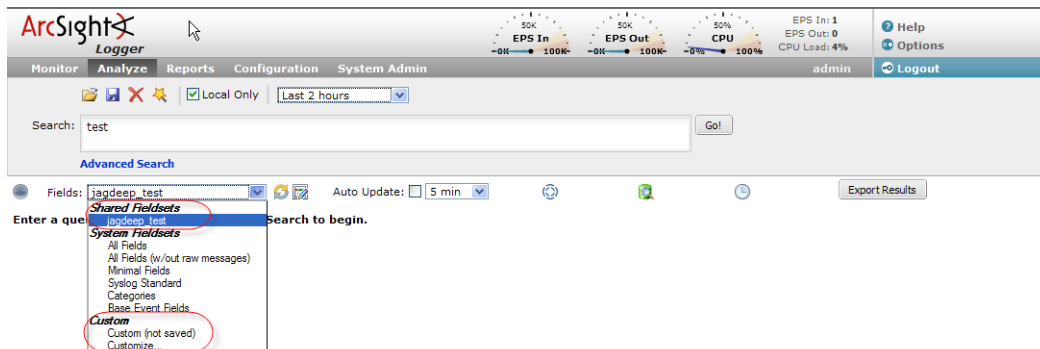
results will list the fields before the asterisk (“*”) first, followed by any other fields in an event. Lastly, the deviceEventClassId and Name fields will be listed.



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 202](#).

Constraints

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

- Devices in a particular device group
- Stored in particular storage groups

- On specific peer Loggers

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

Using constraints can speed up a search operation as they limit the scope of data that needs to be searched. Follow these guidelines when specifying constraints:

- A device group constraint can contain devices or device groups.
- When specifying multiple groups in a constraint, ensure that the group names are enclosed in a square bracket; for example, `_storageGroups IN ["SGA", "SGB"]`.
- Use the following operators to specify constraints in a search query expression:



Use the identifier `_deviceGroup` to also specify individual devices, as shown in the example in the following table.

Metadata Identifier	Example
<code>_deviceGroup</code>	<code>_deviceGroup IN ["DM1", "HostA"]</code> where DM1 is a device group, while HostA is a device.
<code>_storageGroup</code>	<code>_storageGroup IN ["Internal Event Storage Group", "SG1"]</code>
<code>_peerLogger</code>	<code>_peerLogger IN ["192.0.2.10", "192.0.2.11"]</code>

You can apply constraints to a search query in these ways:

- ◆ Typing the constraint in the Search text box

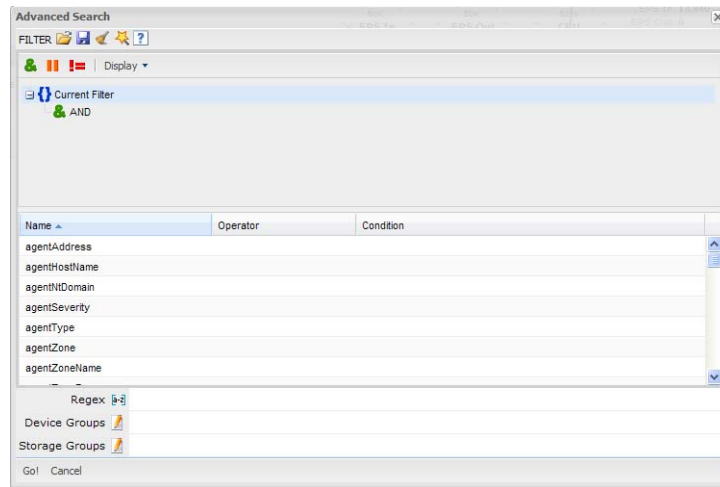
Use Logger's auto-suggest facility to enter a constraint in the Search text box. Once you type `"_s"` (for storage group), `"_d"` (for device group), or `"_p"` (for peerLogger) in the Search text box, Logger automatically provides a drop-down list of relevant terms and operators from which you can select.



If a search query contains constraints and a regular expression, make sure that the constraints are specified before the regular expression. For example, `_peerLogger IN ["10.10.10.10"] name contains abc | REGEX=":\d31"`

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

information about the Search Builder, see [“Using the Search Builder Tool” on page 53](#).



Syntax Reference for Query Expression

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

Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
Syntax	keyword1 boolean_operator keyword2 boolean_operator keyword3...	field_name operator field_value (List of fields in the “Indexing” on page 63 section.) (List of operators in the “Indexed and Non-Indexed Fields” on page 41 section.)	REGEX="<REGEX1>" REGEX="<REGEX2>" ..

Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
Operators	<p>Upper-, lower-, or mixed case boolean operators—AND, OR, NOT. If an operator is not specified, AND is used.</p> <p>To search for literal operator AND, OR, NOT, in an event, enclose them in double quotes.</p> <p>Example: "AND", "or", "Not"</p>	<p>Use any operator listed in the "Indexed and Non-Indexed Fields" on page 41 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. <p>Examples:</p> <pre>message STARTSWITH="NOT" message="LOGIN DID NOT SUCCEED"</pre>	<p> and the operators described in "Regular Expressions" on page 43.</p> <p>Use this operator to AND multiple regular expressions in one query expression.</p>
Nesting (including parenthetical clauses, such as (a OR b) AND c)	<p>Allowed</p> <ul style="list-style-type: none"> Use boolean operators to connect and nest keywords. Metadata identifiers (_storageGroup, _deviceGroup, and _peerLogger), but can only appear at the top level in a query expression). If the query contains a regular expression, the metadata identifiers need to precede the regular expression. 	<p>Allowed</p> <ul style="list-style-type: none"> Use any operator listed in the "Indexed and Non-Indexed Fields" on page 41 section to connect and nest field search expressions. Metadata identifiers (_storageGroup, _deviceGroup, and _peerLogger), but can only appear at the top level in a query expression 	<p>Multiple regular expression can be specified in one query using this syntax:</p> <pre> REGEX="<REGEX1 >" REGEX="<REGEX2 >" ...</pre>
Case sensitivity	<p>Insensitive</p> <p>(Cannot be changed.)</p>	<p>Sensitive*</p> <p>(Can be changed using Tuning options. See "Tuning Advanced Search Options" on page 201.)</p>	<p>Insensitive*</p> <p>(Can be changed using Tuning options. See "Tuning Advanced Search Options" on page 201.)</p>

Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
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>
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>

Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
Space Tab Newline , ; () [] { } " * > < !	Cannot search for these characters. Examples: "John Doe" is invalid	No restrictions. Enclose special character in double quotes. Escape the wildcard character and double quotes. Example: name="John* \"Doe" (matches John* "Doe")	No restrictions. Special regular expression characters such as (,), [,], {, }, ", , and * need to be escaped.
= . : / \ @ - ? # \$ & - %	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.

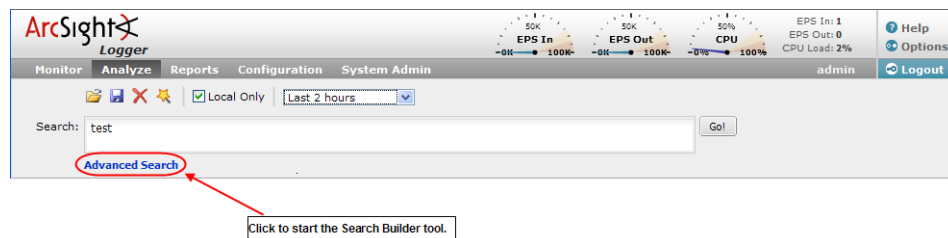
Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
Time format, when searching for a specific timestamp	<p>No specific format. The query needs to contain the exact timestamp string. For example, "10:34:35".</p> <p>Note: The string cannot contain spaces. For example, "Oct 19" is invalid.</p>	<p>Use this format to specify a timestamp in a query (including double quotes):</p> <p><code>"mm/dd/yyyy hh:mm:ss"</code></p> <p>OR</p> <p><code>"yyyy/mm/dd hh:mm:ss timezone"</code></p> <p>OR</p> <p><code>"MMM dd yyyy hh:mm:ss timezone"</code></p> <p>where mm—month dd—day yyyy—year hh—hour mm—minutes ss—seconds timezone—EDT, CDT, MDT, PDT. MMM—First three letters of a month's name; for example, Jan, Feb, Mar, Sep, Oct, and so on.</p>	No restrictions.

Using the Search Builder Tool

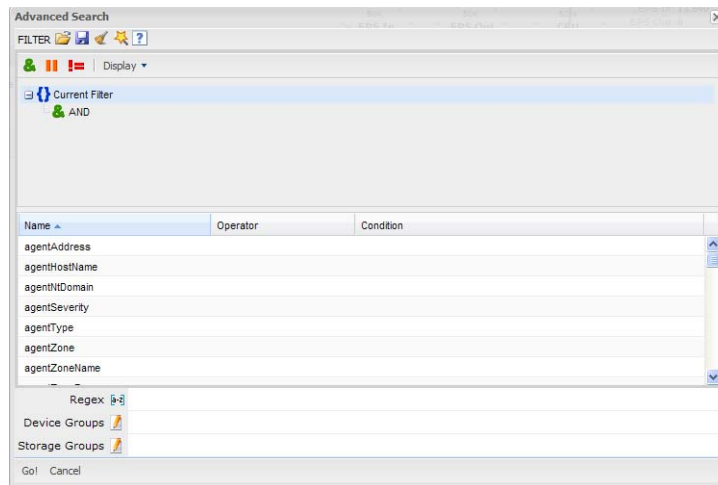
The Logger Search Builder tool is a boolean-logic conditions editor that enables you to quickly and accurately build search queries. The tool provides a visual representation of the conditions you are including in a query. You can specify keywords, field-based conditions, and regular expressions using this tool. In addition, the tool enables you to specify search constraints such as device groups and storage groups (see ["Constraints" on page 47](#)). 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

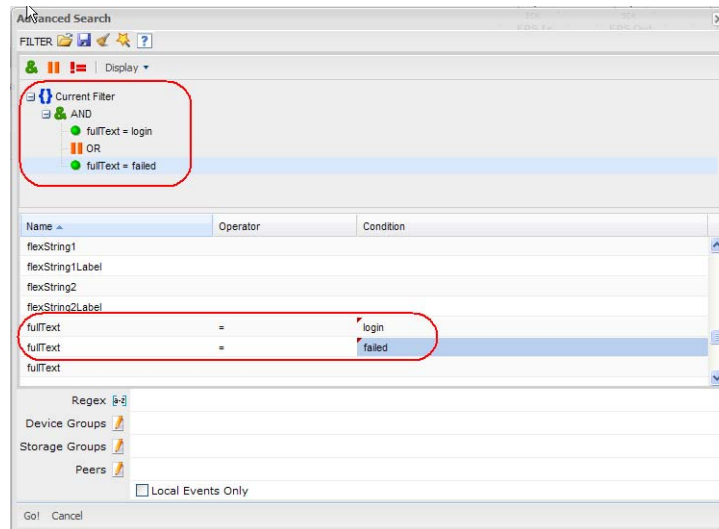
- 2 If you want to load a system or saved filter, or a saved search, click the icon. Select the filter or the saved search from the displayed list and click **Load+Close**.

For more information, see [“Saving Queries \(Saved Filters and Searches\)”](#) on page 67 and [“System Filters/Predefined Filters”](#) on page 68.

- 3 To add a keyword (full-text search) or field condition:

- a Locate the field you want to add under the Name column.

To specify a keyword (full-text search), use the *fullText* field under the Name column, as shown in the following figure.




- b** Click the Operator column associated with the field, select the operator from the displayed list, and press **Enter**.

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

- c** In the Condition column associated with the field, enter a value and press **Enter**.



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

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

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


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

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

- 7** Click **Go**.

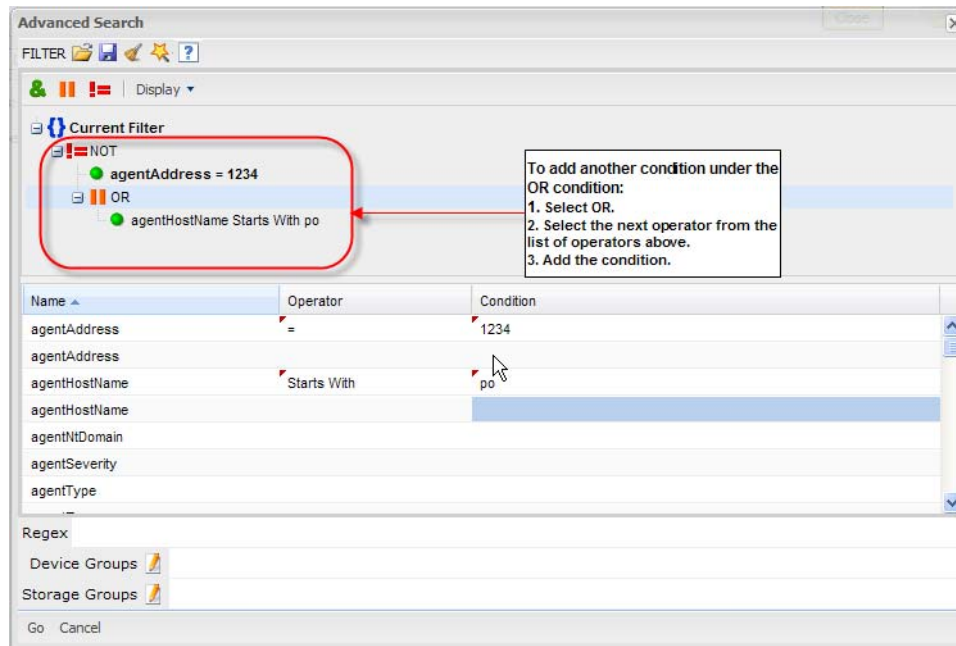
The query is automatically displayed in the Search text box and is ready to be run.

OR

Click the  icon to save the query (referred as Saved Filter or a Saved Search) for a later use. For more information about saving queries, see [“Saving Queries \(Saved Filters and Searches\)” on page 67](#).

Nested Conditions

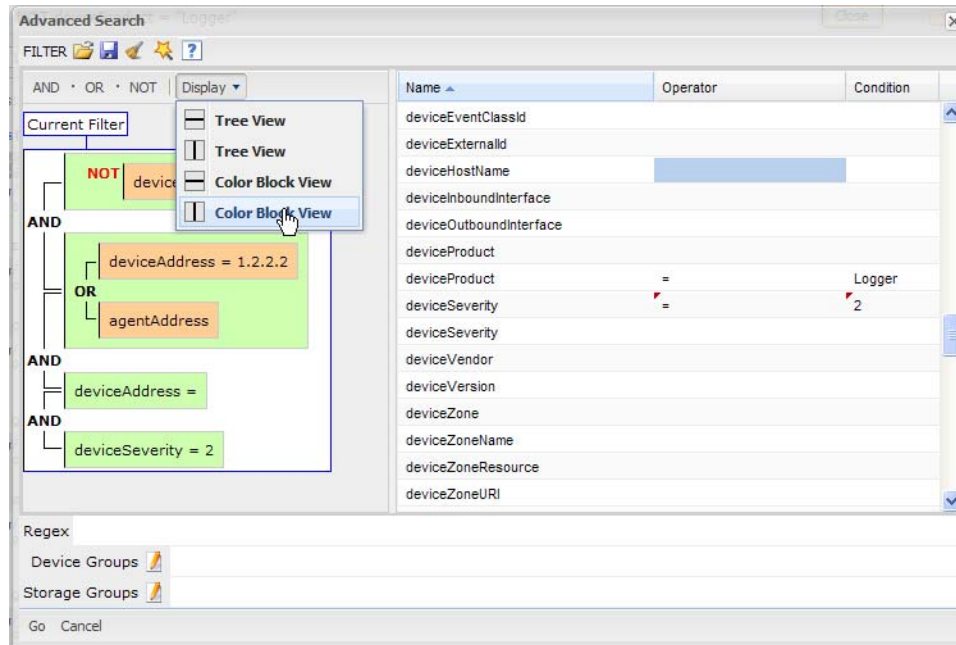
You can create search queries with nested conditions in Search Builder. To do so, click the operator under which you want to nest the next condition and add the condition as described in [“Accessing Search Builder” on page 53](#).



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.

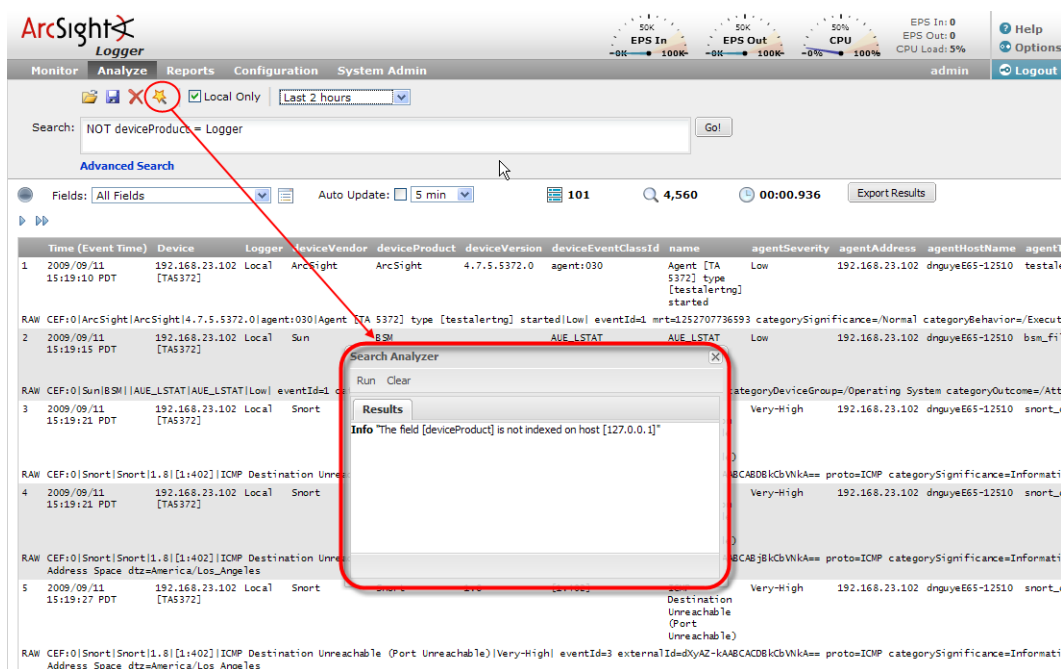
Search Analyzer

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:

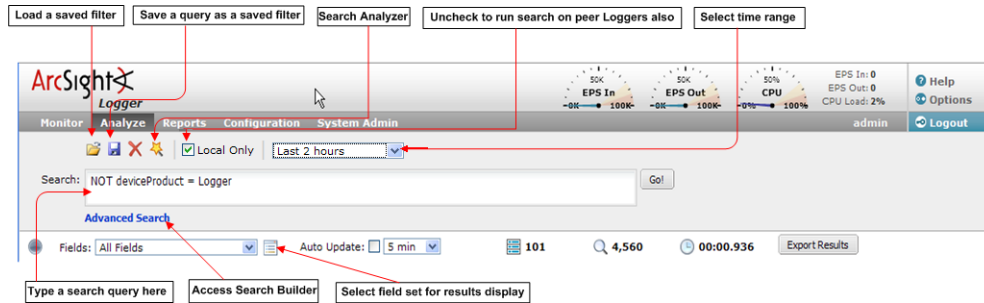
- When you include indexed and non-indexed fields in a query. Therefore, ArcSight recommends that you identify the fields that you will most commonly use in queries and index all those fields.
- When you perform search on data in a time range in which a currently indexed field (included in the query) 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.

Searching for Events on Logger

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



To search for events on Logger:

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


Note: Refer to [“Query Expression” on page 40](#) 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 40](#).

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

- Fields in Logger schema
(See [“Indexing” on page 63](#) for a complete list of fields.)
- Metadata terms (`_storageGroup`, `_deviceGroup`, `_peerLogger`)
Type “_s” (for storage group), “_d” (for device group), or “_p” (for peerLogger) in the Search text box to obtain a drop-down list of constraint terms and operators.
- Regular expression term (`|REGEX=`)


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

- b Click **Advanced** to use the Search Builder tool. (See [“Using the Search Builder Tool” on page 53](#) for more information.) Also, use this option to specify device groups, storage groups, and peer Loggers to which search should be limited.
- c Click the  icon to load a saved filter, a system filter, or a saved search. Select the filter or the saved search from the displayed list and click **Load+Close**.

For more information, see [“Saving Queries \(Saved Filters and Searches\)” on page 67](#) and [“System Filters/Predefined Filters” on page 68](#).

- 3 Use the following default values or change them suit your needs:
 - a **Local Logger:** By default the query is run on the local Logger only. If you want to run the query on the peer Loggers as well, uncheck the “Local Only” field located to the right of the Go! button.
 - b **Time Range:** By default, the query is run on the data received in the last two hours on the Logger. Click the drop-down list to select another predefined time range or specify a custom time range. For more information about time ranges, see [“Time Range” on page 44](#).
 - 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 45](#).
- 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 61](#).

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 67](#).

Searching Peer Loggers (Distributed Search)

When you run a search query, by default, only your local Logger is searched for matching events. However, when specifying a query, you can select an option to run the search on the peer Loggers. You can also select the peer Loggers to which the search should be constrained, as described in [“Searching for Events on Logger” on page 59](#).

If peer Loggers do not have identical storage or device group names, a search query operation skips searching for events for those groups on those peers.

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

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

```
[Peer Logger IP address] Error: Get Query Statistics
```

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

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

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

A query containing keywords can be run only on a Logger running version 4.0. Similarly, a query containing indexed fields can only be run on Loggers running version 3.0 or later. Therefore, before running a query, make sure all Loggers on which it will run are running the supported software version. Running a query on an unsupported version will generate an error message.

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.

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.

To see the next screen of events, click ; or to leap forward by 10 pages of 25 events each (that is, leap forward by 250 events). Once you are past the first screen of events, you can click to go back to the previous screen; or to leap backwards by 10 pages of 25 events each.

The Search Results page also displays the number of events scanned and number of events matching the query and the time it took to run the search.

The screenshot shows the ArcSight Logger Search Results page. At the top, there are status indicators for EPS In, EPS Out, and CPU. Below these are tabs for Monitor, Analyze, Reports, Configuration, and System Admin. The Analyze tab is active, showing a search bar and a 'Go!' button. Below the search bar, there are fields for 'Fields' (set to 'All Fields') and 'Auto Update' (set to '5 min'). To the right of these fields are three statistics: 1,725 (number of matching events), 1,813 (number of events scanned), and 00:05.420 (time taken to run the search). Below these statistics are three buttons: 'Fast Backward', 'Previous', 'Next', and 'Fast Forward'. A table of search results is displayed below the buttons. The table has columns for Time, Event Time, Device, Logger, deviceVendor, deviceProduct, deviceVersion, deviceEventClassId, name, agentSeverity, baseEventCount, and deviceCustomNumber1. The first row shows an event from 19/Aug/2009 at 14:21:00, logged by Local ArcSight, with a name of 'Login Internal Event'. The second row shows a raw event from 19/Aug/2009 at 14:21:00, logged by Local ArcSight, with a name of 'Login Internal Event'. The third row shows a raw event from 19/Aug/2009 at 14:21:00, logged by Local ArcSight, with a name of 'Login Internal Event'. A red line points from the 'Fast Backward' button to the first event in the table. A legend at the bottom right explains the statistics: 1,725 is the number of matching events, 1,813 is the number of events scanned, and 00:05.420 is the time taken to run the search.

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.

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.

The **Shift** and **Ctrl** keys modify this behavior:

- Click—To add the term to query
- **Shift**+Click—To add the negated term to query; for the above example, the query will change to "login and NOT fail"
- **Shift**+**Ctrl**+Click—To replace the current query with the new term

Search results are sorted by receipt time.

Auto Updating Search Results

The Auto Update feature executes the search over specified intervals, updating the search results if new events match the query. Depending on your needs, you can auto update the search results every:

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

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

To auto update search results:

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

Exporting Search Results

You can export the search results to a comma-separated values (CSV) file for further analysis with other software applications. For CEF events, each named value is saved in a separate column.

Typically, you would use a spreadsheet application such as Microsoft Excel to view the exported CSV file. If the number of exported events is larger than 65,535, which is the maximum number of rows MS Excel can process, you might need to use a database management system to process the exported data.

To export events:

- 1 Follow instructions up to [Step 7 on page 55](#).
- 2 Click **Export Results** to display various export options, as shown in the following figure.

Export Options

☒ Save to local disk
 ☐ Export to remote location
 ☐ Save to Logger

☐ Include Summary
☐ Exclude Non-CEF Events
☐ Include Base Events

Fields:

Event Time, Device, Logger, Signature ID, Name, Device Vendor, Device Product, categoryBehavior, categoryDeviceGroup, categoryObject, categoryOutcome

☒ All Fields
 [Clear](#)

[Export](#)
[Cancel](#)

- 3 Select from one of the export options, as displayed in [3 on page 63](#).

Option	Description
Save to local disk	The file is saved to a local system or is it sent to the browser for viewing or saving.
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 written to the Logger's local storage.
Export File Name	Name of the file.
Fields	A list of event fields that will be included in the exported file. By default, all fields are included. You can enter fields or edit the displayed fields by deselecting All Fields.
Include Summary	Include an event count in the exported search results.
Exclude Non CEF Events	Only include CEF events in the exported search results.
Include Base Events	Include base events in the exported search results.

- 4 Click **Export**.

Scheduling an Export Operation

The time it takes to export search results is proportional to the number of events being exported. Therefore, for a large number of events, ArcSight recommends that you schedule the export operation to be performed at a later time by saving the query and time parameters as a Saved Search, and then scheduling a Saved Search Job. For more information about Saved Search Jobs, see ["Saved Search Jobs" on page 198](#).

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 64](#).
- Field-based indexing—Event fields are indexed based on a predetermined schema. See ["Field-based Indexing" on page 64](#).

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

How indexing works

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

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

The following events are not indexed:

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

Full-text Indexing (Keyword Indexing)

For full-text indexing, each event is scanned and divided into keywords and stored on the Logger. Full-text indexing is not enabled by default; you are prompted to enable it at initialization time (described at [“Initialization Sequence” on page 20](#)). 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 66](#).

Field-based Indexing

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

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

Although you can add fields to an index at any time, you are presented with a list of recommended fields during the initialization sequence of Logger (described at [“Initialization Sequence” on page 20](#)). Once you index those fields, indexing is enabled at initialization time. You can add more fields to an index at any time. Once a field has been added, you cannot remove it.



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

Once you enable indexing on a Logger, Logger starts indexing the event metadata fields—event time, Logger receipt time, and device address—for every event in addition to the fields you added to the index. The event metadata fields are also referred to as the “internal” fields and are in addition to the fields you can add through the Logger’s user interface.

The following fields are available for indexing. The fields that ArcSight recommends to you to add during Logger initialization are indicated in **bold** font. In addition to the following fields, the [requestUrl](#) field is available for search queries. However, this field **cannot** be indexed.

Indexable Fields		
agentAddress	deviceCustomDate1Label	flexDate1
agentHostName	deviceCustomDate2	flexDate1Label
agentNtDomain	deviceCustomDate2Label	filePath
agentSeverity	deviceCustomNumber1	flexNumber1
agentType	deviceCustomNumber1Label	flexNumber1Label
agentZone	deviceCustomNumber2	flexNumber2
agentZoneName	deviceCustomNumber2Label	flexNumber2Label
agentZoneResource	deviceCustomNumber3	flexString1
agentZoneURI	deviceCustomNumber3Label	flexString1Label
applicationProtocol	deviceCustomString1	flexString2
baseEventCount	deviceCustomString1Label	flexString2Label
bytesIn	deviceCustomString2	message
bytesOut	deviceCustomString2Label	name
categoryBehavior	deviceCustomString3	priority
categoryDeviceGroup	deviceCustomString3Label	requestClientApplication
categoryObject	deviceCustomString4	requestContext
categoryOutcome	deviceCustomString4Label	requestMethod
categorySignificance	deviceCustomString5	requestUrlFilename
categoryTechnique	deviceCustomString5Label	requestUrlQuery
customerName	deviceCustomString6	sessionId
destinationAddress	deviceCustomString6Label	sourceAddress
destinationDnsDomain	deviceEventCategory	sourceHostName
destinationHostName	deviceEventClassId	sourceMacAddress
destinationMacAddress	deviceExternalId	sourceNtDomain
destinationNtDomain	deviceHostName	sourcePort
destinationPort	deviceInboundInterface	sourceProcessName
destinationProcessName	deviceOutboundInterface	sourceServiceName
destinationServiceName	deviceProduct	sourceTranslatedAddress
destinationTranslatedAddress	deviceSeverity	sourceUserId
destinationUserPrivileges	deviceVendor	sourceUserName
destinationUserId	deviceVersion	sourceZone

Indexable Fields

destinationUserName	deviceZone	sourceZoneName
destinationZone	deviceZoneName	sourcezoneResource
destinationZoneName	deviceZoneResource	sourceZoneURI
destinationZoneResource	deviceZoneURI	startTime
destinationZoneURI	endTime	transportProtocol
deviceAction	eventId	type
deviceAddress	externalId	vulnerabilityExternalId
deviceCustomDate1	fileName	VulnerabilityURI

Guidelines for Field-based Indexing

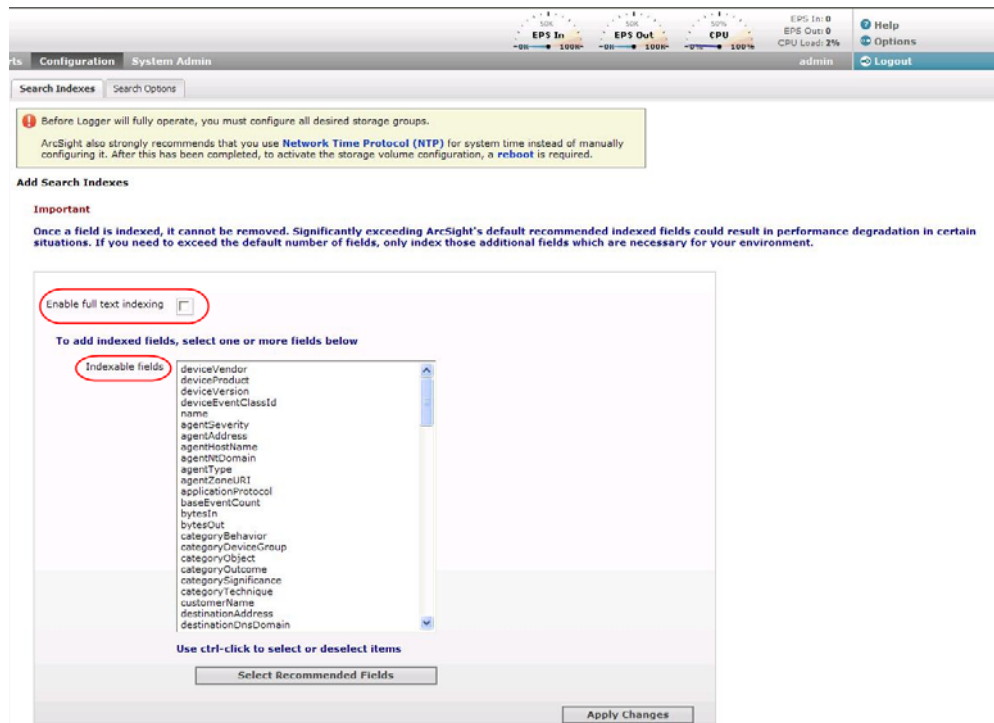
Make sure you are familiar with these guidelines before you index any fields:

- 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 peer Loggers. However, it will run at optimal speed on the local Logger. Therefore, the search performance in such a setup will be slow.
- Although the `requestUrl` field is available for search and report queries, it cannot be indexed. Including this field in such queries will result in the query running slower than a search performed on indexed data.

Enabling Indexing

If you did not enable indexing on Logger at initialization time (described at [“Initialization Sequence” on page 20](#)), 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:

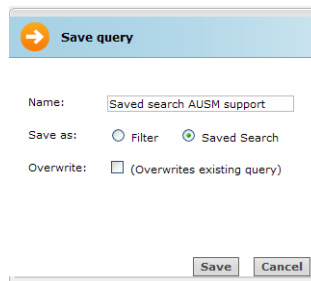
- As a filter
A Filter saves the query expression, but does not save the time range or the field set information.
- As a saved search
A saved search saves the query expression and the time range that you specified.

Saving a Query

To save a query:

- 1 Define a query as described in [“Searching for Events on Logger” on page 59](#) or [“Using the Search Builder Tool” on page 53](#).

- Click the Save icon (💾) and enter a name for the query in the Name field, as shown in the following figure.



A dialog box titled "Save query" with a blue header bar containing a right arrow icon. It contains the following fields and options:

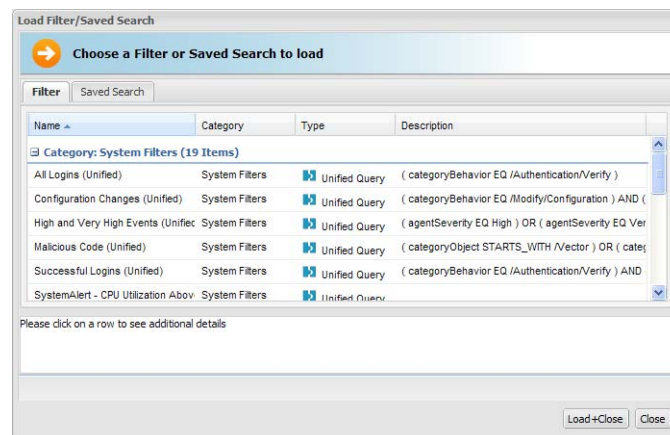
- Name:** A text input field containing "Saved search AUSM support".
- Save as:** Two radio buttons: "Filter" (unselected) and "Saved Search" (selected).
- Overwrite:** A checkbox labeled "(Overwrites existing query)" which is currently unchecked.
- At the bottom right are "Save" and "Cancel" buttons.

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

Using a Saved Filter or a Saved Search

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

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



The "Load Filter/Saved Search" interface shows a tabbed view with "Filter" and "Saved Search" tabs. The "Filter" tab is active, displaying a table of system filters. Below the table is a text area for details and "Load+Close" and "Close" buttons at the bottom right.

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

Please click on a row to see additional details

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.

- 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 appliance 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.

The following is a list of all the system filters.

Field Query (Indexed Search) Filters	Regex Query Filters
All Logins	All Logins (Non-CEF) All Logins (CEF format)
Configuration Changes	System configuration changes (CEF format)
High and Very High Events	High and Very High CEF events
Unsuccessful Logins	Unsuccessful Logins (Non-CEF) Unsuccessful Logins (CEF format)
Successful Logins	Successful Logins (Non-CEF) Successful Logins (CEF format)
Malicious Code	Malicious Code (CEF format) Unusual events (Non-CEF)
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.	
CPU Utilization Above 90 Percent	In addition to these filters, you can define your own alerts based on the system health events listed in “System Health Events” on page 70 .
CPU Utilization Above 95 Percent	
Disk Space Below 10 Percent	
Disk Space Below 5 Percent	
Filter Configuration Changes	
High CPU Temperature	
Power Supply Failure	
RAID Status Disk Failure	
Storage Configuration Changes	
Zero Events Incoming	
Zero Events Outgoing	
Zero Fan Speed	

Using a System Filter

To use a predefined system filter, follow instructions in [“Using a Saved Filter or a Saved Search” on page 68](#).

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 68](#). The predefined system health filters are based on the system health events listed in [“System Health Events” on page 70](#).

- Searching for system health events in Logger's Internal Storage Group, as listed in ["System Health Events" on page 70](#). If a predefined system health filter does not suit your needs, you can create alerts based on the system health events.

To set up notification of system health events

- 1 Configure the Logger's SMTP with the desired e-mail address destination (see ["SMTP Settings" on page 226](#)) or create an SNMP Destination (see ["SNMP Destinations" on page 191](#)) or Syslog Destination (see ["Syslog Destinations" on page 191](#)).
- 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 187](#)).
- 3 Enable the new Alert.

System Health Events

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

Group	Device Event Category	Device Event Class ID
CPU	/Monitor/CPU/Usage	cpu: 100
Disk	/Monitor/Disk/Space/Remaining/Data	disk: 100
	/Monitor/Disk/Space/Remaining/Root	disk: 101
	/Monitor/Disk/Read	disk: 102
	/Monitor/Disk/Write	disk: 103
EPS	/Monitor/Receiver/All/EPS	eps: 100
	/Monitor/Forwarder/All/EPS	eps: 100
Memory	/Monitor/Memory/Usage/Platform	memory: 100
	/Monitor/Memory/Usage/Jvm	memory: 101
Network	/Monitor/Network/Usage/In	network: 100
	/Monitor/Network/Usage/Out	network: 101
Raidcontroller	/Monitor/RAIDController/Configuration/R AID-5	raidcontroller: 100
	/Monitor/RAIDController/Port/p0	raidcontroller: 101
	/Monitor/RAIDController/Port/p1	raidcontroller: 102
	/Monitor/RAIDController/Port/p2	raidcontroller: 103
	/Monitor/RAIDController/Port/p3	raidcontroller: 104
	/Monitor/RAIDController/Sensor/bbu	raidcontroller: 105
Search	/Monitor/Search	search: 100

Group	Device Event Category	Device Event Class ID
Sensor	/Monitor/Sensor/CPU1	sensor: 100
	/Monitor/Sensor/CPU2	sensor: 101
	/Monitor/Sensor/System	sensor: 102
	/Monitor/Sensor/DIMM	sensor: 103
	/Monitor/Sensor/CPU1Core	sensor: 104
	/Monitor/Sensor/CPU2Core	sensor: 105
	/Monitor/Sensor/3.3V	sensor: 106
	/Monitor/Sensor/5V	sensor: 107
	/Monitor/Sensor/12V	sensor: 108
	/Monitor/Sensor/-12V	sensor: 109
	/Monitor/Sensor/Battery	sensor: 110
	/Monitor/Sensor/FAN1	sensor: 111
	/Monitor/Sensor/FAN2	sensor: 112
	/Monitor/Sensor/FAN3	sensor: 113
	/Monitor/Sensor/FAN4	sensor: 114
	/Monitor/Sensor/FAN5	sensor: 115
	/Monitor/Sensor/FAN6	sensor: 116
	/Monitor/Sensor/FAN7	sensor: 117
	/Monitor/Sensor/FAN8	sensor: 118
	/Monitor/Sensor/PowerSupply	sensor: 119

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 201](#).

Alerts

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

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

Viewing Alerts

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

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

When you create Alerts (see [“Alerts” on page 187](#)), 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 44](#) for more detail.

Receiving Alerts for Events

To receive alerts:

- 1 Configure the Logger’s SMTP with the desired e-mail address destination (see [“SMTP Settings” on page 226](#)) or create an SNMP Destination (see [“SNMP Destinations” on page 191](#)) or Syslog Destination (see [“Syslog Destinations” on page 191](#)).



Number of destinations per alert:

- E-mail: Multiple, each separated by a semicolon.
 - 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 67](#).)

- 3 Create an Alert that uses the new Filter and specify match count and threshold (see [“Alerts” on page 187.](#)) Enable the new Alert.

The screenshot shows the ArcSight Monitor interface. At the top, there are tabs for Monitor, Analyze, Reports, Configuration, and System Admin. The user is logged in as 'admin' and can click 'Logout'. Below the tabs, there are filters for 'Show: All Alerts' and 'Within: Last 2 hours'. There are also buttons for 'Go!', 'Export Results', and an 'Auto Update' checkbox set to '5 min' and 'Paused'.

The main section displays 'Alerts: 25' and 'Status: Paused'. It shows the 'Page Start: 14/May/2008 13:11:38 -0700' and 'Page End: 14/May/2008 13:11:39 -0700'. A 'Next >' link is available.

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

Base Event (1 found)

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

The screenshot shows a repeating pattern of these tables, indicating multiple alerts and base events.

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 59](#), [“Understanding the Search Results Display” on page 61](#), [“Viewing Alerts” on page 71](#), and [“Advanced Search Options” on page 71](#).

Chapter 5

Reporting

This chapter describes Logger reporting features.

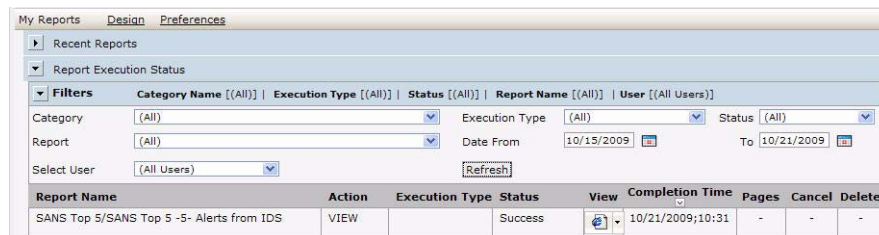
Reports are captured views or summaries of events which you can view from the Logger Reports tab or export for sharing in a variety of file formats. Reporting is an essential tool for communicating the state of your network security to internal and external stakeholders.

["Navigating to Reports" on page 75](#)
["Report Groups" on page 76](#)
["Reports Home Page" on page 81](#)
["Using the Dashboard" on page 82](#)
["Running, Viewing, and Publishing Reports" on page 95](#)
["Designing Reports" on page 106](#)
["Scheduling Reports" on page 155](#)
["Deploying a Report Package" on page 159](#)
["Report Server Administration" on page 160](#)
["Backup and Restore of Report Content" on page 162](#)

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' web interface. At the top, there are tabs for 'My Reports', 'Design', and 'Preferences'. Below the tabs, there is a section for 'Recent Reports' and 'Report Execution Status'. A 'Filters' section allows users to filter reports by Category, Report, Select User, Execution Type, Status, Date From, and Date To. Below the filters is a table of reports. The table has columns for Report Name, Action, Execution Type, Status, View, Completion Time, Pages, Cancel, and Delete. The first row shows a report named '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.

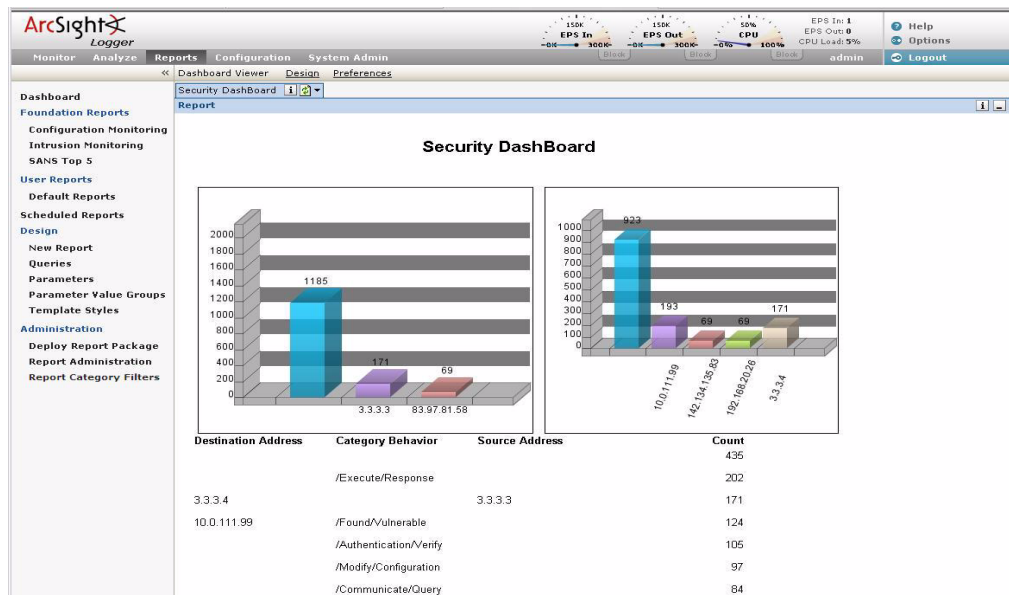


Figure 5-2 Reports Home Page with a Dashboard

For information about designing and selecting dashboard views, see [“Using the Dashboard” on page 82](#).

Report Groups

Logger supports the following report groups:

- **“Foundation Reports” on page 77**—This report group contains ready-made reports that address common security use cases. This report group is displayed by default.
- **“Solution Reports” on page 79**—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 79**—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 80**—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 159](#).)

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 159.](#))

The Foundation Reports include [“SANS Top 5 Reports” on page 77](#), [“Network Monitoring Reports” on page 78](#), [“Intrusion Monitoring Reports” on page 78](#), [“Configuration Monitoring Reports” on page 79](#).

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 155](#).

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 95](#) and [“Task Options on Available Reports” on page 96.](#))

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 95](#).

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 95](#).

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 95](#).

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 159](#).)

For information on deploying Solutions Packages, see [“Deploying a Report Package” on page 159](#). 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 155](#).

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 95](#) and [“Task Options on Available Reports” on page 96](#).)

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 80](#), [“Cross Device Reports” on page 80](#), [“Database Reports” on page 80](#), [“Firewall Reports” on page 80](#), [“Identity Management Reports” on page 80](#), [“IDS-IPS Reports” on page 80](#), [“Network Reports” on page 80](#), [“Operating System Reports” on page 80](#), [“VPN Reports” on page 80](#).

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 155](#).

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 95](#) and [“Task Options on Available Reports” on page 96](#).)

Anti-Virus Reports

These reports provide information on anti-virus activity, such as the anti-virus update status, virus activity by hour, and top infected systems.

For a complete list of reports, click Reports | **Anti-Virus** under the Device Monitoring Reports section on the left panel menu.

Cross Device Reports

These reports provide information on functions that apply to multiple kinds of devices, such as failed login attempts, bandwidth usage by hosts, and accounts created by user, and so on.

For a complete list of reports, click Reports | **Cross Device** under the Device Monitoring Reports section on the left panel menu.

Database Reports

These reports provide information on database errors and warnings.

Firewall Reports

These reports provide information on firewall activity, such as denied connections by port, address, and hour.

Identity Management Reports

This report provides information on the number of connections per user as reported by the Identity Management devices in your network.

IDS-IPS Reports

These reports provides information on activity involving Intrusion Detection and Prevention Systems, such as alert count by device, port, severity, top alert destinations, worm infected systems, and so on.

Network Reports

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

Operating System Reports

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

VPN Reports

These reports provide information on activity involving VPN connections, including authentication errors, connection information such as counts, accepted and denied by address, and so on.

User Reports

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

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

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

Figure 5-3 User Reports

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

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

For information how to run, view, and publish reports, see [“Running, Viewing, and Publishing Reports” on page 95](#).

For information on using the Report Designer to create reports, see [“Designing Reports” on page 106](#).

For information on deploying Custom Packages, see [“Deploying a Report Package” on page 159](#).

Reports Home Page

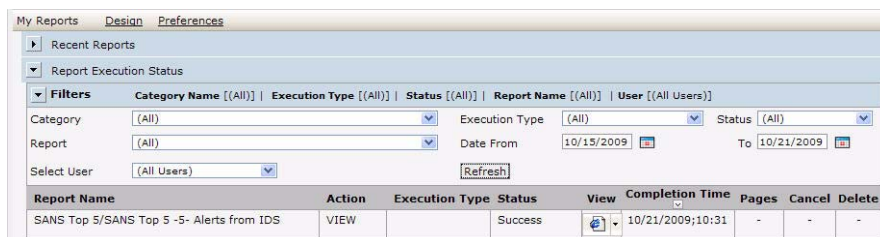
If you click the **Reports** tab from elsewhere in the Logger UI, the Reports Home page is displayed. Also, if you click **Dashboard** on the Reports left panel menu from within Reports, the Reports Home page is displayed.

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 83](#)).

To get started by creating a dashboard to show as your default Reports Home page, see [“Using the Dashboard” on page 82](#), [“Designing Dashboards” on page 83](#), and [“Setting Dashboard Preferences” on page 94](#).

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 are displayed, however, you can restrict the list by defining filter criteria.

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



To get started by running and viewing reports, see [“Running, Viewing, and Publishing Reports” on page 95](#) and [“Scheduling Reports” on page 155](#).

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.

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.

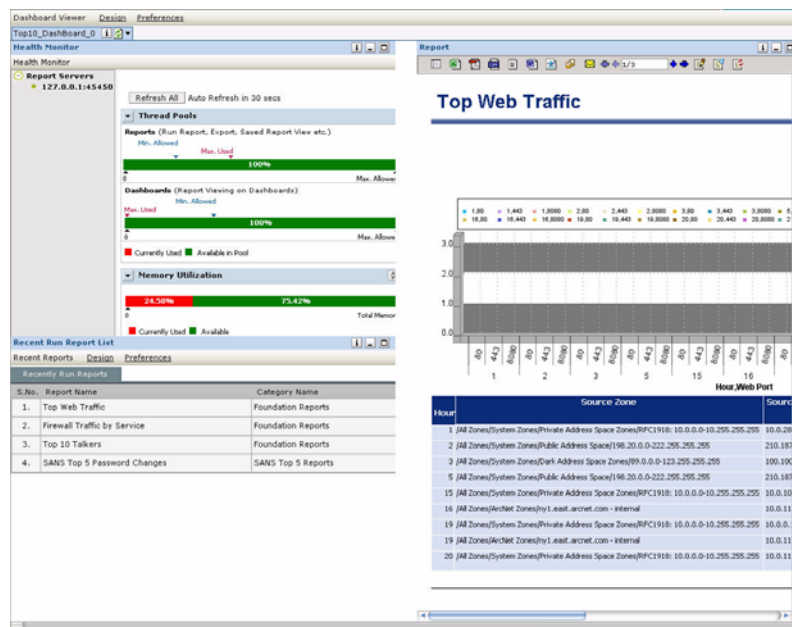


Figure 5-4 Dashboard View with Health Monitor and Reports



Tip

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 83](#) and [“Scheduling Reports” on page 155](#).

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 95](#).

Designing Dashboards

Use the **Dashboard Designer** page to create a new dashboard, name it, add items to it, and design the layout. You can design and save multiple dashboards, but only one at a time can be set as the default Dashboard **View** for the Reports home page. Other dashboards can be saved for later use. Each dashboard can include multiple items (reports, use cases, and Web links).

To access the Reports Dashboard Designer, click **Design** on the Dashboard navigation bar.

Click "Design" to open
Dashboard Designer

Recent Reports **Design** Preferences



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

What items can a dashboard include?


The following information is available for placement on a dashboard:

- **Reports**; any report can be included. The dashboard will show the latest published version of the report. Reports must be published in order for the report data to be accessible to users on the Dashboard View. If no published results are available for a report on a dashboard, the Dashboard View will display a message indicating this. When the report is published, a refresh of the Dashboard view will display the report.
- **Common Use Cases**, including a Report List, Saved Report List, Health Monitor, Recent Run Report List, Quick Job List, Schedule History, and Audit Log.
These are provided as dashboard elements so that users access a use case without leaving the Dashboard View page.
- **External Links**; that is, any URL(s) that you want on-screen as a part of a particular Dashboard View

Quick Start - Creating a New Dashboard


The high-level steps to create a dashboard are described here. A detailed explanation of each of these steps is provided in the topics that follow.

- 1 Add a new, empty dashboard.

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



- 2 Under Dashboard Properties, specify a **Name** for the new dashboard and other dashboard properties, as needed.
- 3 Place items onto the dashboard in the **Widgets** provided in the Layout area.

To do this, click-and-drag an item from the **Dashboard Items** list on the left into an

Empty Widget to the right. Alternatively, click  next to the dashboard item you want to place the item on an empty widget.

You can also click-and-drag an item onto a currently occupied widget if you want to replace an item in a widget with a different one.



To get a new widget, click  (Divide Widget Horizontally) or  (Divide Widget Vertically) on a widget to split it into two widgets. The original widget remains a new empty widget is placed on the dashboard layout.

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

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




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

See ["Selecting a Dashboard View" on page 94](#) 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.

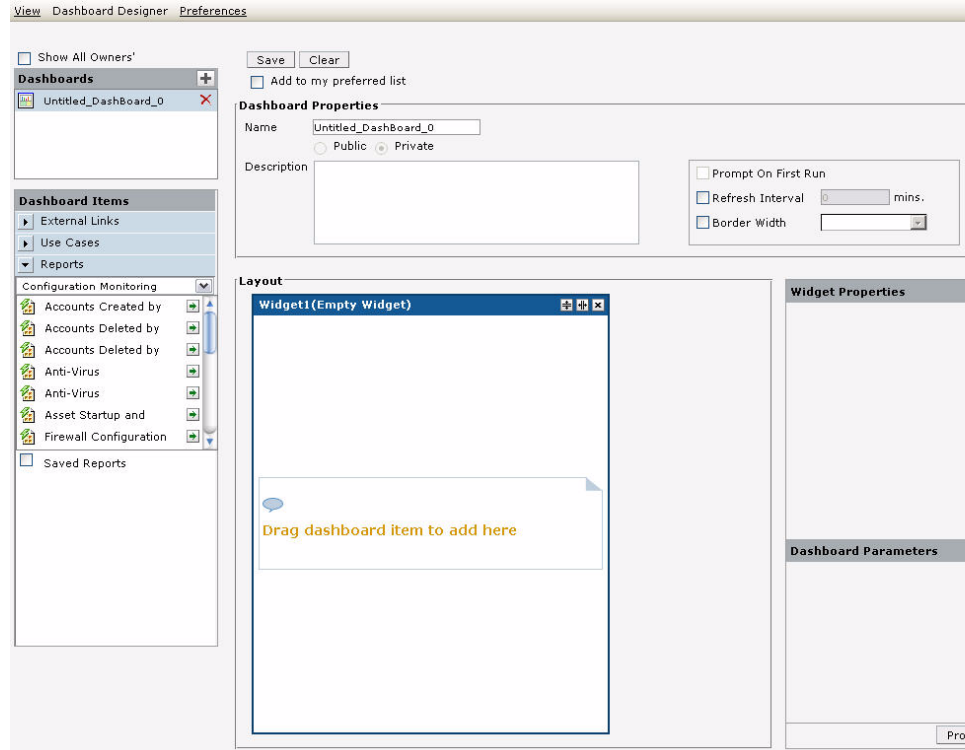


Figure 5-5 New Dashboard Layout

- 2 In the **Name** field, specify a unique name to identify the dashboard.
- 3 Specify Dashboard Properties. (For details, see [“Dashboard Properties” on page 86.](#))
- 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 94.](#))



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

The screenshot shows the 'Dashboard Properties' dialog box. It has a 'Name' field with the text 'Untitled_DashBoard_0'. Below it are radio buttons for 'Public' and 'Private', with 'Private' selected. There is a large 'Description' text area. To the right, there are three checkboxes: 'Prompt On First Run' (unchecked), 'Refresh Interval' (checked) with a value of '0' and the unit 'mins.', and 'Border Width' (checked) with a dropdown menu.

Figure 5-6 Reports Dashboard Properties

The Dashboard Properties are described in the following table.

Table 5-1 Dashboard Properties Description

Property	Description
Name	Name of the dashboard.
Description	Descriptive information about this dashboard.
Refresh Interval	Sets the time in minutes to refresh results for all the reports on dashboard. Note: Reports must be run and published first in order for the results to be accessible on a dashboard view. A <i>refresh</i> or <i>auto-refresh</i> on a dashboard simply updates the dashboard display with the most recently published results; it does not run the report. We suggest using the dashboard refresh interval in conjunction with scheduled reports to ensure that report results are always published and retained long enough to be available to dashboards. (See also, “Scheduling Reports” on page 155.)
Border Width	Check (select) the checkbox to have border around the widgets of the dashboard and select the width from drop-down box.



Creating Widgets

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

To get a new widget, simply split the existing widget either vertically or horizontally, depending on the layout you want. (See [“To get a new widget” on page 87.](#))

You can also delete widgets you do not need. (See [“To remove a widget” on page 87.](#))

To get a new widget

To get a new widget, click  (Divide Widget Horizontally) or  (Divide Widget Vertically) on a widget to split it into two widgets. The original widget remains a new empty widget is placed on the dashboard layout.

To remove a widget

To remove a widget, click  (Remove Widget) on the widget you want to remove.

Placing Dashboard Items on the Layout

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



Figure 5-7 Dashboard Items

To place a dashboard item, click to expand the menu for the type of item you want, click-and-drag an item onto a widget in the Layout area, and specify widget properties as needed. (Widget properties vary depending on the type of item you place on the dashboard.)

The following sections provide more detail on placing each type of dashboard item and setting appropriate widget properties.

Placing a Report on a Dashboard

The following sections describe in detail how to place and configure reports on dashboards, including setting widget properties, report parameters, and dashboard parameters.



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.

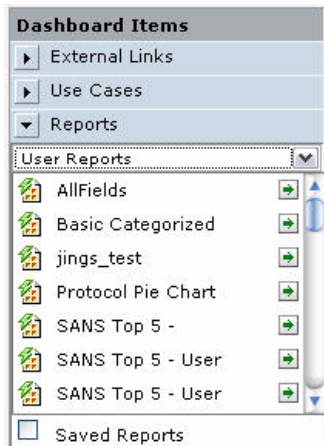



Figure 5-8 Reports under Dashboard Items

- 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 88.](#))

Widget Properties for Reports


Widget Properties	
Report Name	SANS Top 5 -
Refresh Interval (in mins.)	15
Format	HTML
Auto Refresh	YES
Toolbar	MULTIPAGE
Instance Navigation	NO
Link Widgets	...
Description	

Figure 5-9 Widget Properties for Reports on a Dashboard

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

Table 5-2 Widget Properties for Reports on a Dashboard

Property	Description
Report Name	The name of report that occupies this widget.
Refresh Interval (in minutes)	<p>This is the time in minutes. Refresh will take place at the end of specified number of minutes. For example, if you want the report results to refresh every 15 minutes, set the Refresh Interval to 15.</p> <p>Note: Reports must be run and published first in order for the results to be accessible on a dashboard view. A <i>refresh</i> or <i>auto-refresh</i> on a dashboard simply updates the dashboard display with the most recently published results; it does not run the report. We suggest using the dashboard refresh interval in conjunction with scheduled reports to ensure that report results are always published and retained long enough to be available to dashboards. (See also, “Scheduling Reports” on page 155.)</p>
Format	<p>Select the output format in which you want to view the report. Available options are:</p> <ul style="list-style-type: none"> • HTML • Acrobat PDF • Interactive
Auto Refresh	<p>Enables or disables auto-refresh option.</p> <ul style="list-style-type: none"> • Select Yes to refresh the reports as per Refresh Interval. • Select No to view the report generated when dashboard was loaded for the first time. <p>Note: Reports must be run and published first in order for the results to be accessible on a dashboard view. A <i>refresh</i> or <i>auto-refresh</i> on a dashboard simply updates the dashboard display with the most recently published results; it does not run the report. We suggest using the dashboard refresh interval in conjunction with scheduled reports to ensure that report results are always published and retained long enough to be available to dashboards. (See also, “Scheduling Reports” on page 155.)</p>
Toolbar	<p>Specifies Toolbar settings.</p> <ul style="list-style-type: none"> • Select Yes to always show toolbar. • Select No to never show the toolbar. • Select MultiPage to show the toolbar only for multi-page reports. <p>The Multipage setting is applicable to HTML and Interactive output formats.</p>
Instance Navigation	<p>Sets whether to include a report navigation feature on the dashboard.</p> <ul style="list-style-type: none"> • Click Yes to provide a drop-down menu that allows Dashboard users to select a saved report and view it. • Click No if you do not want to provide this feature on the dashboard.

Property	Description
Link Widgets	Click  to bring up a Link Widget dialog in which you can specify a link from any of the charts in the report in this widget to another widget. See “Linking Widgets” on page 90 .
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.

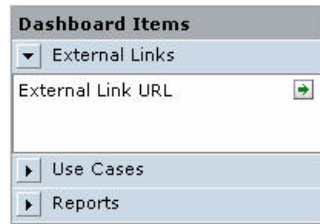
- 5 Click **OK** to save the settings and close the dialog.


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 91.](#))

Widget Properties for Use Cases

 A screenshot of the 'Widget Properties' dialog box. It has a title bar 'Widget Properties'. The dialog contains several fields: 'Name' with the value 'Health Monito...', 'Refresh Interval (in mins.)' with a value of '15', 'Auto Refresh' with a dropdown set to 'YES', and 'Show Scrollbar' with a dropdown set to 'NO'. There is also a 'Description' text area which is currently empty.
Figure 5-12 Widget Properties for Use Cases on a Dashboard

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

Table 5-3 Widget Properties for Use Cases on a Dashboard

Property	Description
Name	The name of use case that occupies this widget.
Refresh Interval (in minutes)	This is the time in minutes. The use case page is refreshed at the specified interval.
Auto Refresh	Enables or disables auto-refresh option. <ul style="list-style-type: none"> • Select Yes to refresh the use case as per Refresh Interval. • Select No to execute only once, when the dashboard is loaded.
Show Scroll Bar	Select Yes to get scroll bar if use case does not fit in widget width.
Description	Description of the widget.

Property	Description
Category	This option appears when Report List, Saved Report List or Quick Job List is placed on widget. Select the category to carry out respective task (get a list of reports in selected category, get a list of saved reports or quick job lists for selected report).
Report	This option appears when Saved Report List or Quick Job List is selected. Select the report for which saved report list or quick job list is to be viewed.

The use cases displayed in the list will depend on the permissions associated with your user group. Other properties are displayed based on the use case.

Placing an External Link on a Dashboard

The following sections describe in detail how to place and configure an external link on a dashboard.

To place a link on a dashboard:

- 1 Under Dashboard Items, click **External Links** bar to expand the list.

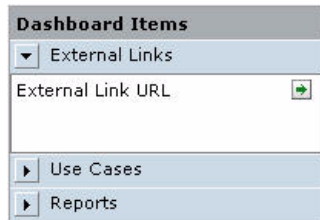



Figure 5-13 External Link under Dashboard Items

- 2 Click and drag a External Link URL object to the widget in which you want to place it.

Alternatively, click  next to the dashboard item you want to place the item on an empty widget.

The External Link URL object is displayed in the widget in the Layout area.

- 3 Set Widget Properties for the URL. (See ["Widget Properties for External Links" on page 92.](#))

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' with a dropdown arrow), 'Show Scrollbar' (set to 'NO' with a dropdown arrow), 'URL' (set to 'www.arcsight.com' with a text input and a small icon), and 'Description' (an empty text area).

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

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

Table 5-4 Widget Properties for External Links on a Dashboard

Property	Description
Name	The name of use case that occupies this widget.
Refresh Interval (in minutes)	This is the time in minutes. The use case page is refreshed at the specified interval.
Auto Refresh	Enables or disables auto-refresh option. <ul style="list-style-type: none"> Select Yes to refresh the URL as per Refresh Interval. Select No to execute only once, when the dashboard is loaded.
Show Scroll Bar	Select Yes to get scroll bar if external link does not fit in widget width.
Description	Description of the widget.
URL	Specify the URL for this widget. If you want to add multiple Web pages to the dashboard, use a different widget for each URL.

Swapping Items on Widgets

You can swap items placed in widgets. To do this, click and drag the item to the widget where you want to place it.

Click and drag an item to a different widget to swap placement of the two items on the page.



Figure 5-15 Swapping Widgets on a Dashboard Design

In the above example, the Recent Run Reports List item is swapped to the position of the the External Link URL, which is then swapped to with the Health Monitor item, which will end up at the top of the dashboard.

Setting Dashboard Preferences

In Dashboard Preferences, you can specify:

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

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

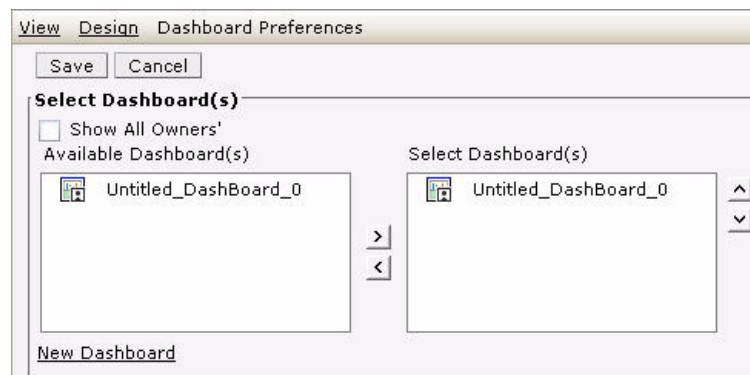


Figure 5-16 Dashboard Preferences

Working with Available Dashboards

The set or subset of dashboards shown under Available Dashboard(s) is based on your user group status and the selection status of **Show All Owners'** checkbox.

For example, it is likely that a user with Administrative status will be able to see more or all dashboards than a user with fewer privileges.

Also, if you limit the view to only your dashboards, the list will not include dashboards designed by other users.

- To access dashboards from all the users (designers), click (checkmark) the **Show All Owners'** checkbox.
- To view only your dashboards, click (uncheck) this checkbox.

Selecting a Dashboard View


Once you have created one or more dashboards, you can select one of them as the default display for the Dashboard **View** page, which also serves as the Reports home page.



You must have at least one dashboard in order to set a preference for the Dashboard View. For a quick summary of steps to create a dashboard, see ["Quick Start - Creating a New Dashboard" on page 84](#).

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 84.

- 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 83](#).

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 83](#).
- 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 76](#))

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

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



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

Best Practices

ArcSight Logger is designed to process events while running a report, but event processing has priority. Running a complex report while the event processing system is under load will result in report timeout rather than dropped events.

ArcSight recommends using the Scheduled Report feature so that reports are run during periods of light load. If an ad hoc report must be run, run it when the system is not under load.

For information on working with scheduled reports, see [“Scheduling Reports” on page 155](#).

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 76](#).)
- 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 90](#).) 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 155](#).)



The Search feature on the Logger “Analyze” page (described in [Chapter 4, Searching and Analyzing Events, on page 39](#)) does not search on resources such as reports. It searches only on events in the database.

Task Options on Available Reports

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

Figure 5-17 Task Options on All Reports

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





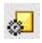



Note

Your access to various reports and report options (view, publish, edit, etc.) depends on the access rights associated with your user role and **Logger Report Group** affiliation. For example; depending on your access rights, you may have privileges to view some reports but not others, to view but not schedule or publish a report, or to view and schedule but not edit a report.

Access rights to report options are configured and managed with the **User/Groups** option on the Logger **System Admin** page.

For more information on Logger Report Group management, see [“Setting Access Rights on Reports” on page 122](#), and [“Groups” on page 251](#) in [“System Admin” on page 219](#).

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 98 and “Quick Run / Run In Background Report Parameters” on page 99.</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 98 and “Quick Run / Run In Background Report Parameters” on page 99.</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 98 and “Run Report Parameters” on page 101.</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 106, “Quick Run / Run In Background Report Parameters” on page 99, and “To publish a report” on page 103.</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 120.</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 77](#), [“Solution Reports” on page 79](#), and [“User Reports” on page 80](#).

About the Pagination of Reports

If a report contains more columns than can be displayed horizontally across a screen using the default width specified in the report query (Reports > Design > Queries), the report is paginated horizontally such that additional columns are displayed on the following pages. For example, if a report contains 45 columns and only 5 can be displayed on each screen, the report would be paginated such that Page 1 displays columns 1 through 5, Page 2 displays columns 6 through 10, Page 3 displays columns 11 through 15, and Page 9 displays columns 40 through 45. Consequently, if the report contained more rows than can be displayed vertically in one screen, the second screen of rows would be displayed starting at Page 10.

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

To run and view a report

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


Note: Even if you selected Run Report in the previous step, you can run a report in the background after setting the Run Report parameters.

The report output is displayed in the specified format (HTML, PDF, or other).

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



Figure 5-18 Results of a Report Run

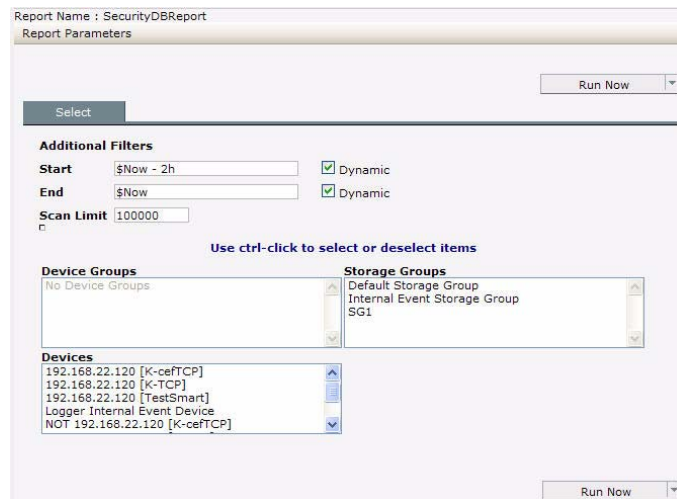
At this point, the results of this report generation is available as a file for viewing only by you. If you close the file without saving or publishing it, the results are no longer available.

If you want to make the results of this run available for others, publish it. To do this, leave the file open, click  (Publish report), and follow the steps in [“Publishing Reports” on page 103](#).

For information about other delivery options available to you at this point, see [“Report Delivery Options” on page 104](#).

Quick Run / Run In Background Report Parameters

When you click or  (Quick Run) or  (Run in Background) for a report, the report will run with the data filters specified in the deployed report. You still get options to select additional filters on timeframe and storage groups over which the report runs.


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

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

Table 5-6 “Quick Run” / “Run in Background” Report Parameters

Option	Description
Start	<p>Specify the starting point for the data gathering from the events database.</p> <p>By default, the start time is specified with a dynamic data expression (\$Now - 2h).</p> <p>You can modify the dynamic expression to specify a different dynamic start time, or disable Dynamic and use the calendar options to specify a fixed start time.</p>

Option	Description
End	<p>Specify the ending point for the data gathering that is some time after the starting point.</p> <p>Keep in mind that large time spans can mean large amounts of data, which can affect system performance.</p> <p>By default, the end time is specified with a dynamic data expression (\$Now).</p> <p>You can modify the dynamic expression to specify a different dynamic end time, or disable Dynamic and use the calendar options to specify a fixed end time.</p>
Scan Limit	<p>Specify the number of events to scan.</p> <p>When you specify a scan limit, the number of events scanned for <i>manually</i> run reports is restricted to the specified limit. Doing so results in faster report generation and is beneficial in situations when you only want to process the latest N number of events in the specified time range instead of all the events stored in Logger.</p> <p>The scan limit is 100,000 by default. If you set the scan limit to 0 (zero), all events are scanned.</p> <p>This setting does not apply to the scheduled reports.</p>
Device Groups	Select the device group(s) on which to run the report query, if any. (See “Selecting Device Groups, Storage Groups, or Devices” on page 100.)
Storage Groups	Select the storage group(s) on which to run the report query. (See “Selecting Device Groups, Storage Groups, or Devices” on page 100.)
Devices	Select the device(s) on which to run the report query. (See “Selecting Device Groups, Storage Groups, or Devices” on page 100.)

Selecting Device Groups, Storage Groups, or Devices

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

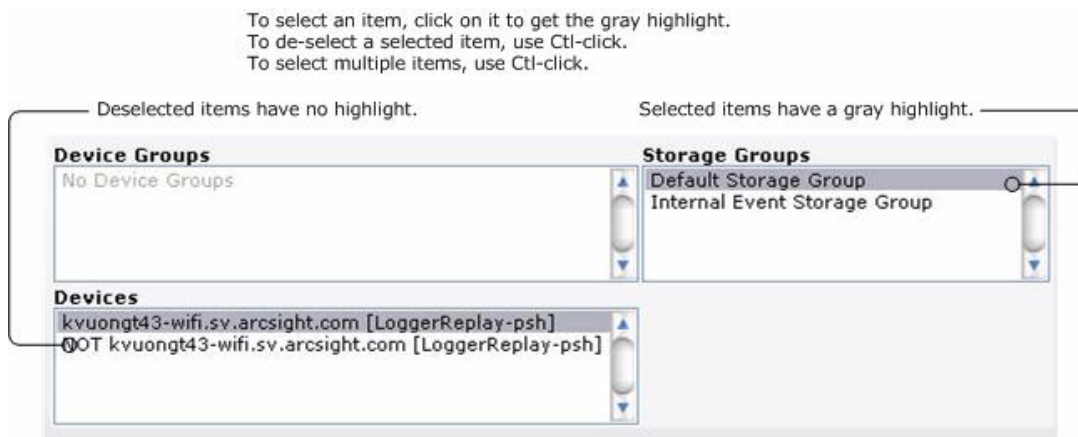
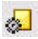


Figure 5-20 Selection Model for “Quick Run” or “Run in Background” Report Scope of Storage and Devices

- Items with a gray highlight are selected and will be included in the report query when the report is run.
- Items that are not highlighted are de-selected and will not be included in the report query.
- To select an item, click on it. To select multiple items in a list, use Ctl-Click.
- To de-select a currently selected item, use Ctl-Click.

Run Report Parameters

When you click  (Run Report) button for a report, you get additional options (beyond what you get for a Quick Run or for Run in Background) to choose a file format, specify pagination, and to modify the data filter criteria for only this run of the report.

If you run the report without specifying any override run-time parameters here, the report is generated with the defaults specified at design time for this report. You can run a report in the background after specifying the Run Report parameters, as indicated in the following figure.

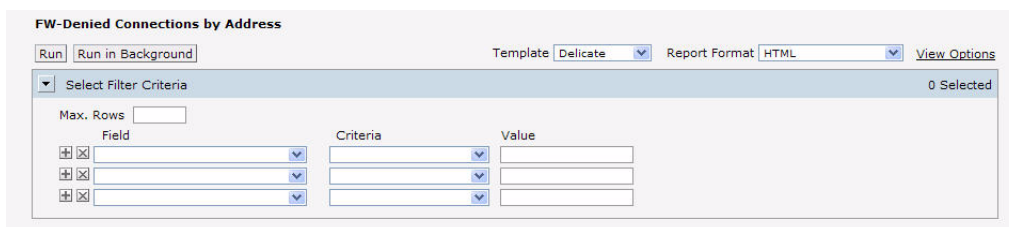


Figure 5-21 “Run” Report Parameters

The following table describes “Run” report parameters.

Table 5-7 Run Report Parameters

Option	Description
Report Format	<p>Specify a file type or “format” option of the output, and toggle on or off the multi-page option if applicable to the chosen file format.</p> <p>ArcSight strongly recommends using the multi-page option for all reports. This option is the default.</p> <p>For descriptions of report format see “Report File Formats” on page 102</p>

Option	Description
Select Filter Criteria	<p>Provides options to define filters, or modify default filters if any are already built in to the report.</p> <p>The filter expression is applied when the report runs, narrowing the focus of the report to the specified criteria.</p> <p>For example, you could set the filter criteria on a report on Top Password Changes to report only on password changes related to specified username(s) or involving specified IP address(es).</p> <p>For details on how to create these filters (with Field, Criteria, and Value fields), see “Select Filter Criteria” on page 112 in “Designing New Reports” on page 110.</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 99](#) 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.)
- Text
- Microsoft Word
- Interactive
- XML
- Raw Text (Comma separated file with headers and data values as received from database.)

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 122](#) 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 98](#)), the output results for that run of the report are saved for subsequent.



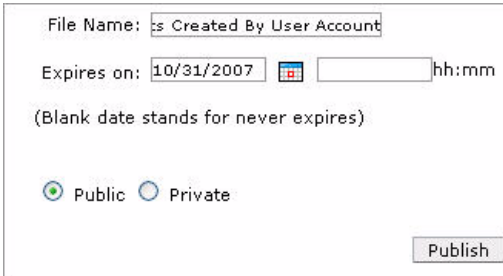
You configure *scheduled reports* to publish after each scheduled run. The publish options for scheduled reports are the same as for *on-demand reports* described here. For more about scheduled reports, see [“Scheduling Reports” on page 155](#) and [“Add Report Job Settings” on page 158](#).

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 106](#).

Report Delivery Options

When you run a report (as described in [“Running and Viewing Reports” on page 98](#)), many options are available to you in terms of delivery options for generated output.

The most common next step is to publish the resulting report (described in [“Publishing Reports” on page 103](#)), but you can also save the report output to a file, e-mail it to other users, refresh the results, change the output format, and so forth.

Refreshing a Report

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


E-mailing a Report

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



You can also configure these same e-mail options on *scheduled reports*, as described in [“Scheduling Reports” on page 155](#) and [“Add Report Job Settings” on page 158](#).

To e-mail a report

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

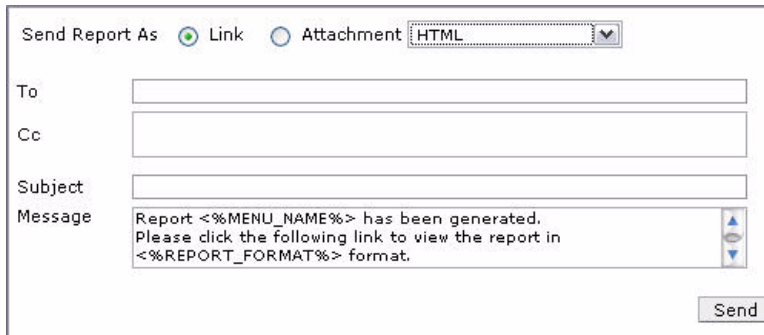


Figure 5-23 E-mail Report Settings

The following table describes the e-mail report options.

Table 5-9 E-mail Report Settings

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







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

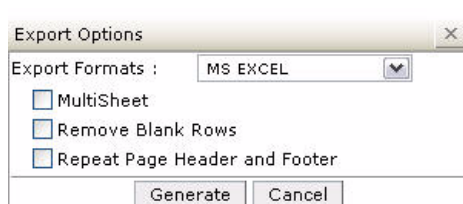
- 3 Click **Send** to send the report.

Exporting and Saving a Report

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

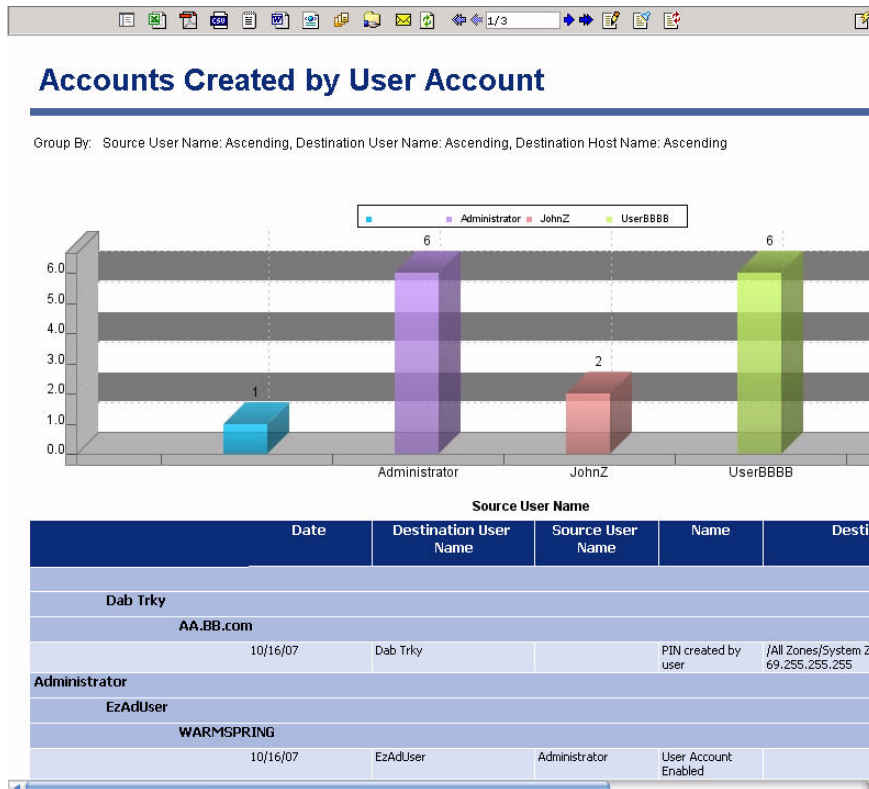
To export and save a report

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




Depending on the Export Format you choose, other settings are displayed as appropriate. Configure the export, and click **Generate**.

- When the report is displayed, you have the option to save it as a file locally or elsewhere just as you would any other file.



Viewing the Output of a Published Report

- Navigate to the report for which you want to view output results. (See [“Finding Reports”](#) on page 96 if you need help locating a report.)
- Click the “Published” button  (Navigate to list of published outputs for this report) next to the report you are interested in.

Saved Report List : Top 10 Talkers

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

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

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


- View report outputs in various formats (HTML, PDF, Microsoft Word, and so on)
- Delete the selected instance of the generated report

Designing Reports

You can use the Logger Report Designer to design simple columnar reports as well as mixed reports with embedded charts and matrices. For columnar reports, the Report Designer provides options for setting up filters, grouping, totals, and sort order to create a full-featured report.

Opening the Report Designer

To open the Report Designer to create a new report from scratch, click Design | **New Report** on the Reports left menu bar.

To open the Report Designer to edit an existing report, click the Edit button  for a report in a reports list. (See [“Report Groups” on page 76](#) and [“Task Options on Available Reports” on page 96](#) for more information on available reports and how to get to their task option buttons, respectively.)

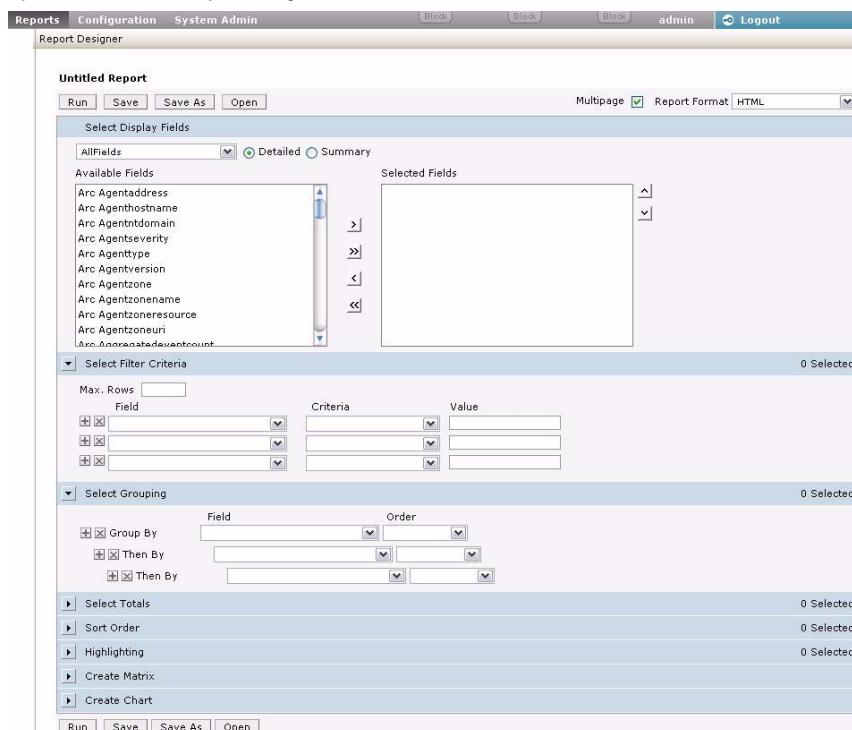


Figure 5-25 Report Designer (click **New Report** or edit an existing report)

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 107](#).

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 110](#).

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 76](#) 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.

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 110.](#))

Figure 5-28 Editing a Report

Designing New Reports

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

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



This brings up the Report Designer with a blank template.

The screenshot shows the 'Report Designer' window with a title bar containing 'Reports', 'Configuration', 'System Admin', and user controls. The main area is titled 'Untitled Report' and includes buttons for 'Run', 'Save', 'Save As', and 'Open'. A 'Multipage' checkbox is checked, and the 'Report Format' is set to 'HTML'. The 'Select Display Fields' section has a dropdown for 'AllFields' and radio buttons for 'Detailed' (selected) and 'Summary'. Below this is a list of 'Available Fields' including Arc Agent address, host name, domain, severity, type, version, zone, zone name, zone resource, and zone URI. To the right is a 'Selected Fields' list. The 'Select Filter Criteria' section shows '0 Selected' and a table with columns for Field, Criteria, and Value. The 'Select Grouping' section also shows '0 Selected' and options for Group By, Then By, and Order. At the bottom are sections for 'Select Totals', 'Sort Order', 'Highlighting', 'Create Matrix', and 'Create Chart', each with a '0 Selected' status. The bottom of the window has 'Run', 'Save', 'Save As', and 'Open' buttons.

The following sections explain how to use the Report Designer.

Report Save, Run, and Template Options

- Click **Run** to test the current version of the report.
- Click **Save** to save the report.
- Click **Save As** to save it under a different name.
- Click **Open** to open another report in the Report Designer.

General Report Settings

Set your preferences for pagination, layout and report output format as described below.

Table 5-10 General Report Design Settings

Option	Description
Template	<p>Select the template to apply to this report. The templates drop-down menu shows supplied templates, and any custom templates you may have added.</p> <p>See “Applying Report Template Styles” on page 154 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 102 .
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)

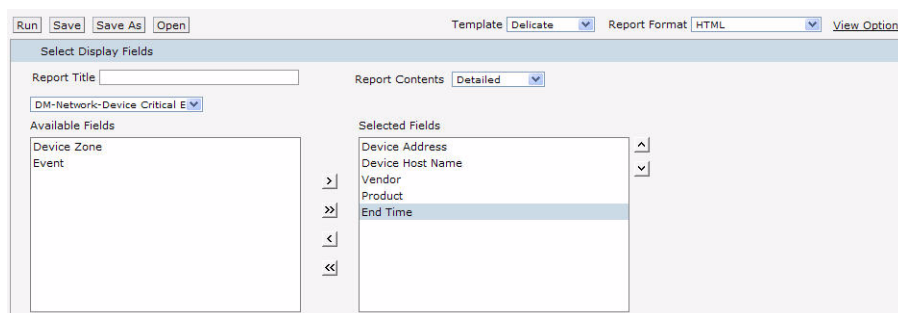


Figure 5-29 Report Display Fields

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 122](#).)



Note

In addition to the fields in the WHERE clause of the query, the fields in the SELECT clause also need to be indexed to yield faster report generation. For more information about indexing fields, see [“Indexing” on page 63](#).


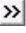

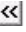


Under **Select Display Fields**, enter a meaningful title for the report (in the Report Title field) and select whether the report contents should be Detailed or Summarized (in the Report Contents field). The report title is the text that appears as the title on top of a report.

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 122](#). All available queries, including new queries you create, show up in the drop-down menu in the Select Display Fields section of the Adhoc Report Designer.

Select Filter Criteria



Field	Criteria	Value
 		
 		
 		

Figure 5-30 Report Filter Criteria

You can set filters on the results of the base query with logical expressions to narrow the focus of the report results. For example, you could set the filter criteria on a report on Top Password Changes to report only on password changes related to specified username(s) or involving specified IP address(es).

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

Filter criteria defined as part of a report design is built in and saved with the report. When other users run the report, they will get the built-in filters by default



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 101](#).



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.

Select Filter Criteria*



Max. Rows

Field	Criteria	Value
<input type="checkbox"/> Time *	Is	<input type="text"/>
<input type="checkbox"/>		<input type="text"/>
<input type="checkbox"/>		<input type="text"/>

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 131](#) under [“Setting up Queries” on page 122](#).

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 114), 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 131 under "Setting up Queries" on page 122.
Criteria	Select a logical operator. (For example, Is, Is Not, Starts With, Ends With, Contains, and so forth.)
Value	Select a value to complete the conditional filter expression.

Select Grouping

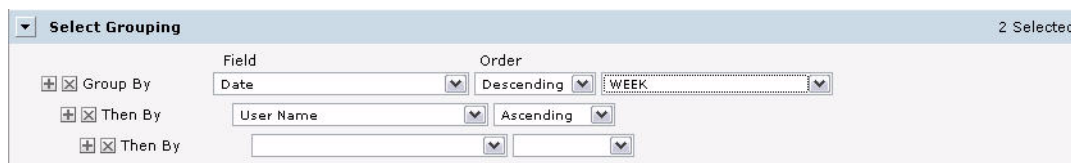


Figure 5-31 Grouping Items by Field in a Report

Define group requirements to arrange the report information into logical groups based on particular fields you are interested in. You can create multiple groupings for report results.

For example, if the report uses a query that includes a Date field, you can group results by date. You could add additional statements to group by "User Name", "Source Address", "Destination Address", and so forth, depending on what other fields are available in the report query.





Note

If you select set Max. Rows under **Select Filter Criteria** (as described in "[Select Filter Criteria](#)" on [page 112](#)) and also specify grouping, you may get a different result than if you just specified Max. Rows without grouping.

To define a group

- 1 Select a field by which you want to group (as described in [Table 5-20 on page 158](#)).
- 2 Select the order of arrangement of group (as described in [Table 5-20 on page 158](#)).

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 116](#).

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

Figure 5-33 Sort Order for Items in a Report

You can have up to three levels of sorting.

To specify a sort order

- 1 In **Field** (on the right of Sort By), select the field on which you want to sort the report.
- 2 In **Criteria** (in the same row), select the sort criteria.
- 3 Repeat [Step 1](#) and [Step 2](#) by providing values in the Then By rows to specify more sorting criteria.



Highlighting

A report can include multiple levels of “highlighting” for specified fields. Highlighted items can serve as visual alerts on generated reports when specified set conditions are satisfied.

Figure 5-34 Highlighting Items in a Report

To set up a highlight

- 1 In **Highlight**, select the field that should be highlighted. Select Entire Row to highlight entire record.
- 2 In **Using Style**, select the style to be applied to highlight it.
- 3 Select **Alert** check box to receive a visual alert on report viewer.
- 4 In **Field**, select the fields which will be evaluated for highlight (alert).
- 5 In **Level**, select the level at which the selected field should be evaluated:
 - ◆ DETAIL evaluates each row (record)
 - ◆ REPORT evaluates at the end of report
 - ◆ Respective groups evaluate at the end of each group
 - ◆ PAGE evaluates at the end of the page
- 6 When REPORT or PAGE is selected in Level, select a Function to be applied.
- 7 Select **Criteria** and specify its **Value**.

Click  (Remove Condition) on the left of the criteria entry to delete an entry. Click  (Add Condition) to add another entry.


Create Matrix


You might choose to include a matrix in your report, since it presents a summary of data. Make sure that the appropriate query object is selected (under “Select Display Fields”).



Figure 5-35 Adding a Matrix to a Report


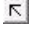
To create a matrix



- 1 To place a field in Row or Column, click the field and drag it to the **Row Fields** or **Column Fields** boxes.
- 2 To place a field as a cell (summary), click the field and drag into the **Summary Fields** box.
- 3 Select a **Function** from the drop-down menu provided for a field placed in **Summary Fields**.
- 4 Optionally, for numeric or date fields in columns or rows, specify a **Group By** function in the drop-down menu provided.
- 5 Optionally, for fields in columns or rows, check **Totals** checkbox to get total row / column.

Select a field and click  to add that field to the matrix as one of the **Column Fields**.

Select a field in Column Fields and click  to remove it from the matrix.

Select a field and click  to add that field to the matrix as one of the **Row Fields**. Select a field in Row Fields and click  to remove it from the matrix.

Select a field and click  to add that field to the matrix as one of the **Summary Fields**. Select a field in Summary Fields and click  to remove it from the matrix.

To move a field up or down, select the field and click  (Move up) or  (Move down), to move the field in the respective direction.

To remove all settings and contents of the current matrix, click **Clear Matrix**.

Create Chart

For pictorial representation of summary data, you can add a chart on your report. Make sure that the appropriate query object is selected (under "Select Display Fields").

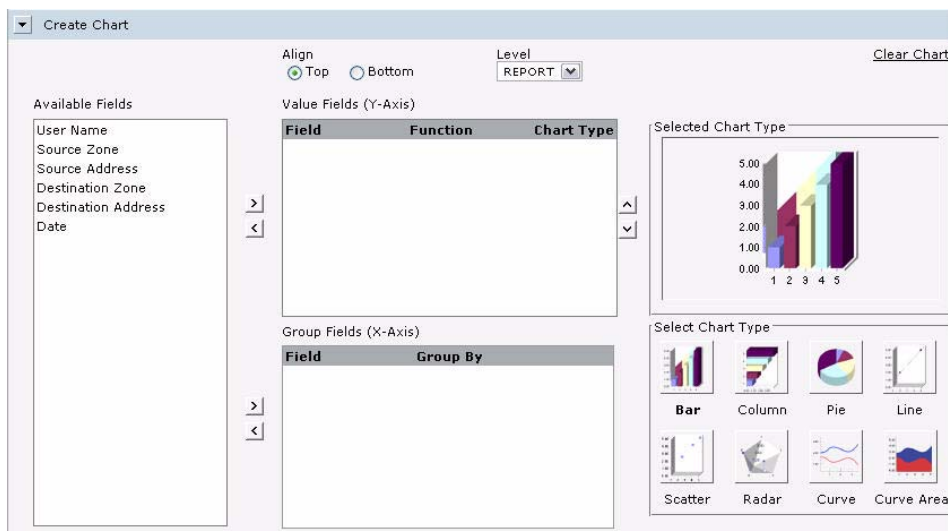


Figure 5-36 Adding a Chart to a Report

For pictorial representation of summary data, you may choose to have a chart on your report. Make sure that the right query object is selected (under Select Display Fields).

Chart Placement

Chart Placement is important when the chart is placed on the report along with other component. Specify chart placement preference using the Align option:

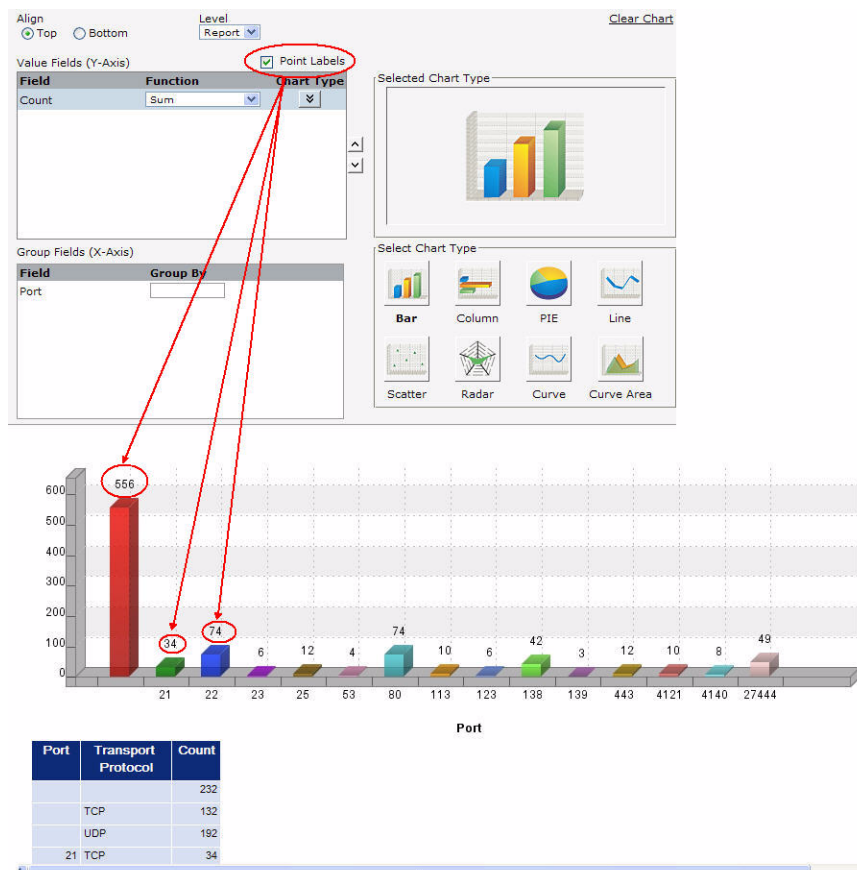
- Select **Top** to place the chart above other components.
- Select **Bottom** to place the chart below other components.
- In **Level**, select PAGE to plot chart having page level data. Select REPORT to plot chart from data that has come from entire report.

Chart Type


Select the chart type by clicking button (image) from **Select Chart Type** area. The image corresponding to the chart you select is displayed in the **Selected Chart Type** box at the top.

Select Point Labels



Select this setting to show the number of matches for a value of a field in a chart, as shown in the following figure.




Set Value Fields (Y-Axis)

- 1 Click and drag the Field in **Value Fields (Y-Axis)** box, or use the  button (Add field) to add the selected field.
- 2 Select summary function for the field.

- 3 To select a chart type different than the selected one, click the button on the right to open a box having chart types. Select the type you need.

Follow steps 1 through 3 above for each attribute to be placed as series. To re-position fields, select a field and click  (Move up) or  (Move down) as needed.


Set Group Fields (X-Axis)

- 1 Click and drag the field in **Group Fields (Y-Axis)** box, or use the  button (Add field) to add the selected field.
- 2 Select the method to group (for Numeric or date type).

You can specify groups in numeric fields. For example, to have groups of 10, specify 10 in Groups box.

You can specify groups in date fields. From the drop-down box select from Day, Week (Sunday to Saturday), Month, Quarter (Jan-Mar, Apr - Jun, Jul - Sep, Oct - Dec), Year.



To remove fields from Value fields (Y-Axis) or Group Fields (X-Axis), drag them out of the respective box or use the  button (Remove field) on selected fields.

To remove all settings and contents of the current chart, click **Clear Chart**.

Editing a Report

You can use the Report Designer to edit existing User Reports. (The supplied reports are not editable.)

To edit an existing report

- 1 From any Report list, click the Edit button  (Customize report) for the report you want to edit.

This brings up the Report Designer for the selected report.

- 2 Modify the report as needed (via the settings described in [“Creating New Reports” on page 107](#)).
- 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 121](#)).
- 4 Click **Save**.



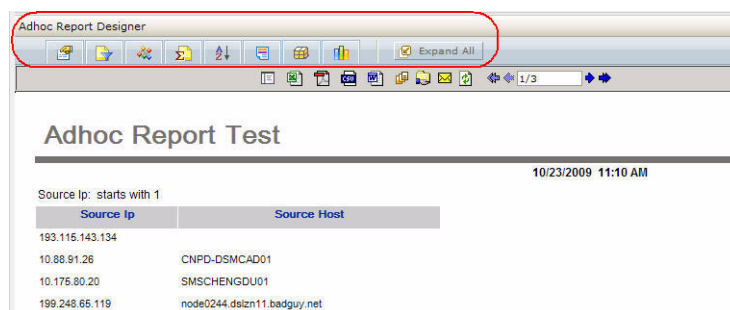
If a user has access rights to “view, run, and schedule all reports”, you can create **private** reports. If you do not have permissions to edit a **public** report that you want to modify but you do have permissions to create private reports, then you can save the public report as a private one and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

For more about publishing a report as “public” or “private”, see [Table 5-8 on page 103](#). For more about “access rights” on reports, see [“Setting Access Rights on Reports” on page 122](#).

See also [“Quick Start: Base a New Report on an Existing One” on page 107](#).


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 111 for more information.
	Specify filter criteria. See “Select Filter Criteria” on page 112 for more information.
	Specify grouping. See “Select Grouping” on page 114 for more information.
	Specify the summary (total) fields. See “Select Totals” on page 116 for more information.
	Specify sort order. See “Sort Order” on page 116 for more information.
	Set up highlighting. See “Highlighting” on page 117 for more information.
	Include a matrix. See “Create Matrix” on page 117 for more information.
	Create a chart. See “Create Chart” on page 118 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 96.](#))

Access rights on report options are configured and managed with the User/Groups option on the Logger System Admin page.

For more information on System Admin User/Group management, see [“Groups” on page 251 in Chapter 7, System Admin, on page 219.](#)

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 145.](#)

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 123](#), the query “SANS Top 5 - Password Changes” is selected.

The screenshot displays the 'Query Object List' window. On the left, a list of queries is shown, with 'SANS Top 5 - Password Changes' selected. The main area shows the details for this query, including its name, SQL statement, and field configurations. The SQL statement is as follows:

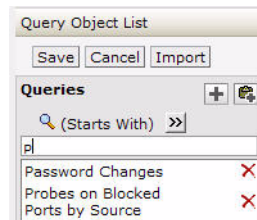
```
SELECT events.arc_destinationUsername "User Name",
events.arc_sourceZoneURI "Source Zone",
events.arc_sourceAddress "Source Address",
events.arc_destinationZoneURI "Destination Zone",
events.arc_destinationAddress "Destination Address",
events.arc_endTime "Date"
FROM events
WHERE events.arc_categoryBehavior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

The 'Fields' section on the right shows the configuration for the 'User Name' field, including its caption, data type (CHAR), width (10), and alignment (Left). The 'Lookup Values' section is also visible, with options for SQL, Predefined, and User Defined SQL.

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 76](#)).

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 59](#) 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



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**.



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.)

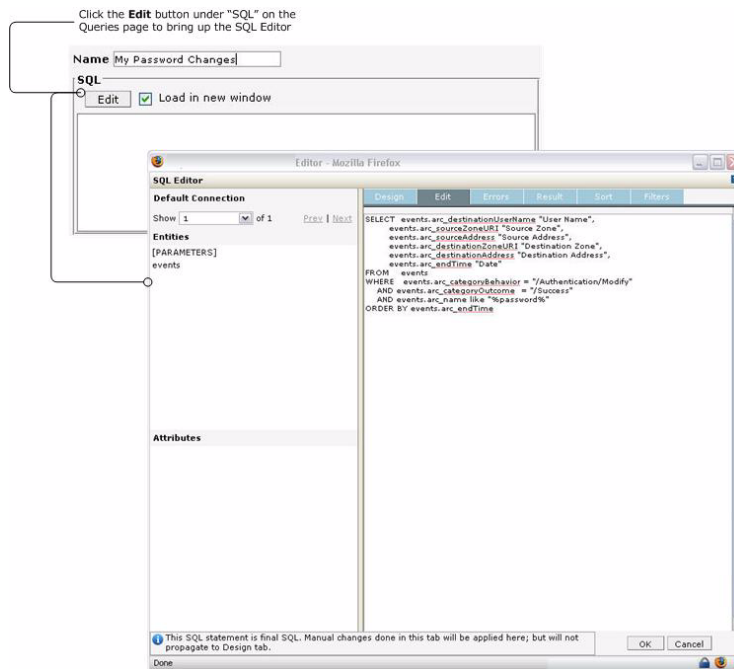


Figure 5-38 Query SQL Editor

- 4 Use the SQL Editor to define the query statement. (See [“Defining SQL in the Editor”](#) on page 137.) Report queries are case insensitive.



Caution

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

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

- 5 Click **OK** to temporarily save the SQL statement for the query.

The SQL you defined is displayed in the SQL box on the main Query Object List page.

Similarly, any fields you defined in the SQL Editor are displayed in the Fields list on the Query Object List page.

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



Caution

You must click **Save** on the main Query Object List page to save updates made in the SQL Editor as part of the query. If you navigate away from this page without clicking Save, edits you made in the SQL Editor since the previous Save will be lost.

Field Attributes and Properties

To set Field attributes, select a field under **Fields** and edit the properties associated with the that field.



Figure 5-39 Query Field Attributes


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

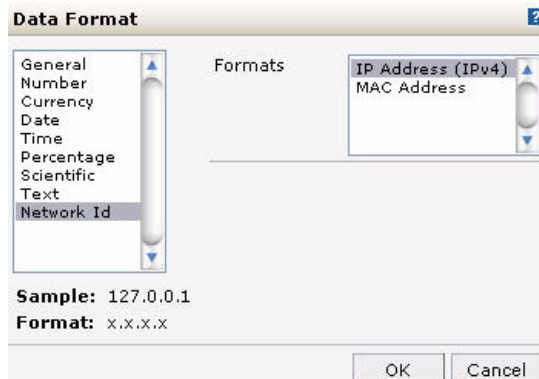
Table 5-13 Query Field Attributes

Option	Description
Field	Name of field (as received from data source).
Caption	The text that will appear as a caption when this field is selected for placement on the report.
Width	Number of characters for the selected field.
Align	Sets alignment for the selected field.
Hidden	Hides the associated field so that it is not available to be placed on report. This field will also not available for sorting as well as filtering.
Data Type	Sets the data type for field from Date, Character or Number. This is especially useful when field selected is XML type data source and you need to set it as number or date. Similarly when a field that is character (having numeric value) is supposed to be used in calculation.

Specifying Output Format for a Field

If you specify the output format for a query field here, at run-time the report output will adhere to the specified formatting.

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




- 3 Select the appropriate format and provide necessary values for that format.

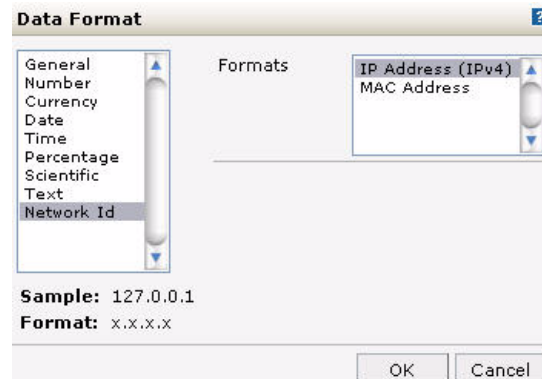
- 4 Click **OK** to accept the changes and close the dialog.

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

Specifying Input Format for a Field

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

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



- 3 Select the appropriate format and provide necessary values for that format.
- 4 Click **OK** to accept the changes and close the dialog.

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

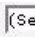

Format			
Width	<input type="text" value="15"/>	Output Format	<input type="text" value="x.x.x.x"/> 
Align	<input type="button" value="Right"/>	Input Format	<input type="text" value="x.x.x.x"/> 

Grouping Fields

When fields in a query object are grouped, they are displayed within a group header in the Report Designer. All fields in the group can be selected or removed from the report with a single click. Once groups are created, fields can be assigned to groups.

To create groups

- 1 In Group Label drop-down box 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.

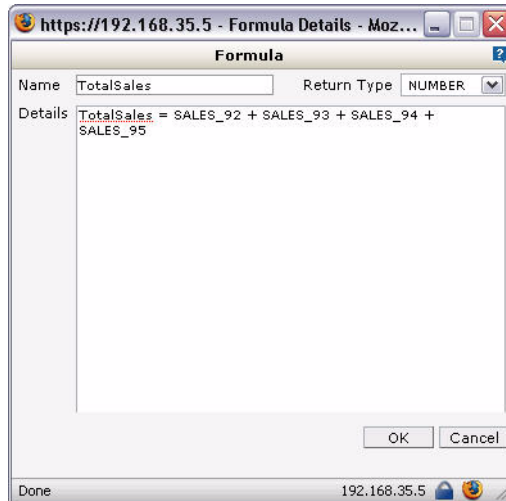
Formula Fields

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.

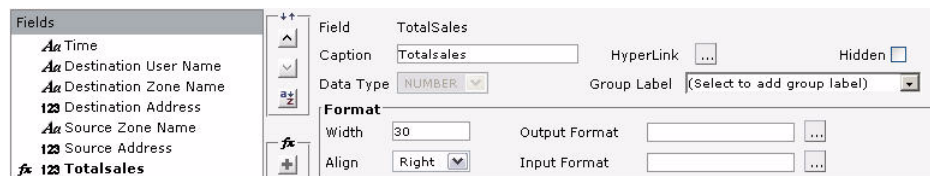
To create a formula

- 1 Click  (Add New) under the  tool palette (next to the Fields list) to get the Formula dialog.
- 2 Specify the formula and click **OK** to save the formula and close the dialog. (See [“Syntax for Formulas” on page 129.](#))





- ◆ 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 129.](#))
- 3 Click **OK** to save the work and close the dialog box.

The new formula is listed in the Fields list. Formula names shown in the list are pre-fixed by  to indicate they are formulas.



Positioning Formulas in Fields List

Select a formula and click  (Move up) or  (Move down) as needed to shift the position of the selected formula in the list.



Formulas further down in the list can use the formulas above them. Avoid the opposite; formulas higher in the list should not use the formulas below them.

Syntax for Formulas

The general syntax for formula is:

`FormulaName = formula`

where, `FormulaName` is the same as specified in the **Name** field on the Formula dialog.

In general, use JavaScript syntax to create formulas.

A formula can include:

- Field names
- Variables (custom or supplied)
- "if" and "nested if" constructs
- logical operators

For formulas that contain multiple statements, use a semicolon ";" as a separator between two statements.

Examples

```
NewForm1 = var a = 5 ; b = 3 ; if (a!=b) { f = a } {NewForm1=f}
```

```
TotalAmount = var total ; if (unitprice < 10 ) {total = unitprice*quantity} else {total = unitprice} {TotalAmount = total}
```

Importing Field Attributes

You can import field attributes from other Logger Report queries and apply the imported attributes to the currently selected field in your query. Leveraging attributes from existing queries can save time and re-work, and also serve as a learning tool.

You can import the following field attributes from one query into another:

- Captions
- Format (including Width, Alignment, Input, and Output formats)
- Data Types
- Hidden properties
- Group Labels
- Hyperlinks
- Lookup Values

You can select a field from which to import attributes from any of the saved query objects on the Logger Reporting server. Imported attributes can come from one field in another query, or from multiple fields.

To import field attributes

- 1 On the Queries list, select the query object into which you want to import field attributes (the “local” query you are editing), and click **Import** to bring up the Import Attributes dialog.

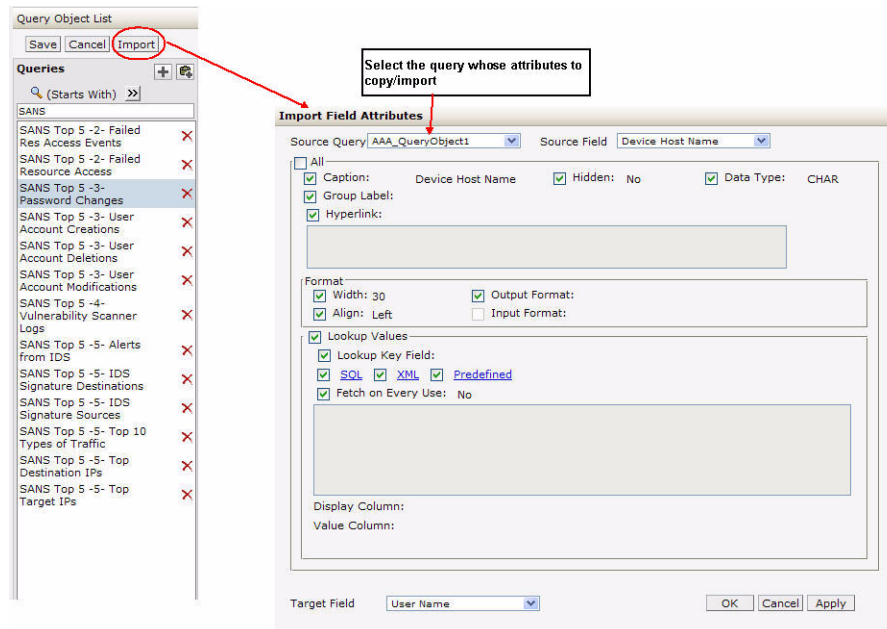


Figure 5-40 Importing Attributes from One Query into Another

- 2 From the **Query** drop-down menu on the Import dialog, select the query object from which you want to import field attributes (the “remote” query with the attributes you want to copy).

The Field drop-down menu is populated with the fields in the selected query.

- 3 In the **Field** drop-down menu, select the field in the remote query whose attributes you want to import (copy).
- 4 Select field attributes to import by clicking (checking) checkboxes for attributes you want.
- 5 From **Target Fields** drop-down menu, select the target field in your local query to which you want to copy the attributes.



Note

For successful field attribute import, consider the data types of source and target fields. For a valid import, data types for source and target fields must generally match.

Lookup values will not import if the data type of the target field is NUMBER.

- 6 Click **Apply** to save current selections and keep the dialog open.



Caution

A field attribute import cannot be revoked. Please make sure you are importing the right attributes before you click **Apply** or **OK**.

Click **Cancel** to abandon selections made after last Apply button and close the dialog. (Clicking Cancel will not revoke changes already applied.)

Click **OK** to save (apply) current selections and close the dialog.

- 7 To import selected field attributes to another target field, repeat these steps with a different target field selected.

To select from different fields in the same query, or different queries, choose different options for **Query** and **Field** at the top of the dialog.

Mandatory Filtering

You can provide built-in filters for a query when you want users to apply one or more filters when designing and running reports that use that query. Building in mandatory filtering at the query level can save unnecessary data transfer from the server database during report run time.


You can configure mandatory filtering in either of these ways:

- Mandating filtering on *any field*. Report designers can decide which field to filter on at report design time.
- Mandating filtering on a *specific field*. Report designers are required to filter on the specified fields at report design time.

To configure a query for mandatory filtering

- 1 Select (check) the **Mandatory Filtering** checkbox to enable mandatory filtering.



- 2 To specify a field for mandatory filtering, choose the field you want from the **On Field** drop-down menu. If you do not want to specify a field for mandatory filtering now, leave it as **Any**.
- 3 Click  (Add Filter) to get another row for mandatory filtering, and repeat [Step 2](#) above.

To remove a field filter

To remove a mandatory filter field, click  (remove button) next to the row you want to remove.

To disable mandatory filters

To disable mandatory filters (but not remove the specified fields), uncheck the **Mandatory Filtering** checkbox. (Click on it if it is enabled to toggle it off.)

Effect of Mandatory Filtering on Report Design

Mandatory filtering comes into play during report design time with regard to selecting filter criteria. (See [“Select Filter Criteria” on page 112](#) under [“Designing Reports” on page 106](#).)

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.

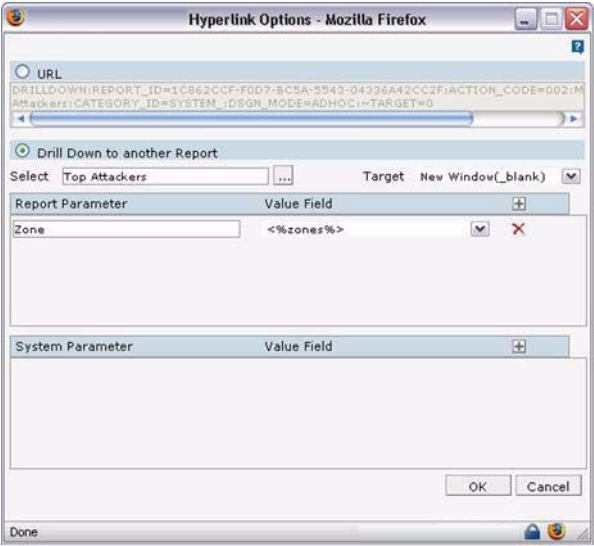


Figure 5-42 Making a Field in a Report a Hyperlink

- 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).
Drill Down to Another Report	<ul style="list-style-type: none">Select Drill Down to another ReportChoose 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 134.

- 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 102)
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

Parameter	Description and Values
Show HTML Toolbar	<p>If the target report is published or viewed in HTML:</p> <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 dialog box titled "Lookup Values". At the top, there is a checked checkbox labeled "Lookup Values". Below it, "Lookup Key Field" is set to "Destination User Name" with a dropdown arrow. 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, there is a section labeled "Values" containing a large, empty rectangular area for listing 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.

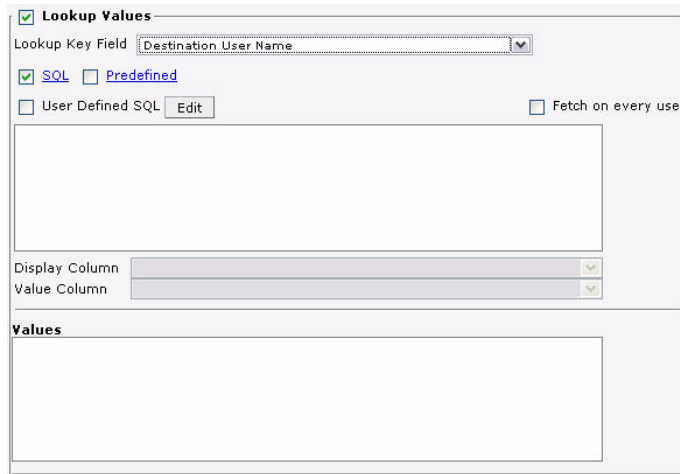


Figure 5-45 Setting SQL Lookup Values in a Query

- 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 124](#).

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 122](#)) and click **Save**.

Deleting a Query Object

You can remove custom queries, but not supplied queries provided with Logger or add-on Solution packs.

To remove a query

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

Defining SQL in the Editor

Each report is built on an SQL query of the Logger databases. SQL (Structured Query Language) is an ISO based standard programming language for retrieving and updating information in a database. ArcSight Logger supports SQL queries, and provides an interactive, SQL Editor in which to define SQL statements.

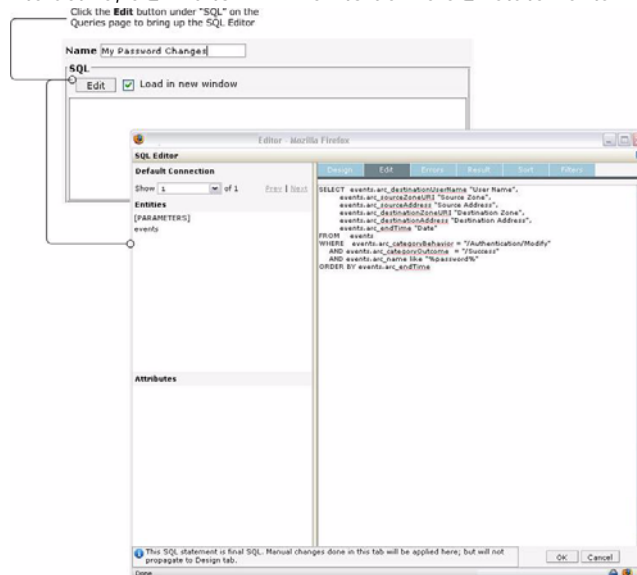


Figure 5-46 Accessing SQL Editor on the Reports | Queries page

Entities and attributes for the selected entity are listed on the left side of the SQL Editor. The right side of the SQL Editor provides tabs showing information related to the selected statement.

Table 5-14 SQL Editor Tabs

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

List of Database Objects

The SQL Editor shows the **Default Connection** to the database that provides the database objects list. ArcSight Logger Reporting provides a single type of object or *entity*, which is an *events* table. When you click on **events** (under Entities), event fields (attributes) are shown under **Attributes**.

Design Tab

You can design simple SQL queries on the **Design** tab using “drag-and-drop”.

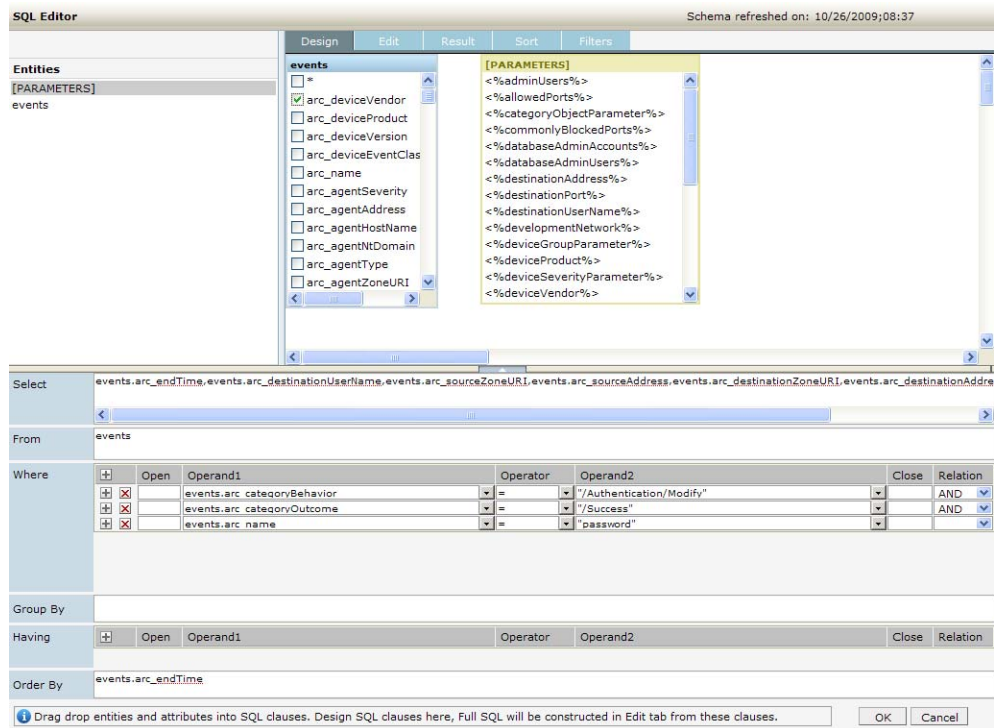


Figure 5-47 SQL Editor: Design Tab

To create an SQL query statement using the Design tab:

- 1 Under **Entities** on the left side of the editor, click **events** to select the “events” entity.

The list of event attributes are shown under **Attributes**.

- 2 Click and drag event attributes from the **Attributes** list on left side of the editor to the **Select** box on the right. The associated values are automatically displayed in the **From** clause.

Repeat these steps to select other attributes from different entities.




The **events** entity must be selected (under Entities on the top left) in order for the event attributes to show up under **Attributes**. If no attributes are displayed, make sure you have “events” selected in the Entities list on the left side of the SQL Editor.

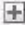

Select

The Select box shows the attributes selected for a given entity.

Where

The Where area shows the “where” clause for the query.

- To get a row at the top, click  (Insert first condition) in the left-most cell of column header.

- To get a row below current row, click  (Add a condition) in the row below which you want to add a row for condition. A row is inserted in the row below the respective row.
- To remove a condition, click  (Remove this condition) in the row for the condition you want to remove.
- To specify a where clause, form a condition by selecting Operand1, Operand2 and Operator.
- To join conditions, create two conditions, and select a relation in the right-most column of the first condition (of the two being joined).
- To group conditions, specify opening brace and closing brace in the right row.

Group By

In the Group By clause you can provide grouping criteria for the SQL statement. To place an entity in Group By, click the entity in the Entity List and drag it in the box below Group By.

Having

To build a "Having" clause, use the same settings as described in the "Where" clause. See ["Where" on page 139](#).



Note

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

Order By

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



Caution

An order-by report query that involves millions of events can fail to run and display the following error messages: `"The server is too busy, try again later"`.

Therefore, 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.

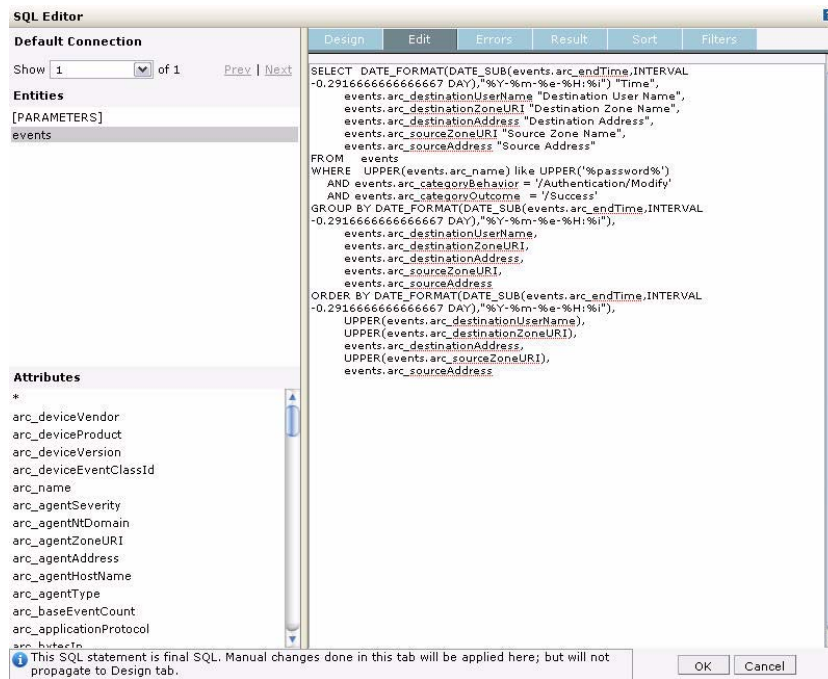


Figure 5-48 SQL Editor: Edit Tab

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.

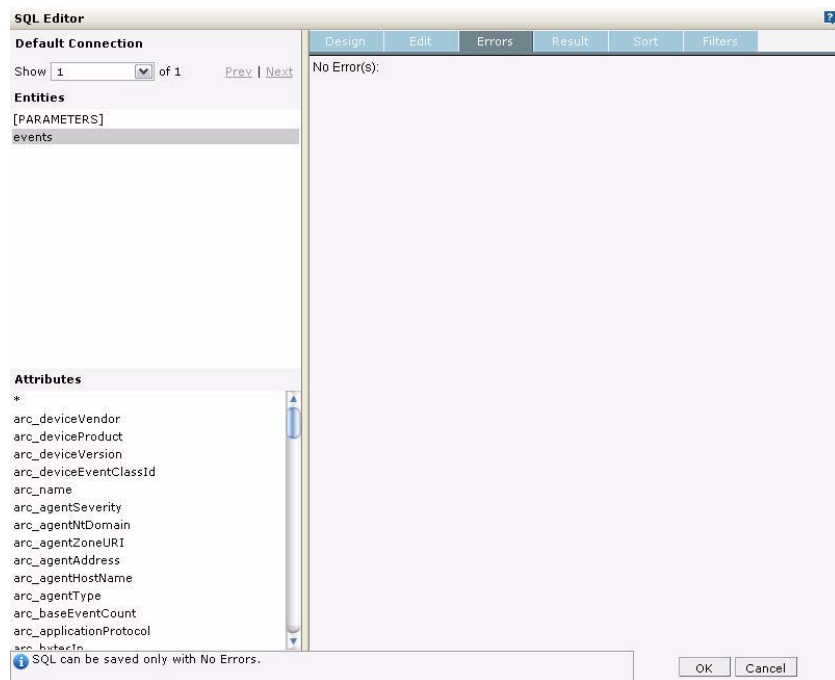


Figure 5-49 SQL Editor: Errors Tab

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.

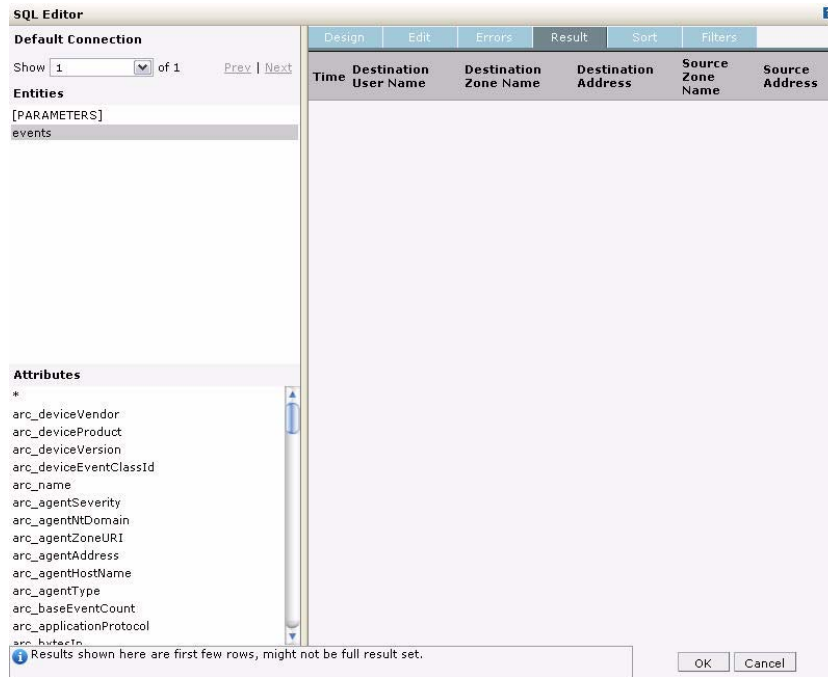


Figure 5-50 SQL Editor: Result Tab

Sort Tab

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

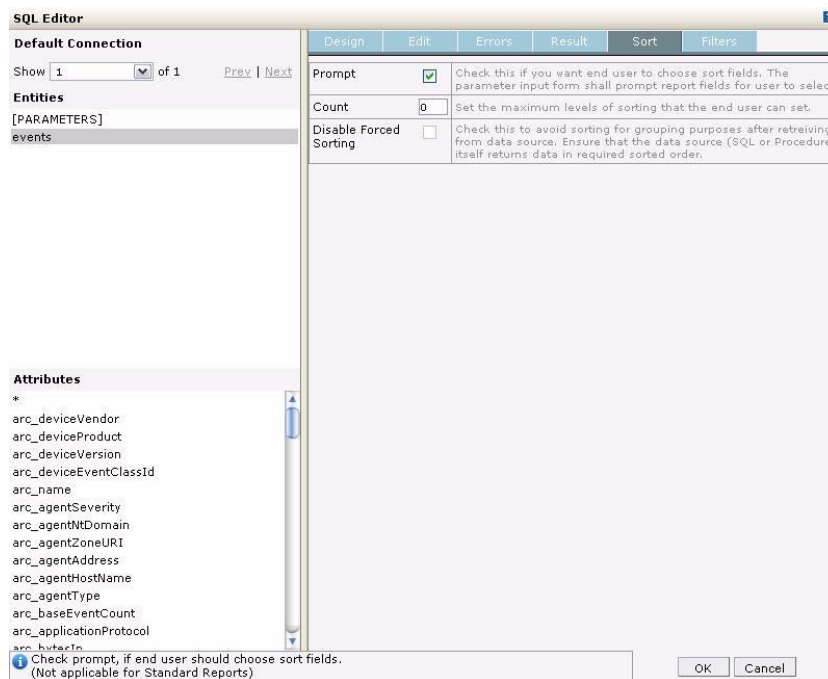


Figure 5-51 SQL Editor: Sort Tab

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.

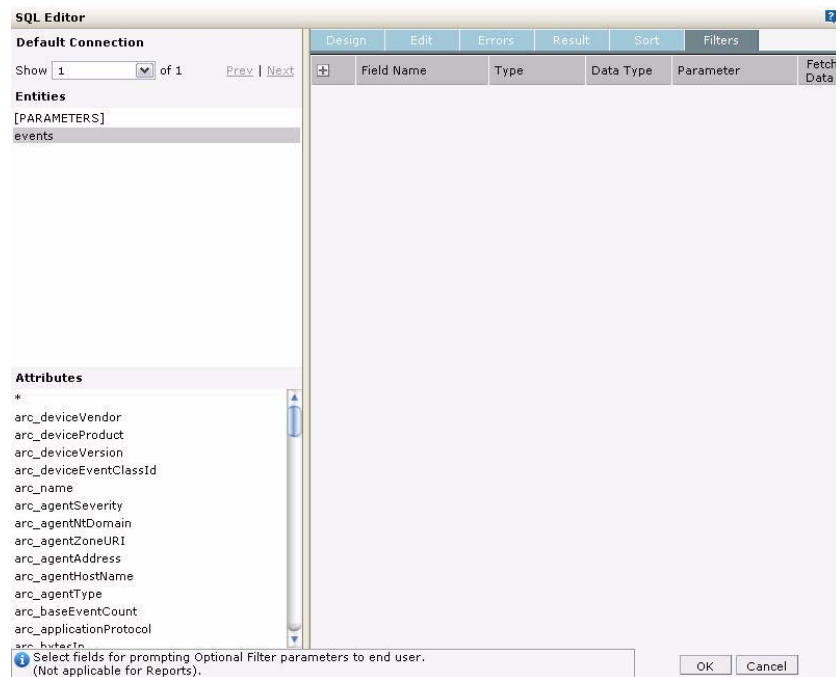




Figure 5-52 SQL Editor: Filters Tab

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 122.](#))

Parameters are stored on server and so can be used in one or more report and query objects.

To view and work with Logger Report parameters, click Design | **Parameters** on the Reports left menu bar.

The screenshot shows the 'Parameters' configuration window. On the left, a list of parameters is shown, with 'categoryObjectParameter' selected. The right pane displays the configuration for this parameter. The 'Name' field is 'categoryObjectParameter', 'Prompt' is 'Resource Type', 'Data Type' is 'CHAR', and 'Size' is '300'. The 'Default Value' is '/Host/Application/D...'. The 'Pre Defined List' is populated with a hierarchy of values: Host, Host/Operating System, Host/Application, Host/Application/Database, Host/Application/Database/Data, Host/Application/Service, Host/Application/Service/Email, Host/Application/Service/Instant, Host/Application/Service/MMS, Host/Application/Service/Peer to Peer, and Host/Application/Service/Phone. The 'Multi Select' checkbox is checked, and the 'Select Default Values' section shows 'Selected' as the chosen option.

Figure 5-53 Report Parameters Object List

Creating New Parameters

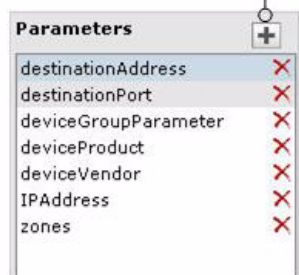


You can search for an existing parameter. To do so, either

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

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

Click to add a parameter



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

The parameter is added to the Parameters list.


Setting Parameter Name, Data Type, and Default Values

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


Figure 5-54 Parameter Name, Data Type, and Default Values

Specify the parameter unique ID, display name, data type, size, format, and default value as described in the table below.

Table 5-16 Parameter Name, Data Type, and Default Values

Option	Description
Name	Provide a name to uniquely identify this parameter.
Prompt	Parameter name displayed on-screen to the user at report run time.
Data Type	Specify type of value the user must provide at report run time: <ul style="list-style-type: none"> CHAR - Value may include alphabetical characters, numbers and special characters. NUMBER - Value may include digits and decimal points DATE - A date or part of a date, like day, month, or year BOOLEAN (For more information, see "To set up a BOOLEAN parameter:" on page 149.)
Size	Specify number of characters or digits this parameter should accept. <p>Note: This is only applicable to CHAR and NUMBER data types, not for Boolean or Date type parameters.</p>
Format	Select the appropriate format in which user should provide value for this parameter. Click  to open a Data Format dialog box. Based on the format you have selected, a format string will appear in the entry box.
Default Value	Specify a default value that is appropriate in most cases to provide for this parameter at report run time. <p>The default value will be automatically selected at report run time. The user can change the default value, if needed. If the user does not change it, the report will run using the default value you specify here for this parameter.</p>

Default Value for Date Type Parameter

For a date type parameter, the **Default Value** field provides an drop-down menu and a calendar. Click the calendar  to provide an explicit date, or select one of these dynamic variable values from the drop-down menu:

- CURRENT_DATE
- MONTH_START_DATE
- YEAR_START_DATE

You can also set a default date that is *relative* to any of the above three dynamic variable dates.

For example, to set a default date as 3 days after CURRENT_DATE, specify CURRENT_DATE + 3.

To set a default date as 5 days before MONTH_START_DATE, specify MONTH_START_DATE - 5 .

To provide a value that is relative to a dynamic variable, select one of the dynamic variables, then type a suffix to it in the Default Value field by adding + or - and the number.

Default Value CURRENT DATE + 3 ±Days

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

Defining Input Type

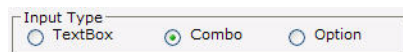
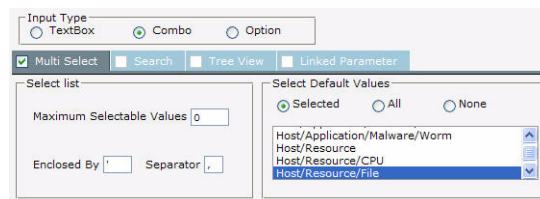


Figure 5-55 Parameter Input Type

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

Table 5-17 Input Type

Option	Description
Text Box	Select "Text Box" input type if you want the user to type the value for the parameter.
Combo	Select "Combo" if you want the user to select one value or multiple values from a drop-down menu. Select the Multiselect checkbox so that user can select multiple values from the box.



See ["Setting Multiple Default Values" on page 151](#) to configure other settings for this option.

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

Setting up Boolean Parameters

Parameters that have a Boolean “Data Type” are represented to the user as checkboxes (the input type) and have only two states:

- Checked (chosen at run time)
- Unchecked (de-selected at run time)

To set up a BOOLEAN parameter:

- 1 Select **Data Type** as BOOLEAN.
- 2 In the **Values** area, for **Checked** specify the value to be passed when the user chooses this option at run time (selects/checks the checkbox presented).
- 3 In **Unchecked** specify value to be passed when the user does not choose this option at run time (de-selects/leaves the checkbox unchecked).

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

Setting Various Run Time Behaviors

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

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

Table 5-18 Parameter Options

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

Option	Description
Restrict to List	<p>This setting is applicable for parameters with Input Type of Combo. Select (check) the Restrict to List checkbox here to force user input of a parameter value from the available run-time options only.</p> <p>If Restrict to List is <i>not selected</i> in the parameter definition you create here, the user can specify a value or can select value(s) from available options.</p>
Pass Values Using Tables	This setting is applicable for Multi Select . Select this checkbox when you want to pass parameter values through a table. This is done especially when the number of values that can be passed (total number of bytes of selected values) as part of the SQL is more than allowed.
Forced	Select this checkbox if you want to restrict parameter values to a pre-specified list of values.

Setting the Data Source List

Specify values for Checkbox, Combo and Option input type. Values can be predefined only.

To Set Predefined Values

Pre Defined List

Display Name	Value
Host	/Host
Host/Operating System	/Host/Operating System
Host/Application	/Host/Application
Host/Application/Database	/Host/Application/Database
Host/Application/Database/Data	/Host/Application/Database/Data
Host/Application/Service	/Host/Application/Service
Host/Application/Service/Email	/Host/Application/Service/Email
Host/Application/Service/Instant	/Host/Application/Service/Instant
Host/Application/Service/MMS	/Host/Application/Service/MMS
Host/Application/Service/Peer to	/Host/Application/Service/Peer to
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 148](#)), 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 146](#)) and click **Save**.

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

Deleting a Parameter

- 1 On the **Reports** right panel menu, click **Parameters** to bring up the Parameters Object List.
- 2 In the **Parameters** list, select the name of parameter that you want to delete.

- 3 Click  (Delete) next to the parameter you want to remove.
- 4 Click the **Save** button to make the change permanent. (If you do not click **Save**, the query object will not be deleted.)

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

Configuring Parameter Value Groups

Some reports may require users to provide multiple run-time values that would be easier to select if they were grouped. For example, a report that requires a user to select more than one country name might be a good candidate for parameter value groups. Users might find it difficult to select a few country names from a single, long list of countries.

As an alternative, the query designer could create parameter value groups for the Americas, Europe, Asia, Africa, and so forth; each with lists of countries belonging to those continents or areas. At report run time, when the user selects a group, values belonging to the group are pre-selected. Users do not have to manually select countries in, for example, Europe or Asia, for every report run. Selections are saved from one report run to the next.

Using parameter value groups as a part of your query design strategy can save users time and reduce error at report run time.

To view and work with Logger Report parameter value groups, click Design | **Parameter Value Groups** on the Reports left menu bar.

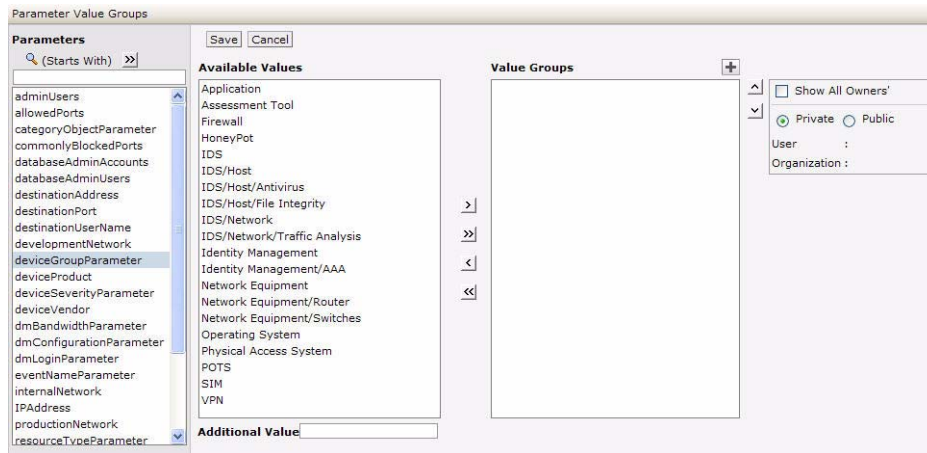



Figure 5-57 Parameter Value Groups



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

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

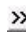
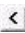
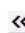
Table 5-19 Parameter Value Groups





Option	Description
Parameters	Lists all the parameter objects.
Available Values	Lists available values for the selected parameter.
Value Groups	Lists groups created and the values selected within a group. An icon appears on the left of a Private group.
Show All Owners	If selected, displays groups created by all users. Such groups will have  icon in the group title.
Option buttons Private and Public	Select Private to list the groups you have set for you only. Select Public if you wish to list the groups you have set for everyone.

To create a group

- 1 Click  (Add Group) next to the **Value Groups** box. A group is created and listed under Value Groups with a default name (based on the currently selected parameter in Parameters list).
- 2 In the Value Groups list, edit the new group name as needed. (Double-click the name to edit it, if it is not already in edit mode.) Double-click the name again to set it, or click outside the box.
- 3 Add the values you want in the group by selecting a value in **Available Values** list and clicking  (Add value to selected group) button. The selected value is added to the selected group in the Value Groups list.
- 4 Repeat [Step 3](#) for each value you want to add to the group.

If a value that you want to add to a group is not listed in Available Values list, specify the value in **Additional Value** field (under Available Values) press Return key. The custom value is added to the currently selected group.

Select an Available Value and click  to add all the values to the selected group in Value Groups, click  to remove the selected value from Value Groups, and click  to remove all the values from Value Groups box.


Select a group and click up  and down  arrows to move the selected group up or down. Select a value and click up  and down  arrows to move the selected value up or down (within the group).

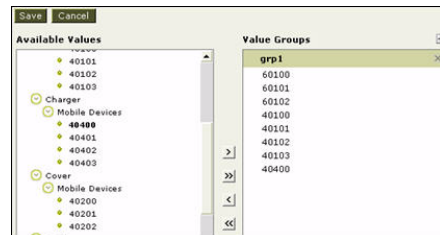
Click **Save** button to save the work.




If the name of a group is changed by a user, the values under that group will be removed from the "Selected Values" group of that user's preferences.


To create a tree view parameter

To select a value, click the leaf node and click  button.



To select all values in a branch (only for a multi-select parameter), click the respective branch and click the  button. All the values under that branch will be selected.

To make changes in name of a group, double-click the group name to make it editable. Specify a new name and click outside the box.

To delete a group, click  in the title of group you want to delete.

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

Applying Report Template Styles

Logger Reports use a style file (.sty) to generate report output per a specified format. The style file defines the look and feel, arrangement, orientation, and so on, of the report output.

You can modify any of the style files from the Logger Reports Template Styles page or you can define a new style to suit your needs.



A report layout file (.irl) defines factors like paper size and static controls. Starting with Logger v4.0, you can define your own layout files. See ["Defining a New Template" on page 155](#) for more information.

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

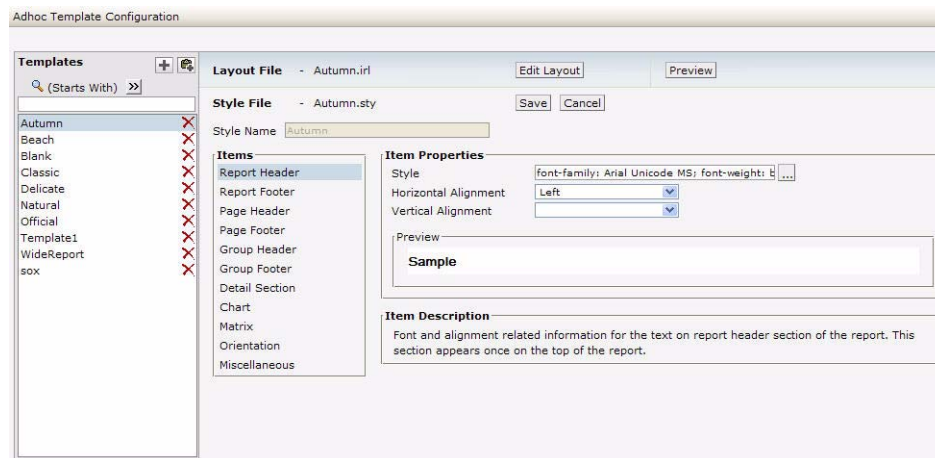


Figure 5-58 Report Template Styles Configuration

Defining a New Template




Note

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**.
- 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 96](#).

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 225](#) 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.

				
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 156](#).

**Note**

The job name is not editable once the scheduled report job is created.

Other settings can be modified with an edit, and work the same way as on the Add a Report Job page described in [“Scheduling a Report” on page 156](#).

To remove a scheduled report

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

**Tip**

Removing the report from Scheduled Reports list here deletes the *scheduled job*, not the report itself nor any instances of its previously published output.

Scheduling a Report

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 225](#) 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 * 

Delivery Operations

Select Delivery Options

☐ Email ☒ Publish


File Name: ☒ Suffix Timestamp Format:

☒ Public ☐ Private

Valid Upto

☒ after generation

☐ End of this

☐ Date: 

Report Parameters

No Parameters

Start: ☒ Dynamic

End: ☒ Dynamic

Device Groups

No Device Groups

Storage Groups

Default Storage Group
Internal Event Storage Group

Devices

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

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

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

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

- ◆ **Email** - For details on setting e-mail delivery options, see [“E-mailing a Report” on page 104.](#)

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 103.](#)

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 158.](#)
- 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 104. • Publish - For details on setting publishing options, see "Publishing Reports" on page 103.
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 102.</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 99.</p>

Deploying a Report Package

You might obtain additional sets of reports from ArcSight to address new security scenarios, add packaged solutions, or enhance your current coverage with updated reports. You can use the Deploy Report Package page to load and deploy packages of new reports onto your Logger system.

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

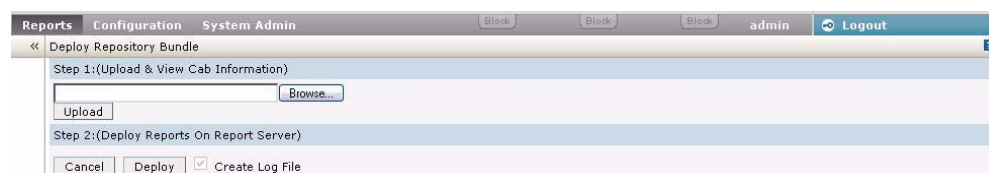


Figure 5-60 Deploy Report Package

A report package (or “cab” file) can contain several types of reporting resources, including:

- Categories and reports
- Organization information
- Schedules
- Portal properties and server properties
- Parameter objects
- Query objects
- Adhoc report templates
- Printer settings
- Database connections

To upload and deploy report package

- 1 In the entry box provided under Step 1, specify the reports package file name and with its full path. Click the **Browse** button to locate the file.
 - 2 Click **Upload**.
- The content is uploaded and information is displayed about the included categories and reports. (A legend is provided below these steps).
- 3 If you want to create log of the deployment process, click (check) the **Create Log File** option.
 - 4 Click **Deploy** to continue with the deployment process.

(Or click Cancel to discontinue with deployment process.)

Status information is displayed about the objects in the package being deployed.

A legend is displayed just below the Deploy button. Information about each of the components in the package is displayed in respective tabs.



Note

Overwrite behaviors are determined when the package was created.

For example, protocol on whether or not an object in the deployed package will overwrite an existing object on the system, and under what circumstances, is determined at package creation time. Therefore, these settings on package deployment are not available to you at deploy time.

A log file will be created if the “Create Log File” checkbox was selected.

The contents of the deployed reports package is available on the respective Logger Reports pages. Solution Reports will be listed under **Solution Reports** on the left panel menu. For more information about these types of reports, see [“Solution Reports” on page 79](#).

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: 50</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 194](#) 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 206](#).

Chapter 6

Configuration

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

In this chapter:

["Devices" on page 163](#)
["Event Archives" on page 166](#)
["Storage" on page 168](#)
["Event Input/Output" on page 173](#)
["Alerts" on page 187](#)
["Scheduled Tasks" on page 193](#)
["Filters" on page 194](#)
["Saved Searches" on page 197](#)
["Search Optimization" on page 200](#)
["Peer Loggers" on page 202](#)
["Configuration Backup and Restore" on page 206](#)
["System Maintenance" on page 208](#)
["Retrieve Logs" on page 215](#)
["Exporting and Importing Content" on page 215](#)

Configuration

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



Note

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

Devices

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

Devices

A Device is a named event source, comprising an IP address (or hostname) and Receiver name. Two Receivers can receive events from the same IP address, so IP address alone is insufficient to identify a Device. Devices can be added to Device Groups, and Device Groups can be referenced in filters and queries. Receivers perform *autodiscovery* by automatically creating a Device for each source IP address. Devices created by autodiscovery are named for their hostname, or if the hostname cannot be determined, their IP address.

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

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

Figure 6-1 Devices page

Maximum number of devices that can be defined on Logger: No limit.

To pre-define a Device

Autodiscovery creates Devices automatically, but you can also pre-define them manually.

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

To edit a Device

One reason for editing a Device is to replace the default name created by autodiscovery (the IP address or hostname) with a more meaningful one.

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

To delete a Device

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

- 2 Locate the Device to be deleted and click the delete icon (✖) on that row. Deleting a Device does not block the source IP address from sending events. If new events are received, autodiscovery recreates the Device.
- 3 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Device.

Device Groups

Device Groups allow you to categorize named source IP addresses called Devices. The Device Groups page, shown in Figure 6-2, lists all Device Groups with edit and delete icons and includes the ability to create new Device Groups.



Define Devices first (or allow autodiscovery to create them), then categorize them into Device Groups, or define empty Device Groups in order to set up Storage Rules and assign Devices to the Device Group later.

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 let you save the events for any day in the past, not including the current day. Archive Storage Settings must be configured before Event Archives can be created. Archive Storage Settings specify the NFS mount, CIFS mount, or SAN to which all event archives will be written. See [“Archive Storage Settings” on page 168](#) for information.

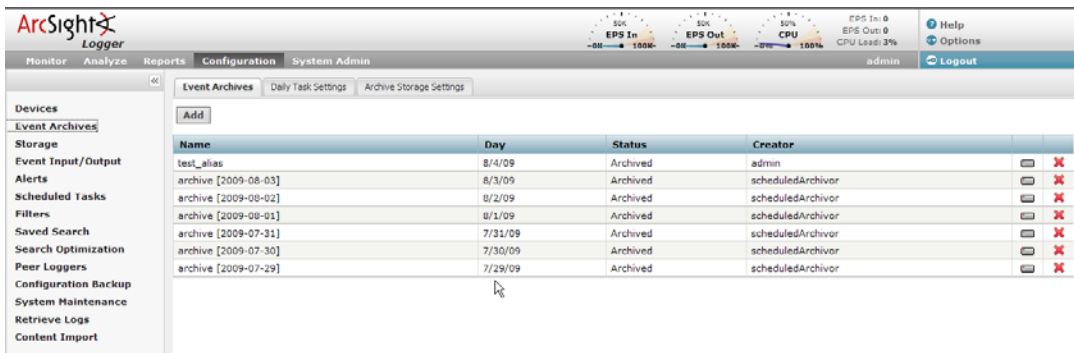
You can also schedule a daily archive of the events. (See [“Scheduled Event Archive” on page 167](#) for information.)

Once events are archived, they are no longer included in search operations. To include those events in search operation, you must load the archive back to the Logger. When an Event Archive is loaded, its events are included in searches, but the archive itself remains on the remote storage. When a query that includes indexed fields is run on archived events, it runs slower than when the data was not archived because the index data on Logger is not archived with events. Therefore, when event archives are loaded, indices are not available.

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

Event Archives

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



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 Event Archives sub-tab is selected. A table displays the current event archives with columns for Name, Day, Status, and Creator. The table lists several archives, all with a status of 'Archived' and created by 'scheduledArchiver'. Each row has a delete icon (X) in the rightmost column.

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

To add an Event Archive

- 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 Name and an Alias for the new Event Archive and specify the day in the form m/dd/yy, where m is month number, dd is the day of the month (with a leading zero if necessary), and yy is the two-digit year number.

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

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

To load or unload an Event Archive

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Archives** in the left panel.
- 3 Locate the Event Archive to be loaded or unloaded and click the edit icon (🔧) on that row. When an Event Archive is loaded, its events are included in searches, but the archive itself remains on the remote storage. When an Event Archive is unloaded, it is available for loading, but its events are not included in searches.



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

To delete an Event Archive

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Archives** in the left panel.
- 3 Locate the Event Archive to be deleted and click the delete icon (✖) on that row.
- 4 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Event Archive.

Scheduled Event Archive

You can schedule a daily event archive and specify what hour of the day it should run. Scheduled event archives appear on the archive list on the Event Archives tab.

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 225](#) before you schedule an event archive.

To schedule a daily event archive

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Archives** in the left panel.
- 3 Click the **Daily Task Settings** tab in the right panel.
- 4 Select a “Time For Daily Archive to Start” from the pulldown list. Scheduled archives must start on the hour, and midnight and 1:00 AM are not on the list.
- 5 Click **Save** to schedule daily event archive, or click on another tab or sub-menu to cancel.

Archive Storage Settings

Event Archives are saved to a specific NFS or CIFS mount point, or SAN.



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

To perform one-time Archive Storage Setting setup

- 1 Create the NFS or CIFS mount point. (See [“Storage” on page 231](#) and [“CIFS Settings” on page 231](#).)
- 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 Select the NFS or CIFS mount point, or SAN for Event Archive output and enter a file path. Click **Save**.



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

Storage

Different storage groups support the implementation of multiple retention policies. Each group can have a different policy, and storage rules determine which storage group is used for events from specific Device Groups. See [“Devices participate in Retention Policy” on page 20](#). 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 21](#)) even if you do not need all of them because **you cannot add storage groups after the Logger has been initialized**, although you can decrease or increase the size of a Storage Group later.

Storage Groups

Storage Groups support multiple retention policies by defining a maximum size (Maximum Size) and number of days (Maximum Age) to retain events. Once events are older than the specified Maximum Age or there are more events than the storage group will hold (as specified by Maximum Size), the oldest events are deleted at the next retention cycle. The retention process triggers periodically on Logger, therefore, events might not be deleted immediately when events gets older than Maximum Age or the storage group size exceeds the Maximum Size limits.

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



Once a Storage Group is created, it cannot be deleted however its size can be increased or decreased any time. **Storage Groups can only be created during the Logger initialization phase**, described in [“Initialization Sequence” on page 20](#). (See [“To edit \(including resizing\) a Storage Group” on page 169](#) 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 admin Help Options Logout

Storage Groups Storage Rules Storage Volume

Edit Storage Group

Important: Please be aware that if you reduce the maximum age from the current value you will lose all events that are older than the new maximum age.

Reducing the maximum storage group age may take a few minutes to complete.

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

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

Name: SGC

Maximum Age (Days): 30

Maximum Size (GB): 50

Current Size (GB): 24

Save Cancel

If you choose to delete data to reduce the storage group size, follow these steps:

- a Set the Maximum Age value to the number indicated in the above message. Doing so, triggers deletion of events.
 - b Refresh the Edit Storage Group screen. When the “Current Size” value is less than or equal to the storage group size you want to set, go to the next step. Otherwise, keep refreshing the screen periodically.
- Note:** The “Current Size” value changes as data is deleted, which can take some time. Therefore, you need to wait before proceeding to the next step.
- c Set the Maximum Size value to suit your needs.
 - d If you wish, restore the Maximum Age setting (that you changed in Step b) to the original value.

If you choose **not** to delete data, go to the next step to exit the procedure.



If there is sufficient space to reduce the storage group size, you can change it without modifying the Maximum Age value (to modify the retention policy to delete data).

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

Storage Rules

Storage Rules create a mapping between Device Groups and Storage Groups. This relationship implements retention policy based on the source of incoming events. All events from firewall devices, for example, can be subject to a short retention period. To accomplish this, manually assign the firewall devices to a single Device Group and then

create a Storage Rule that maps that Device Group to the 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


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

Parameter	Description
Storage Group	Select a Storage Group from the drop-down list. The Storage Groups must already be set up before any Storage Rules are added.
Device Groups	Select one or more Device Groups to associate with the specified Storage Group. You may associate several Device Groups with a single Storage Group.
Priority	An integer that indicates the new rule's priority. The number must be unique for each Storage Rule. The smaller the number, the higher the rule's priority.

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

Figure 6-4 Storage Rules page

To edit or reorder a Storage Rule

- 1 Click the **Configuration > Settings** tab, click **Storage** in the sub-menu, and click the **Storage Rules** tab.
- 2 Find the Storage Rule to be edited in the table.
- 3 Click the Edit icon (). Change the information in the form--for example, change the priority value to reposition the Storage Rule in the table--and click **Save**.

To delete a Storage Rule

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

Storage Volume

Storage Volume settings allows you to specify where Logger will store events. Logger can store events locally (on the storage provided with Logger), on a Network File System, or a Storage Area Network (SAN). This decision must be made when the Logger is first initialized. For more information, see [“Initialization Sequence” on page 20](#).



Note

Storage volume cannot be extended after initialization.

To specify storage volume settings

- 1 Click the **Configuration > Settings** tab, then click **Storage** in the sub-menu. Click the **Storage Volume** tab.
- 2 Enter the following values:

Parameter	Description
Mount Location	Choose Local (to store events on Logger) or any remote file system mounts. (To establish remote file system mounts, see “Storage” on page 231 .)
File Path	If Mount Location is not Local, specify the root folder on the remote file system in which to store event data.
Maximum Size	Enter a maximum event data 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.

If the storage has already been configured, the screen will be read-only.

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:

Path:

Maximum Size (GB):

Preallocation Amount (%):

Save

Figure 6-5 Storage Volume page

Event Input/Output

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

Receivers

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

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

Creating a receiver is a three-step process. First, create a named receiver of a certain type. Receiver type cannot be changed after the receiver is created. New receivers are initially disabled. Second, add type-specific parameters. Receiver parameters are documented in [Table 6-1, "Receiver Parameters," on page 175](#). Third, 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 6-6 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



Caution

Before creating a Receiver of type File Receiver, set up a Network File System mount. See [“Storage” on page 231](#).



Note

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 24](#) and specify:


Logger IP or hostname

Port 443 (port must be 443)


Receiver name

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

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input/Output** (left panel) > **Receivers** tab (right panel).

- 3 Click **Add**.
- 4 Enter a name for the new receiver and choose UDP, TCP, File Transfer, File Receiver, or SmartMessage type.
- 5 Click **Next** to edit receiver parameters listed in Table 4-1.
- 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 221.</p>

Parameter	Receiver Types	Description
Character Encoding	All except File Transfer	Select a character encoding, such as US-ASCII, Big5, or EUC-KR, from the pulldown list. CEF UDP, CEF TCP, and SmartMessage receivers must use US-ASCII or UTF-8 encoding.
Port	UDP, TCP, File Transfer	The default port is 514. (For SmartMessage receivers, configure the SmartConnector for port 443.)
User	File Transfer	A user on the host with privileges to view and read the source log files. If the protocol is FTP, you can specify the special user, "anonymous."
Password	File Transfer	The password of the specified User. The password must not be empty, even in the case of anonymous FTP (although in this case, the password will be ignored.)
FilePath	File Transfer	<p>The path to the log file(s) to be read. Use wild cards like ? and * (for example, "*.log" or "Log-??.txt". Separate directories with forward slashes ('/').</p> <p>Separate multiple file specifications with commas.</p> <p>Example: /security/logs/*/, /security/log?/admin/special/</p>
Schedule	File Transfer	<p>If no schedule is specified, the File Transfer will occur just once.</p> <p>Choose Everyday or Days of Week from the first pulldown menu.</p> <p>If Everyday, select Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the text box.</p> <p>If Days of Week, enter the days (day 1 is Sunday) in the text box. Then choose Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the second text box.</p> <p>For example, to read log files every day at 2 a.m., select Everyday in the first pulldown menu, then choose Hour of Day from the second pulldown menu and enter 2 in the text box. To read the log files every day at 2 a.m. and 3 p.m., enter 2,15 in the text box.</p> <p>For another example, to read log files Tuesdays and Thursdays at 10 p.m., select Days of Week from the first pulldown menu and enter 3,5 for days. Then choose Hour of Day from the second pulldown menu and enter 22 in the text box.</p> <p>Make sure you are familiar with the information in "Impact of Daylight Savings Time Change on Logger Operations" on page 225 before you schedule a file transfer.</p>
Zip Format	File Transfer	Choose gzip, zip, or none.
RFS Names	File Receiver	<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 231. To mount CIFS shares, see "CIFS Settings" on page 231.</p> <p>For more information about SAN, see "SAN" on page 235.</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.
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 format	File Receiver or File Transfer	Required if the log file contains timestamps in the same format for each event. If not specified (or if the Date/time location regex is blank), each event in the file will be stamped with the date that the file itself was first seen by Logger (not its file system timestamp). See Step Table 6-3 for a list of format specifiers. The default is <code>''</code> (no timestamp in log file).
Date/time zone	File Receiver or Transfer	Required if the timestamp in the log file does not specify a time zone. This parameter is ignored if either Date/time format or Date/time location regex are blank. The default is the time zone configured on the Logger (System Admin > Settings > Platform > Time/NTP).

Parameter	Receiver Types	Description
Date/time location regular expression	File Receiver or 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 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 178](#).

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 178](#) show how date and time patterns are interpreted in the U.S. locale. The given date and time are 2001-07-04 12:08:56 local time in the U.S. Pacific Time time zone.

Table 6-2 Date/time examples

Source	Date and Time Pattern
2001.07.04 AD at 12:08:56 PDT	yyyy.MM.dd G 'at' HH:mm:ss z
Wed, Jul 4, '01	EEE, MMM d, ''yy
12:08 PM	h:mm a
12 o'clock PM, Pacific Daylight Time	hh 'o'clock' a, zzzz
0:08 PM, PDT	K:mm a, z
02001.July.04 AD 12:08 PM	yyyyy.MMMMM.dd GGG hh:mm aaa
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

Symbol	Meaning	Presentation	Examples
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: Depends on your hardware platform. Contact your ArcSight Sales Engineer, Professional Services Engineer, or Customer Support for details.



Tip

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

Figure 6-7 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 6-7](#) 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 180](#). Click **Save**.
- 7 New forwarders are initially disabled. Click the disabled icon (🚫) to enable the new forwarder.


Table 6-4 Forwarder Parameters

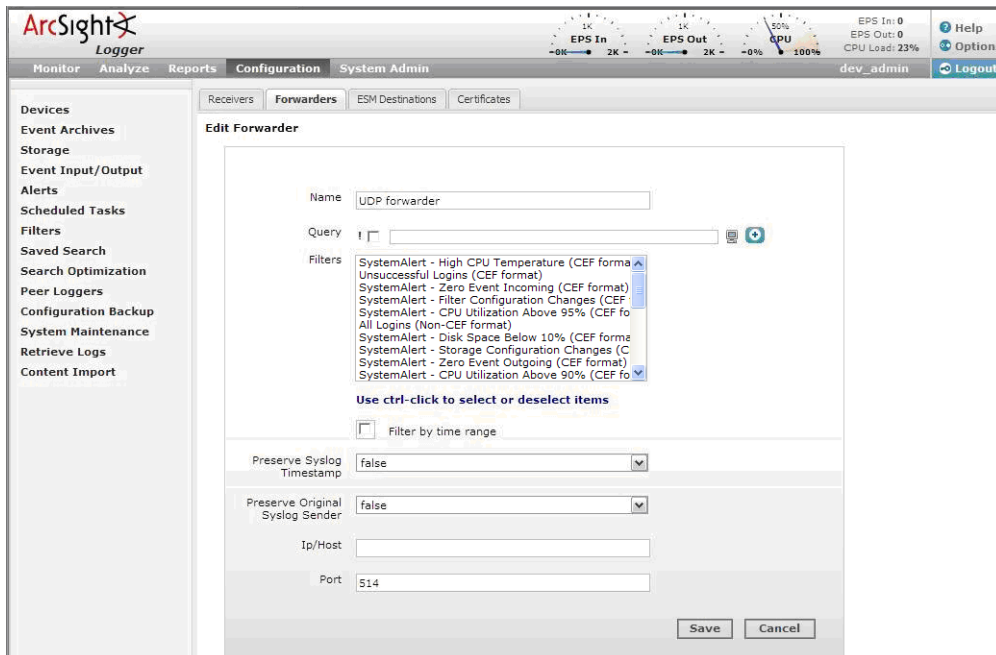
Parameter	Forwarder Types	Description
Name	All	The name of the Forwarder, used for reporting and status monitoring.

Parameter	Forwarder Types	Description
Type		<p>The Type of a Forwarder cannot be changed after the forwarder is created.</p> <p>UDP</p> <p>TCP</p> <p>Connector Forwarder</p> <p>ArcSight ESM (CEF)</p>
Query Terms	All	Specify the events to be forwarded. See “Searching for Events on Logger” on page 59 . Forwarder queries can be constrained by Device Groups and Storage Groups, but not by Peers. See Figure 6-8 .
Filter	All	A filter that specifies which events to forward. (See “Filters” on page 194 .) ESM forwarders always filter out non-CEF events.
Filter by time range	All	<p>Check this box to specify a time range of events to be sent by the forwarder. When this box is checked, Start and End date and time fields appear.</p> <p>Start must be earlier than End. Specifying a time in the future changes that field to the current time. For example, specifying a Start of the current day at 7 a.m. and an End of current day at 7 p.m. will produce events with timestamps from 7 a.m. to the time the filter is saved (that is, earlier than 7 p.m.).</p> <p>Once a forwarder has forwarded all events within a time range, it will no longer forward events. The <code>\$now</code> value, if specified in a time range, is not treated as a variable. Instead, the time when the forwarder was created or updated is assigned to <code>\$now</code>. For example, if the time when forwarder was created was 1:45 p.m. and the time range specified in the forwarder is 10 a.m. to <code>\$now</code>, all events between 10 a.m. and 1:45 p.m. will be forwarded. After events within that time range have been forwarded, no additional events will be forwarded.</p>
Source Type	Connector	<p>Select from the pulldown list of log file types, including:</p> <p>Apache HTTP Server Access</p> <p>Apache HTTP Server Error</p> <p>Juniper Steel-Belted Radius</p> <p>Microsoft DHCP Log</p> <p>IBM DB2 Audit</p> <p>Note: Source Type must be the same in Receiver, Forwarder, and SmartConnector. See “Forwarding Log File Events to ESM” on page 186.</p>
Preserve Syslog Timestamp	UDP, TCP, ESM	Set to true to preserve the syslog timestamp. The default is false--the timestamp is the Logger receipt time.
Preserve Original Syslog Sender	UDP, TCP, ESM	Set to true to preserve the original sender. The default is false--the sender is this Logger.
IP/Host	UDP, TCP, Connector	The destination to receive forwarded events

Parameter	Forwarder Types	Description
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 183.)

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 6-8](#) is displayed.




The screenshot shows the 'Edit Forwarder' configuration window in the ArcSight Logger interface. The window has a sidebar on the left with navigation options like Devices, Event Archives, Storage, Event Input/Output, Alerts, Scheduled Tasks, Filters, Saved Search, Search Optimization, Peer Loggers, Configuration Backup, System Maintenance, Retrieve Logs, and Content Import. The main area is titled 'Edit Forwarder' and contains several sections:

- Name:** A text field containing 'UDP forwarder'.
- Query:** A text field with a search icon and a dropdown arrow.
- Filters:** A list box containing several system alerts, such as 'SystemAlert - High CPU Temperature (CEF format)', 'SystemAlert - Unsuccessful Logins (CEF format)', 'SystemAlert - Zero Event Incoming (CEF format)', 'SystemAlert - Filter Configuration Changes (CEF format)', 'SystemAlert - CPU Utilization Above 95% (CEF format)', 'All Logins (Non-CEF format)', 'SystemAlert - Disk Space Below 10% (CEF format)', 'SystemAlert - Storage Configuration Changes (CEF format)', 'SystemAlert - Zero Event Outgoing (CEF format)', and 'SystemAlert - CPU Utilization Above 90% (CEF format)'. Below the list is a note: 'Use ctrl-click to select or deselect items'.
- Filter by time range:** A checkbox that is currently unchecked.
- Preserve Syslog Timestamp:** A dropdown menu set to 'false'.
- Preserve Original Syslog Sender:** A dropdown menu set to 'false'.
- Ip/Host:** A text field.
- Port:** A text field containing '514'.
- Buttons:** 'Save' and 'Cancel' buttons at the bottom right.


Figure 6-8 Specifying Query Terms, Filters, and other Forwarder parameters.

- 4 Edit the information in the form, as described in [Table 6-4 on page 180](#), 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 259](#)) that are not normalized or categorized.

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

To setup Logger to forward events to an ArcSight ESM Manager

- 1 Copy the server SSL certificate file from an ESM Console or other component that is already communicating with the target Manager, and upload the certificate file to Logger, as described ["Uploading a Certificate to the Logger" on page 185](#).

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 184](#).

- 3 Create an ESM Forwarder that refers to this ESM Destination. (See [“Forwarders” on page 179](#)).

The screenshot shows the 'Add ESM Destination' form with the following fields and values:

- ESM Destination Name: n111-h248
- Connector Name: n111-h248
- Connector Location: /All Connectors/Devices
- Logger Location: QA Lab
- IP/Host: n111-h248
- Port: 8443
- User Name: admin
- Password: [masked]

Buttons: Save, Cancel

Figure 6-9 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 185](#) before adding it as a destination on the Logger. If the certificate file does not exist on the Logger, you will not be able to create an ESM destination.

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

Click **Alerts** in the left panel if you are creating an ESM destination for forwarding Alerts.



Note

The ESM Destinations tab located under Event Input/Output and Alerts in the left panel is the same and contains all ESM destinations configured on a Logger. The tab is accessible from two UI locations to ease configuration.

- 3 In the **ESM Destinations** tab, click **Add**. The page shown in [Figure 6-9](#) is displayed.
- 4 Enter the following parameters:

Parameter	Description
Destination 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 appliance. If you do not want to specify a location, enter “None.”
IP or Host	The ESM Manager to which the Forwarder will direct events.

Parameter	Description
Port	Typically 8443.
Login	The name of an existing User of the ESM Manager with administrator privileges.
Password	The password for the Login user.

- 5 Click **Save**.

To delete an ESM Destination

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

Click **Alerts** in the left panel if you are deleting an ESM destination for forwarding Alerts.



The ESM Destinations tab located under Event Input/Output and Alerts in the left panel is the same and contains all ESM destinations configured on a Logger. The tab is accessible from two UI locations to ease configuration.

- 3 In the **ESM Destinations** tab, locate the ESM Destination to be deleted and click the delete icon (✖) on that row.
- 4 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the ESM Destination.

Uploading a Certificate to the Logger

You need to upload a valid server SSL certificate file for the ESM Manager that you are establishing as a Logger destination.

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

- From the Manager's keystore
- From the ESM Console's truststore
- From the truststore of one of the SmartConnectors that communicates with the Manager

Use the `keytoolgui` utility to export a Manager's certificate as described in the "Using Keytoolgui to Export Certificate" procedure in the *ArcSight ESM Administrator's Guide*. For detailed information about keystore, truststore, their locations on the Manager, ESM Console, and the SmartConnectors, see the *ArcSight ESM Administrator's Guide*.

Once you have exported a certificate for your Manager, copy it to the machine from which you connect to your Logger.

If your Manager has FIPS 140-2 mode enabled, run this command to export the Manager's certificate from the Manager's `<ARCSIGHT_HOME>/bin` directory:

```
arcsight runcertutil -L -n managerkey -r -d
<ARCSIGHT_HOME>/config/jetty/nssdb -o
<absolute_path_to_manager.cert>
```

This command generates the `manager.cert` file, the Manager's certificate, in the location that you specified in the above command.



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.



The ESM Destinations tab located under Event Input/Output and Alerts in the left panel is the same and contains all ESM destinations configured on a Logger. The tab is accessible from two UI locations to ease configuration.

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

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

Forwarding Log File Events to ESM

Logger can read events from a log file and forward those events to an ArcSight Logger Streaming SmartConnector that sends the events on to ArcSight ESM. Unlike other events that Logger receives, such as syslog, SmartMessage, or CEF, log file events must be parsed to determine event timestamp.

To forward log file events to ESM, configure the Receiver, Forwarder, and SmartConnector to accept the same Source Type (as described in [“Receiver Parameters”](#) on page 175).

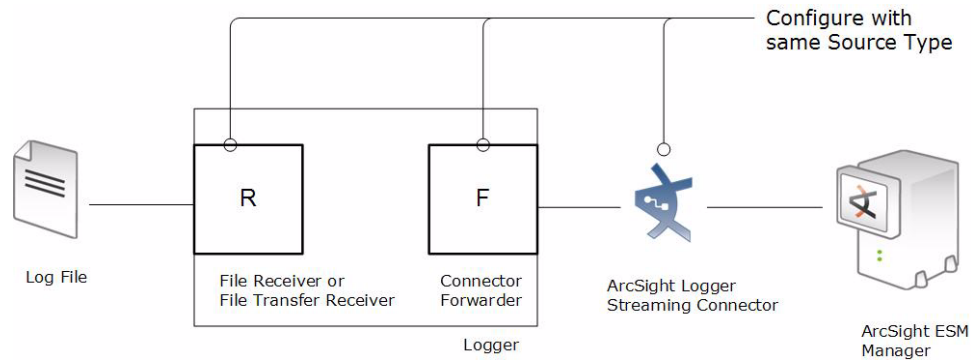


Figure 6-10 Configure the Receiver, Forwarder, and SmartConnector with the same Source Type to use Logger to forward log file events to ArcSight ESM.

Alerts

Alerts respond to internal events or specified event patterns with optional notification. Internal events include storage capacity warnings or, on some hardware models, CPU temperature warnings. Event patterns are specified events that occur above a particular frequency (a threshold number of events in a specified time period).

An alert is triggered if a specified number of matches occur within the specified threshold (time interval in seconds). An alert notification is sent through previously configured destinations—e-mail addresses, SNMP server, Syslog server, and ESM Manager.

When an alert is triggered, an audit event is logged indicating that an alert has fired. Secondly, Logger creates an alert event containing the trigger event. This alert event is then sent to the specified destinations.

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]
```

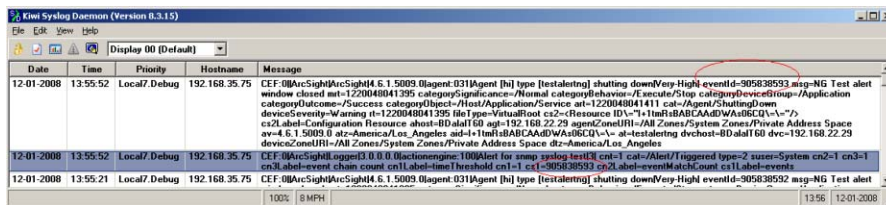
```

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:



- Non-CEF events do not contain event IDs. If you need to associate such base events with their trigger alert, send such events to Logger through a connector.
- All SNMP alerts are sent as SNMP traps; therefore, trigger alerts and their associated matched (base) events are received as SNMP traps on an SNMP destination.
- SNMP uses UDP to send packets. As a result, the order in which alerts arrive at an SNMP destination is not guaranteed. Use the event IDs in the trigger alert to identify its associated base events.

Similarly, when Syslog events are sent using UDP, the order in which the trigger alert and matched events arrive is not guaranteed.

Maximum number of Alerts you can create: no limit

Maximum number of Alert destinations you can set: no limit on e-mail destinations; however, you can only set one SNMP, one Syslog, and one ESM destination.

Maximum number of Alerts you can enable at one time: 5

To view Alerts

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

- Click **Alerts** in the left panel. The Alerts list is displayed, as shown in [Figure 6-11](#).

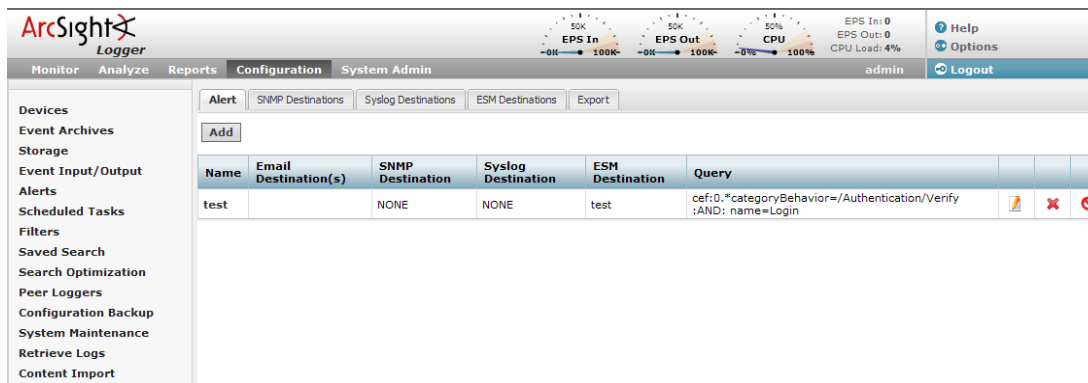


Figure 6-11 Alert list

To add an Alert

To create an Alert, you will need to specify a query or filter, event aggregation values (Match Count and Threshold, described below), and one or more notification destinations. If the new Alert will send notification using an SNMP, Syslog, or ESM destination, set up those destinations before creating the Alert. To configure the e-mail destination, see [“SMTP Settings” on page 226](#). See also [“SNMP Destinations” on page 191](#), [“Syslog Destinations” on page 191](#), and [“ESM Destinations” on page 183](#).

When you create an alert, it is in disabled state. You can enable it using instructions in [“To Enable or Disable an Alert” on page 190](#).

- Click **Configuration** from the top-level menu bar.
- Click **Alerts** in the left panel.
- Click **Add**. The page shown in [Figure 6-12 on page 190](#) is displayed.
- 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 Alerts. For more information about specifying a regular expression query, see [“The Need to Search Events” on page 39](#).

- 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.



To be notified if 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.

- 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**.

Figure 6-12 Add Alert dialog

To Enable or Disable an 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 an 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 6-12 on page 190](#) appears. Remember that only alphanumeric characters can be used in an Alert name.

To Remove an 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.

SNMP Destinations

SNMP Destinations describe how Alert notifications should be sent using Simple Network Management Protocol (SNMP). Set up SNMP Destinations before creating Alerts that will use them.

To Add an SNMP Destination

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **SNMP Destinations** tab in the right panel.
- 4 Click the **Add** button.
- 5 Enter parameters:

Parameter	Description
SNMP Destination Name	A name for this destination.
Connector Name	The SmartConnector name.
Connector Location	The physical location of the SmartConnector machine. If you do not want to specify a location, enter "None".
Logger Location	Optional comment describing Logger's physical location.
SNMP Host	Host name or IP address.
SNMP Port	162, by default.
Community Name	SNMP community name.

- 6 Click **Save** to create the new SNMP Destination.

To Remove an SNMP Destination

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **SNMP Destinations** tab in the right panel.
- 4 Locate the SNMP Destination to be removed and click the remove icon (✖) on that row.
- 5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the SNMP Destination.

Syslog Destinations

Syslog Destinations describe how Alert notifications should be sent using the comparatively simple Syslog protocol. Set up Syslog Destinations before creating Alerts that will use them.

To Add a Syslog Destination

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **Syslog Destinations** tab in the right panel.
- 4 Click the **Add** button.

- 5 Enter parameters:


Parameter	Description
Name	A name for this destination
Type	UDP or TCP Syslog. This choice cannot be edited later.

- 6 Click **Next**. Enter the secondary parameters:


Parameter	Description
Send Syslog Timestamp	True or false (default is false). If false, the syslog message will have the current Logger time.
Send Original Syslog Sender	True or false (default is false). If false, the syslog message will be sent from Logger.
Ip/Host	Host name or IP address
Port	Port (default is 514)

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

To Edit a Syslog Destination

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **Syslog Destinations** tab in the right panel.
- 4 Click the Edit icon (). You can edit the parameters of the Syslog Destination except its type.
- 5 Click **Save** to make the changes, or **Cancel** to return to the Syslog Destination table.

To Remove a Syslog Destination

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **Syslog Destinations** tab in the right panel.
- 4 Locate the Syslog Destination to be removed and click the remove icon () on that row.
- 5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Syslog Destination.

ESM Destinations

ESM Destinations describe how Alert notifications should be sent to an ArcSight ESM Manager. Set up ESM Destinations before creating Alerts that will use them. If an ESM Manager uses a signed SSL certificate, you will need to load it on the Logger.

To setup Logger to send alerts to an ArcSight ESM Manager

- 1 If the ESM Manager uses a certificate, copy the server SSL certificate file from an ESM Console or other component that is already communicating with the target Manager,

and upload the certificate file to Logger, as described [“Uploading a Certificate to the Logger” on page 185](#).

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 184](#).

Export

See [“Exporting and Importing Content” on page 215](#).

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 225](#) 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 6-13](#), 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**.

Scheduled Tasks Currently Running Tasks Finished Tasks					
Filter by Job Type 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 6-13 Scheduled Tasks page

To add a Scheduled Task

Scheduled Tasks can be created for:

- Saved Search (See [“Saved Search Jobs” on page 198](#))
- File Receivers and File Transfer Receivers (See [“Receivers” on page 173](#))
- Event Archive (See [“Event Archives” on page 166](#))
- Configuration Backup (See [“Configuration Backup and Restore” on page 206](#))

To delete a Scheduled Task

- 1 Click the **Configuration > Settings** tab, click **Scheduled Tasks** in the sub-menu, then click the **Scheduled Tasks** tab.
- 2 Locate the Scheduled Task to be deleted and click the delete icon (✖) on that row.
- 3 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Scheduled Task.

Currently Running Tasks

The Currently Running Tasks page displays the Scheduled Tasks that are running at the present time. The table shows task name, type, and the date and time that the task started.

To view Currently Running Tasks

- 1 Click the **Configuration > Settings** tab, click **Scheduled Tasks** in the sub-menu, then click the **Currently Running Tasks** tab.
- 2 Filter the list by selecting a specific type of Scheduled Task from the drop-down list, or select **All**.

Finished Tasks

The Finished Tasks page acts like a log of all Scheduled Task runs, with the most recently finished tasks on top.

To view Finished Tasks

- 1 Click the **Configuration > Settings** tab, click **Scheduled Tasks** in the sub-menu, then click the **Finished Tasks** tab.
- 2 Filter the list by selecting a specific type of Scheduled Task from the drop-down list, or select **All**.

Filters

The Filters section has three tabs: Filters, Search Group Filters, and Export.

Filters

The Filters page provides a convenient place to manage filters. There are two types of filters: shared and Search Group. You can also create and delete shared filters on the Analyze page. Shared filters are optional; they provide a way to focus on events of interest. Search Group filters are not optional—they limit the events that users in a particular user group are able to see as an access control mechanism, those users can not opt to remove the filter. Search Group filters can also be used to limit the events processed by a category of reports (see [“Using Report Category Filters” on page 162](#)).

A set of predefined filters, also known as system filters, exist on your Logger as well. For more information about system filters, see [“System Filters/Predefined Filters” on page 68](#).

To create a filter

- 1 Click the **Configuration > Settings** tab, click **Filters** in the sub-menu, then click the **Filters** tab.
- 2 Click **Add**. The page shown in is displayed.
- 3 Enter a name for the new filter in the Name field.

Filter names are case-sensitive.

- 4 If you are creating a shared filter, select **Unified**.
For Search Group filters, select **Search Group**. Additionally, non-administrator users cannot create Search Group filters.
- 5 Click **Next**.
- 6 If you selected Unified method in the previous step, enter the query for the new filter.
Click Advanced Search to use the Search Builder Tool to create the query. For details about using the Search Builder Tool, see [“Using the Search Builder Tool” on page 53](#).
For instructions on creating a query, see [“Searching for Events on Logger” on page 59](#).
The new filter will eliminate events that do not match the query.
- 7 Click **Save**.



If you created a Search Group filter, make sure that you associate it to a user group, as described in [“Search Group Filters” on page 195](#).

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 162](#).

- **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.


Filters Search Group Filters			
Name	Filter	Description	
Default Logger Search Group	NONE	The default search group allows both local and distributed searches.	

Figure 6-14 Search Group Filters Page




In the Search Group Filters page (shown in [Figure 6-14](#)), 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 194](#). To create, edit, or delete User Groups, see [“User Groups” on page 251](#).



Only users that are members of a System Admin group can assign Search Group Filters. For more information, see [“User Groups” on page 251](#).

To associate a Search Group Filter with a User Group

- 1 If the User Group that you want to associate with the Search Group Filter does not exist, create a new User Group of type Search Group. For instructions, see [“User Groups” on page 251](#).
- 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 194](#). 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 6-14](#) 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 215](#).

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.

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 225](#) before adding a Saved Search.

Saved Search

The Saved Search tab displays all Saved Searches and supports Adding, Editing, and Deleting Saved Searches.

To add a Saved Search

- 1 Click the **Configuration > Saved Search**.
- 2 Click **Add** and enter the following parameters:

Parameter	Description
Name	A name for this Saved Search. This name will be used for exported output files, with the Saved Search date and time appended.
Start Time	Absolute date and time of earliest possible event. Or check Dynamic to specify the start time relative to the time when the Saved Search job is run.
End Time	Absolute or Dynamic date and time of latest possible event, as described above.
Query Terms	Enter the query in the text field or select one or more Filters from the list below the text field.
Local Search	Check this box to limit the Saved Search to the local Logger box. If the Local Search box is left unchecked, the Saved Search will include all Peer Loggers as well as the local Logger.

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

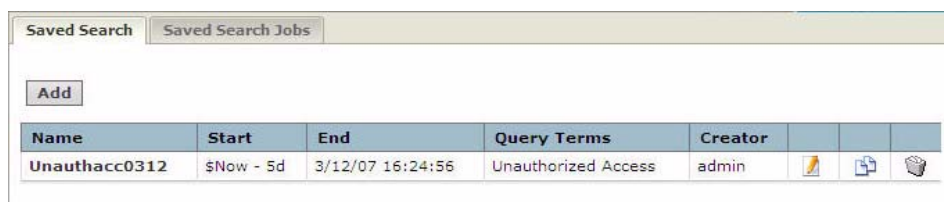




Figure 6-15 Saved Search page

To edit a Saved Search

- 1 Click the **Configuration > Saved Search**.
- 2 Find the Saved Search to be edited in the table.

- 3 Click the Edit icon (). Change the information in the form and click **Save**.

To delete a Saved Search

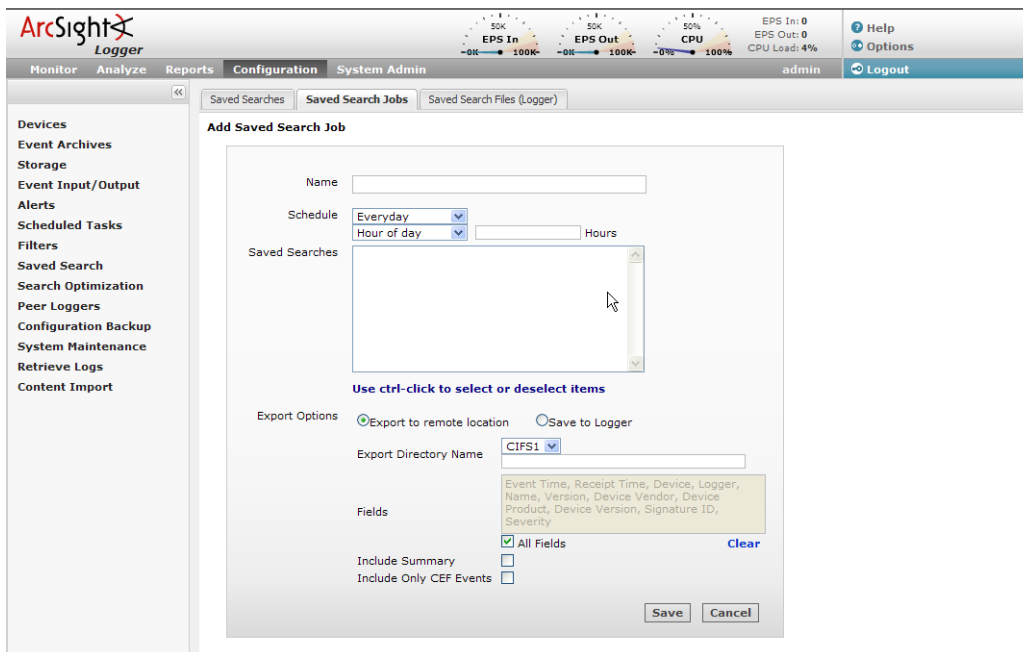
- 1 Click the **Configuration > Saved Search**.
- 2 Find the Saved Search to be deleted in the table.
- 3 Click the Delete icon (). Confirm the delete.

Saved Search Jobs

A Saved Search Job schedules a Saved Search to be run at a later time. Before you can create a Saved Search Job, you must have created or saved at least one Saved Search.

To add a Saved Search Job

- 1 Click the **Configuration > Saved Search > Saved Search Jobs** tab. The screen shown in [Figure 6-16](#) is displayed.



The screenshot shows the 'Add Saved Search Job' page in the ArcSight Logger interface. The page has a sidebar on the left with navigation links: Monitor, Analyze, Reports, Configuration, and System Admin. The main content area is titled 'Add Saved Search Job' and contains a form with the following fields and options:

- 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'.
- Saved Searches:** A list box showing available saved searches.
- Export Options:** Radio buttons for 'Export to remote location' (selected) and 'Save to Logger'.
- Export Directory Name:** A dropdown menu set to 'CIFS1'.
- Fields:** A list of fields including Event Time, Receipt Time, Device, Logger, Name, Version, Device Vendor, Device Product, Device Version, Signature ID, and Severity. The 'All Fields' checkbox is checked.
- Include Summary:** A checkbox.
- Include Only CEF Events:** A checkbox.
- Buttons:** 'Save' and 'Cancel' buttons at the bottom right.

Figure 6-16 Saved Search Jobs page

- 2 Click **Add** and enter the following parameters:


Parameter	Description
Name	A name for this Saved Search Job.

Parameter	Description
Schedule	<p>Chose 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> <p>Make sure you are familiar with the information in “Impact of Daylight Savings Time Change on Logger Operations” on page 225 before adding a Saved Search.</p>
Saved Searches	<p>Select one or more Saved Searches to perform on the specified schedule. Click to select a Saved Search. Ctrl+Click selects or un-selects a Saved Search without affecting the other selected Saved Searches. Shift+Click selects a range of Saved Searches.</p> <p>If a single Saved Search is selected, the output will detail all specified events. If more than one Saved Search is selected, a summary report will be produced, with an event count for each Saved Search specified. In the single case, each output row will describe an event; in the multiple case, each output row will describe the results of a Saved Search.</p>
Export Options	<p>Select from the following 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 200) feature to access those files.
Export Directory Name	Folder path for the output file.
Fields	Edit the list of fields desired for output or check the All Fields box. Click the Clear link to clear the text box.
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 259) events in the output.


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

To edit a Saved Search Job

1 Click the **Configuration > Saved Search > Saved Search Jobs** tab.

- 2 Locate the Saved Search Job to be edited and click the edit icon () on that row.
- 3 Change the parameters of the Saved Search Job.
- 4 Click **Save** to update the Saved Search Job, or **Cancel** to abandon your changes.

To delete a Saved Search Job

- 1 Click the **Configuration > Saved Search > Saved Search Jobs** tab.
- 2 Locate the Saved Search Job to be deleted and click the delete icon () on that row.
- 3 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Saved Search Job.

Saved Search Files

Access Saved Search results that were Saved to Logger with the Saved Search Files command. Saved Search Files can be retrieved (streamed to the browser) or deleted.

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

Figure 6-17 Saved Search Files page

Search Optimization

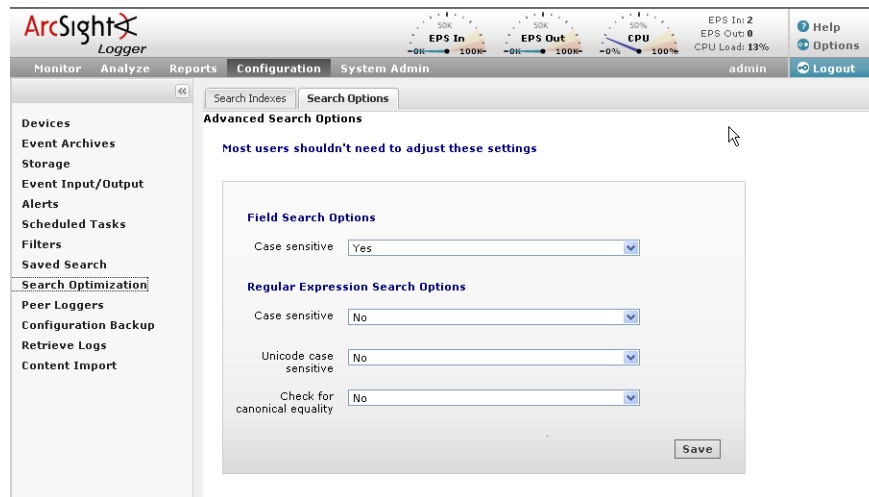
The search optimization option enable you to:

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

Add Search Indexes

See ["Indexing" on page 63](#) for more information.

Tuning Advanced Search Options



The following table lists the advanced search options you can view and configure. These options support i18n choices. If you change any of these options, you will need to reboot your Logger for those changes to take affect.

Option	Description
Case sensitive	<p>Defaults: Yes, for field query; No, for regular expression.</p> <p>Full-text search (keyword search) is case insensitive. You cannot change its case sensitivity.</p> <p>When this option is set to No, searching for "login" will find "login," "Login," and "LOGIN".</p> <p>Note: Case-sensitive search only applies to the local Logger. Peer loggers will continue to use case-insensitive search.</p> <p>Set this option to Yes to increase local query performance.</p>
Unicode case sensitive	<p>Default: No</p> <p>Set to Yes if non-English events should be compared in a case-sensitive way.</p> <p>This option only applies to the regular expression search method.</p> <p>Note: ArcSight strongly recommends that you do not change this option.</p>
Check for canonical equality	<p>Default: No</p> <p>Set to Yes if non-English events should be compared using locale-specific algorithms.</p> <p>This option only applies to the regular expression search method.</p> <p>Note: ArcSight strongly recommends that you do not change this option.</p>

To change any of the above options, click **Configuration > Search Optimization > Search Options** tab (selected by default).

Deleting Custom Field Sets



You need to have the “Edit, save, and remove fieldsets” privilege to delete a custom field set.

Note

To delete a custom field set:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Search Optimization** from the left panel.
- 3 In the Fieldsets tab, identify the field set you want to delete and click the delete (✖) icon.
- 4 Confirm the deletion.

Peer Loggers

A Logger can establish peer relationships with one or more Loggers to enable distributed event searches.

When two Loggers peer with each other, one Logger initiates the relationship. The initiator Logger sends the credentials to authenticate itself to the other Logger, called the remote Logger from hereon. If the authentication succeeds, a peer relationship is established between the two Loggers. (The remote Logger must have the credentials for the initiator Logger configured on it for authentication to succeed.)

Adding a peer on a Logger is a bi-directional process. That is, when Logger A adds peer access for Logger B, Logger B automatically adds peer access for Logger A. Similarly, if you delete the peer access for B on A, the peer access for A is automatically deleted on B.

Peer Loggers can authenticate using any of these methods:



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

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

- User name and password

A user name configured on the Logger is used for authentication
- Peer Authorization ID and Code

Authorization ID and Code generated on a remote Logger are used by the initiator Logger to peer with it. The generated ID and Code are specific to the initiator Logger because the IP address of the initiator is used to generate the ID and code, and can be

used only for peering from the initiator. Therefore, this method is more secure than using user name and password.



If the peer-relationship initiating Logger has version 3.0.x installed and the other Logger is running v4.0 GA with SSL Client Authentication (CAC) enabled, you can still use this authentication method. Enter the generated Authorization ID in the User Name field and the Code in the Password field on the v3.0.x Logger.

Guidelines

You should be aware of these guidelines when peering Loggers:

- You can peer a Logger to one or more remote Loggers.
- Peer Loggers can run different versions. However, if a peer Logger is not running v4.0, make sure it is running at least v3.0.x or later; versions earlier than 3.0.x are not supported for peering.
- Currently, report generation across peer Loggers is not supported.
- If the remote Logger is configured for SSL Client authentication (CAC), you must configure an authorization ID and code on the initiator Logger. If the initiator Logger is v3.0.x, enter the authorization ID and code in the User Name and Password fields.

There are no special authentication requirements for FIPS-enabled Loggers. Such Loggers can use any of the allowed authentication methods.

- Peer loggers cannot be edited, however you can delete and readd a peer.
- A user must belong to the Logger Search User Group with "Search for events on remote peers" privilege set to Yes.
- Users performing search operations on peers have the same privileges on the peer that they have on the Logger they are logged in.

For example, UserA is restricted by a search group filter to only search for events in which deviceVendor is set to "Cisco". When UserA performs a search operation across LoggerA's peers, the same constraint (to search events where deviceVendor = "Cisco") is applied on all peers.

- If user name and password are used for authenticating to a remote peer Logger, the credentials are only used one-time, during the peering relationship set up. After a relationship has been established, the credentials are not saved (on the Peer Loggers page) and the peers do not authenticate periodically. Therefore, if the user name or password used to establish a relationship is changed at a later date or the user name is deleted, peering relationship is not broken. However, if you delete the peering relationship or it breaks for other reasons, you will need to enter the updated credentials to re-establish the relationship.

The following example illustrates the steps you need to follow to set up peering between two Loggers.

Logger A

Logger B

- 1 Select the Logger that will initiate the establishment of the peering relationship. In this example, Logger A will initiate the relationship.
- 2 If Logger B is configured to use user name and password authentication, go to [Step 3](#).
If Logger B is configured to use SSL Client Authentication (CAC), go to [Step 4](#).

Logger A**Logger B**

- 3 Set up a user name and password that Logger A will use to authenticate itself when peering with this Logger, as described in [“Users” on page 256](#).
- 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 205](#).
- 5 Add Logger B's information, as described in [“To add a peer Logger” on page 204](#):

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' form. This form includes fields for 'Peer Host Name', 'Peer Port' (set to 443), 'Peer Login Credentials' (selected with a radio button), 'Peer User Name', 'Peer Password', 'External IP Address' (set to 192.168.36.42), and 'Local Port' (set to 443). A note states: 'Following fields are for local (currently connected) logger and are optional. This needs to be changed only seldomly.' '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.

Parameter	Description
Peer Login Credentials	Select Peer Login Credentials for password-based authentication with the remote Logger.
Peer Authorization Credentials	<p>Select Peer Authorization Credentials for SSL client authentication with the remote Logger. (See “SSL Client Authentication (CAC Authentication)” on page 240.)</p> <p>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.</p>
If you selected Peer Login Credentials...	
Peer User Name	The user name to use when connecting to the remote Logger.
Peer Password	The password for the user on the remote Logger.
If you selected Peer Authorization Credentials...	
Peer Authorization ID	Enter the authorization ID of the other Logger to which this Logger is initiating a peering relationship. (See “To generate Authorization ID and Code for configuring a peer relationship” on page 205 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 205 for more information.)
These fields need to be updated in rare circumstances. For more information, read the description of each field in this table.	
External IP Address	<p>In most cases, the value in this field matches the IP address you use to connect to this Logger from your browser, and you do not need to do anything.</p> <p>However, if the IP address does not match (for example, when the Logger is behind a VPN concentrator), change the value to match the IP address with which you connect to this Logger.</p>
Local Port	Make sure the value of this field is set to 443.

- 4** Click **Save** to add the new Logger, or **Cancel** to quit.


To 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

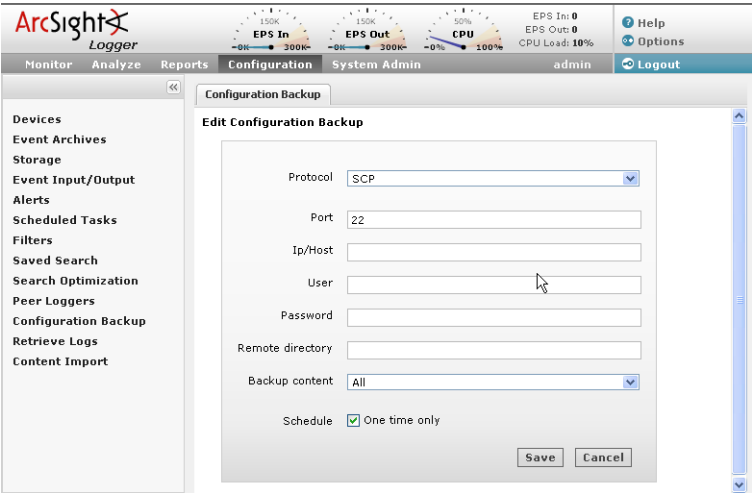
You can use the backed up content to:

- 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)



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 225.</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 hardware model and the version running on it is 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 207.

System Maintenance

Certain operations on Logger, such as database defragmentation, require that Logger be in a maintenance state—a state in which operations related to data on the Logger are not

running. Maintenance mode enables you to place the Logger in such a state. When a Logger is in maintenance mode:

- Events are not processed
- Reports are not generated
- Search cannot run
- Scheduled jobs do not run

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

Admin Rights		
Reboot		
Reboot Appliance		<input checked="" type="radio"/> Yes <input type="radio"/> No
Update		
Update Appliance		<input checked="" type="radio"/> Yes <input type="radio"/> No
Enable Maintenance Mode		<input checked="" type="radio"/> Yes <input type="radio"/> No
System Information		
Process Status		<input checked="" type="radio"/> Yes <input type="radio"/> No
RAID Controller		<input checked="" type="radio"/> Yes <input type="radio"/> No

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

Not Allowed

Another user has placed Logger in maintenance mode.

During this time, only maintenance operations may be performed by that user.

Although it is not recommended, you may [reboot](#) Logger to resume normal operation.

You can [refresh](#) this page or report the problem to your Administrator

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

MAINTENANCE MODE IN EFFECT

Use a valid username and password to reboot the ArcSight Logger appliance.

ArcSight Logger Login

Username

Password

Login

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

Entering Maintenance Mode

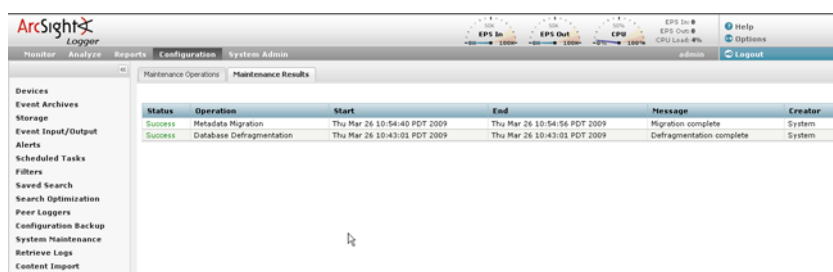
You cannot place a Logger in maintenance mode directly. A Logger can enter maintenance mode only when you perform an operation that requires it to be in that mode. For example, when defragmenting database, the user interface prompts you to enter Logger in maintenance mode, as illustrated in “[Database Defragmentation](#)” on page 210.

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.

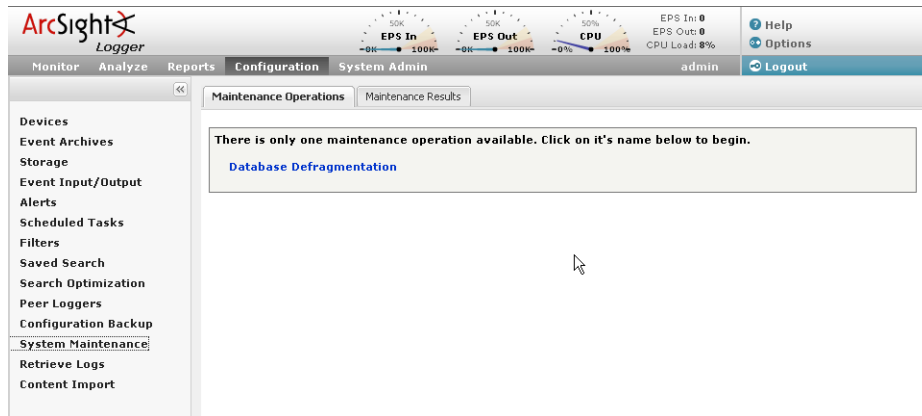
- Long pauses in the receiver and forwarder operations

Starting with this release, you can defragment a Logger that exhibits the above listed symptoms. Make sure that you have read the following guidelines before starting the defragmentation process.

Guidelines for Database Defragmentation

- Ascertain that the Logger symptoms are not due to issues related to network infrastructure such as network latency or unexpected load on the Logger.
- The Logger system needs to be placed in maintenance mode before defragmentation can begin. As a result, most processes on the Logger are stopped—no events are processed or scheduled jobs run, and most user interface operations are unavailable. For more information about maintenance mode, see [“System Maintenance” on page 208](#).
- A minimum amount of free disk space is required on your system to run database defragmentation. The utility automatically checks for the required free space and displays a message if sufficient disk space is not found.
- Although you can defragment as needed, if you are using this utility too often (such as on a system that was defragmented over the last few days), contact ArcSight Customer Support for guidance.
- If the defragmentation process fails at any point, the Logger returns to the same state that it was in before you started defragmentation. You can safely reboot the Logger and restart the process from the beginning.

Defragmenting a Logger



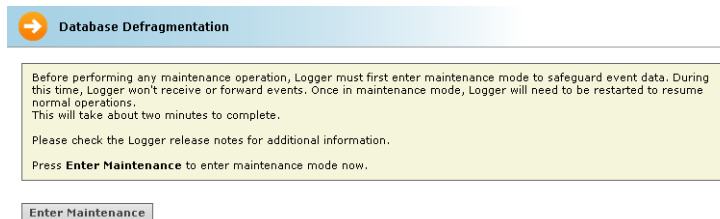
To defragment a Logger:



- You can perform this process only if you have the “Enable Maintenance Mode” privilege set to Yes (System Admin > User/Groups > Manage Groups > System Admin Group).
- If the defragmentation process fails at any point, reboot the Logger and restart the process from the beginning.

- 1 Click **Configuration > System Maintenance**.
- 2 Click **Database Defragmentation**.
- 3 Click **Enter Maintenance** so that the Logger can enter maintenance mode.

For more information about maintenance mode, see [“System Maintenance” on page 208](#).



- 4 A minimum amount of free storage is required for the database defragmentation process to proceed. Therefore, Logger performs a check to determine free storage when entering maintenance mode.

If required amount of free storage is found and Logger successfully enters maintenance mode, the following screen is displayed. Click **Begin Defragmentation** to start the defragmentation process.

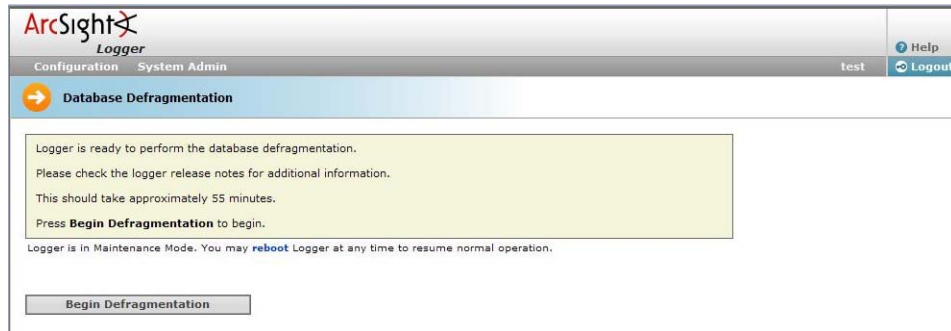
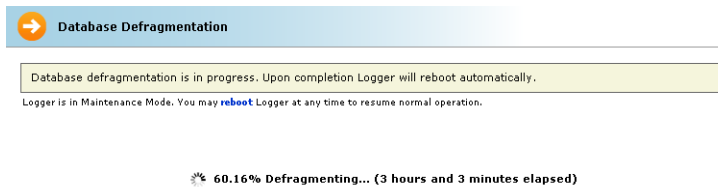


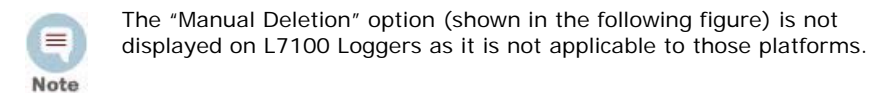
Figure 6-18 Begin Database Defragmentation

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.



If the required storage is not found, Logger prompts you to free sufficient space, as shown in the following example:



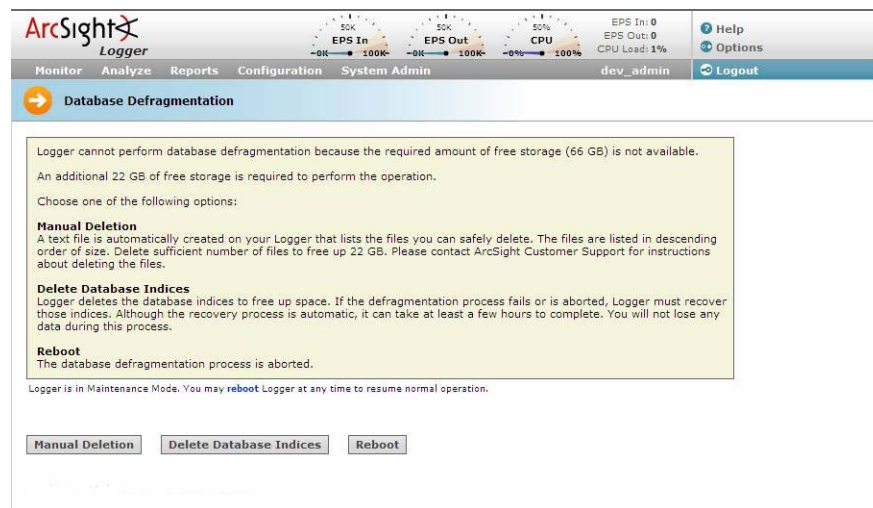
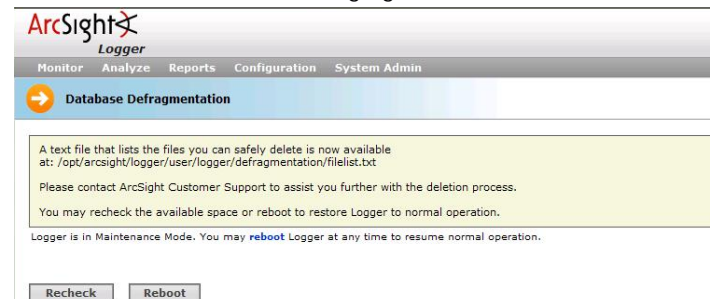


Figure 6-19 Required storage for Database Defragmentation is not available

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 213](#). 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 6-18 on page 212](#) is displayed. Click **Begin Defragmentation** to proceed further.

If sufficient storage is still not found, the screen in [Figure 6-19 on page 213](#) is displayed. Choose from the listed options to create additional space. See [“You](#)

can choose from one of the following options:" on page 213 for more information.



Note

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

◆ **Delete Database Indices**

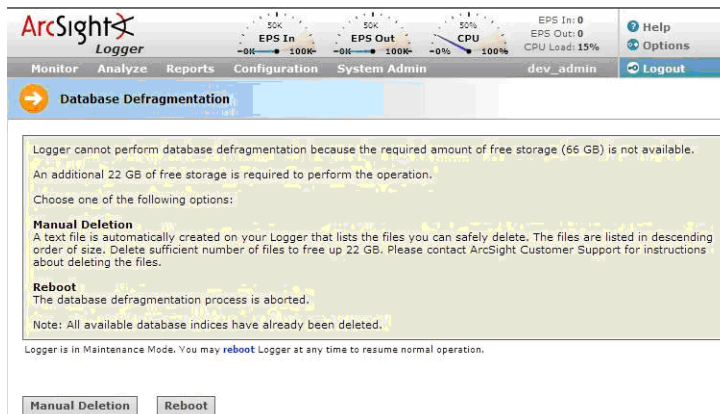
Logger automatically deletes a sufficient number of database indices, starting with the largest index, to free up the required amount of storage. If sufficient space becomes available after deleting database indices, defragmentation proceeds further automatically.

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



Note

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



Follow these steps to proceed:

i Click **Manual Deletion**.

A text file is created on your Logger that lists the files you can safely delete. The files are listed in descending order of size in a text file.

ii Click **Reboot**.

Logger exits the maintenance mode.

iii Contact ArcSight Customer Support for instructions on manually deleting the files.

You can delete sufficient number of files to free up storage.

- iv After deleting the files, restart the defragmentation process from [Step 1 on page 211](#).



Note

If the defragmentation process fails or is aborted at any time, Logger must recover those indices. Although the recovery process is automatic, it can take at least a few hours to complete. You will not lose any data during this process.

◆ **Reboot**

The database defragmentation process is aborted and Logger returns to the state it was in before you started the defragmentation utility.

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 6-20](#) appears.

- 2 When the Summary Status is Completed, click **Download** to retrieve the system log files are compressed into a single zip file.

Retrieve Snapshot Status		
Summary		
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 6-20 Retrieve Logs provides snapshot status.

Exporting and Importing Content

Starting with Logger v3.0, you can export and import content 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. When you need to use that content for any of the situations described previously, simply import the saved content.

The Logger content that you can export and import is Alerts and Filters.

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.

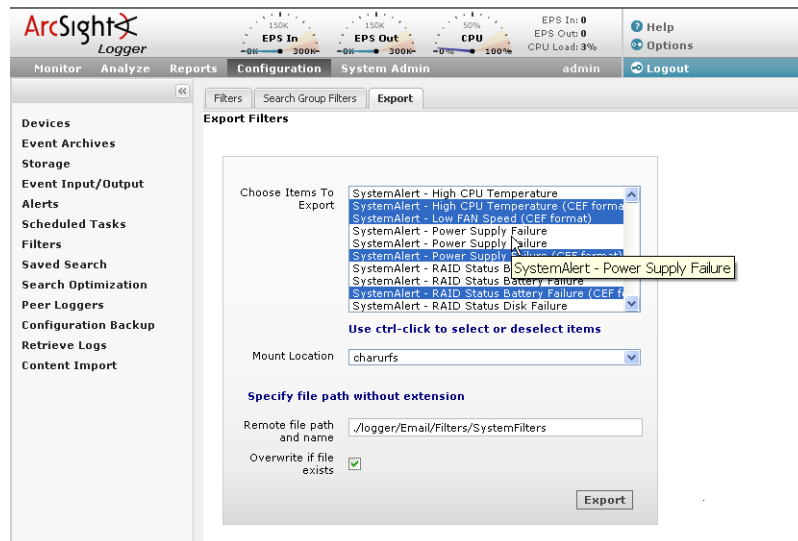
Importing Guidelines

- 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.

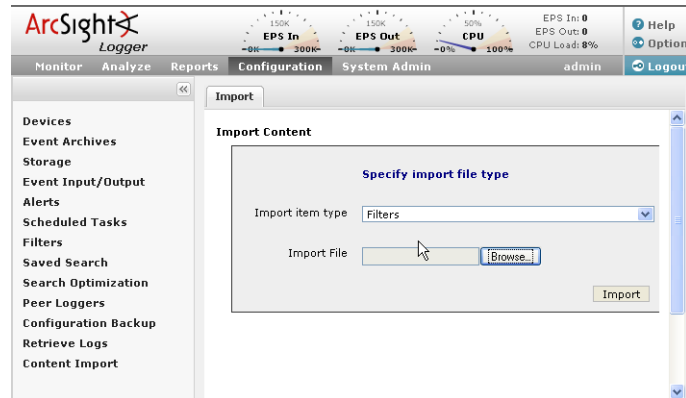
Exporting Content



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 Logger, go to Step 6.
 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 231](#).
- 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 provides access to system and platform settings. You create and manage Users in the System Admin tab, as well.

In this chapter:

- ["Reboot" on page 219](#)
- ["DNS Settings" on page 220](#)
- ["Hosts" on page 221](#)
- ["Network" on page 221](#)
- ["Time/NTP" on page 223](#)
- ["SMTP Settings" on page 226](#)
- ["Static Routes" on page 227](#)
- ["License & Update" on page 227](#)
- ["Process Status" on page 228](#)
- ["Support Login" on page 229](#)
- ["Logs \(Audit and Error\)" on page 230](#)
- ["Audit Forwarding" on page 230](#)
- ["Storage" on page 231](#)
- ["SAN" on page 235](#)
- ["Security" on page 238](#)
- ["Users/Groups" on page 247](#)

System Admin

The System Admin tab, like the Configuration tab, has an associated sub-menu. On the pages associated with the different sub-menu choices, you can configure network, storage, and security settings. In addition, the System Admin tab is where user accounts are managed.

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.



Caution

During reboot, Logger is not able to receive events. Events may be lost while the Logger reboots, unless SmartConnectors are used. SmartConnectors cache events when destinations like Logger are temporarily unavailable.

DNS Settings

ArcSight Platform Settings

DNS Hosts Network Time/NTP SMTP Static Routes

DNS Settings

Please enter DNS Servers

Primary IP Address
0.0.0.0

Secondary IP Address
0.0.0.0

Search Domains
localdomain

Update Settings

Figure 7-1 Domain Name Servers page

To change DNS settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.
- 3 In the **DNS** tab on the ArcSight Platform Settings page, enter new values for the IP address of the primary and secondary DNS servers, or edit the list of search domains.
- 4 Click **Update Settings** to make the changes, or click another tab or sub-menu to cancel. You must reboot Logger for the changes to become effective. See ["Reboot" on page 219](#).

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.

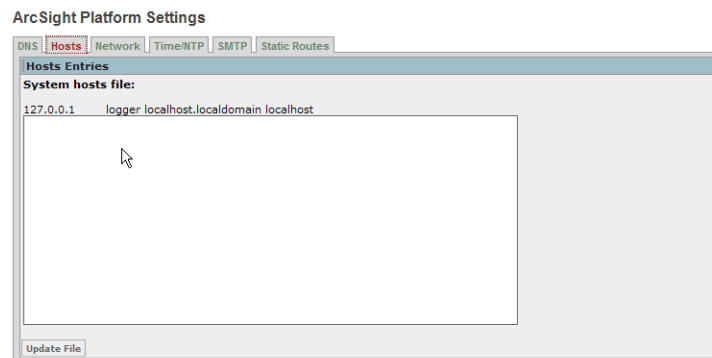


Figure 7-2 Hosts tab allows direct editing of etc/hosts file

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 Connector Appliance for the changes to take effect. See ["Reboot" on page 219](#).

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 219.)

ArcSight Platform Settings

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

Network Settings

Note: Settings take effect after reboot.

System Hostname
localhost

Default Gateway
192.168.35.1

☐ Automatically route outbound packets
(interface homing)

NIC 'ETH0'
IP Address
192.168.35.35
Mask
255.255.255.0
Speed/Duplex
Auto (recommended) ▼

NIC 'ETH1'
IP Address
192.168.36.35
Mask
255.255.255.0
Speed/Duplex
Auto (recommended) ▼

Update Settings

Figure 7-3 Network Settings page

To change network settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.
- 3 In the **Network** tab on the ArcSight Platform Settings page, enter new values for the following fields.

Parameter	Description
System Hostname	<p>The network host name for this Logger. A meaningful name will help, for example, when making a set of Loggers aware of each other.</p> <p>This name must be identical to the domain specified in the Certificate Signing Request, described in “Generating a Certificate Signing Request” on page 239.</p>
Default Gateway	The IP address of the default gateway.
Automatically route outbound packets	<p>When this feature is enabled (checked box), the response packets are sent back on the same Logger interface on which the request packets had arrived. Doing so can improve performance as the routing decisions do not need to be made (using the default gateway information and static routes) to send packets out from the Logger. If you have default gateway and static routes configured, they are ignored when this feature is enabled.</p> <p>When this feature is disabled (unchecked box), the default gateway and static routes (if configured) are used to determine the interface through which the response packets should leave the Logger.</p> <p>If you configure only one network interface, this setting does not provide any additional benefits.</p>

Parameter	Description
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	Choose a speed and duplex mode, or let Logger automatically determine the network speed: Auto (recommended) 10 Mbps - Half Duplex 10 Mbps - Full Duplex 100 Mbps - Half Duplex 100 Mbps - Full Duplex 1 Gbps - Full Duplex

- 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 219](#)) to commit changes to network settings.
- It is important that the System hostname is resolvable by DNS and that it resolves to the Logger's IP address. Performance is significantly affected if DNS cannot resolve the host name.

Time/NTP

The Time/NTP settings page enables you to configure system time, date, local timezone, and NTP servers. The times displayed for Logger operations such as searches, reports, and scheduled jobs are in the Logger's local time zone.



Tip

Because precise time stamping of events is critical for accurate and reliable log management, ArcSight strongly recommends using an NTP server.

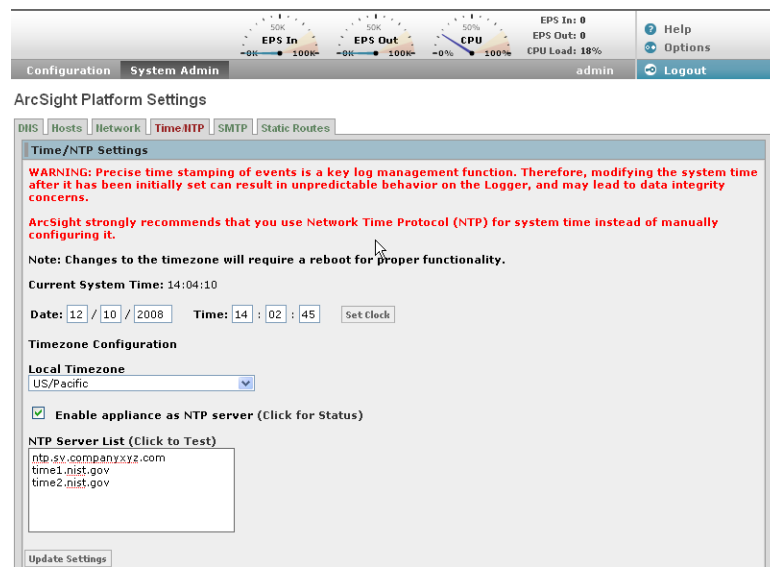


Figure 7-4 Time Settings page

To change the current Logger time:



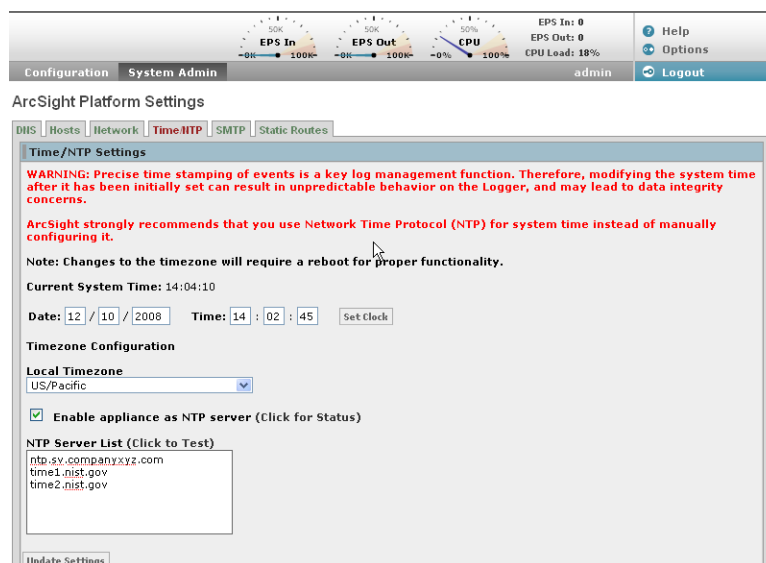
Caution

Modifying the system time after it has been initially set can result in unpredictable behavior on the Logger, thus compromising data integrity.

ArcSight strongly recommends that you use Network Time Protocol (NTP) for system time instead of manually configuring it. However, if you need to change the system time manually, please contact ArcSight Customer Support for guidance.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.
- 3 In the **Time/NTP** tab on the ArcSight Platform Settings page, enter new values for hour, minute, second, month, day, or year.
- 4 Click **Set Clock** to set the Logger clock to the new values.

To change time configuration:



- 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 219](#).

Impact of Daylight Savings Time Change on Logger Operations

Scheduled operations on Logger such as reports, event archives, and file transfers are impacted when system time is adjusted on the Logger at the start and end of the daylight saving time period (DST). The operations scheduled for the hour lost at the start of DST

(for example, on March 8, 2009) are not run on the day of time adjustment. Similarly, operations scheduled for the hour gained at the end of the DST (for example, on November 1, 2009) are run at standard time instead of the DST time.

Examples:

- A report scheduled to run at 1 a.m. DST on November 1, 2009 will run at 1 a.m. standard time, which is an hour later than the DST time on that day.
- A report scheduled to run at 2 a.m. on November 1, 2009 will run at 2 a.m.; however, due to time adjustment, an hour later than it ran on the previous day (October 31, 2009).
- A report scheduled to run at 2 a.m. on March 8, 2009 will not run.

SMTP Settings

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

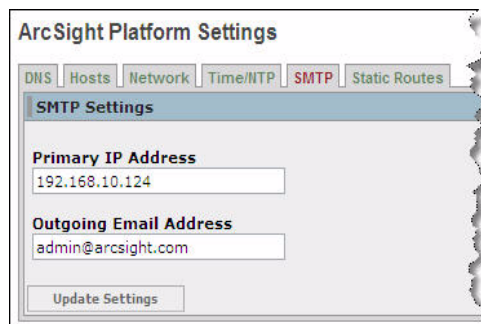


Figure 7-5 Simple Mail Transfer Protocol (SMTP) settings

To change SMTP configuration:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.
- 3 In the **SMTP** tab on the ArcSight Platform Settings page, enter new values for the following fields.

Parameter	Description
Primary SMTP Address	Enter the IP address of the SMTP server that will process outgoing e-mail.
Outgoing Email Address	The e-mail address that will appear in the From: field of outbound e-mail.

- 4 Click **Update Settings** to make the changes, or click another tab or sub-menu to cancel. Changes take effect immediately; reboot is not required.

Static Routes

Advanced users can specify static routes for either or both network adapters. The Static Routes page displays a table of all specified static routes.

ArcSight Platform Settings

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

Static Routes Settings

Add Static Route

Network Adapter:

Dest. Type:

Destination:

Subnet Mask:

Gateway:

Current Static Routes

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface	Last Update	Action
<input type="button" value="Add Static Route"/>									

Figure 7-6 Static Routes page

To add a static route:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** from the System section.
- 3 In the **Static Routes** tab on the ArcSight Platform Settings page, click **Add Static Route**.
- 4 Enter new values for the following fields.

Parameter	Description
Network Adapter	Choose the network interface card (NIC).
Destination Type	Select Network or Host.
Destination	Specify the IP address for the static route destination.
Subnet Mask	Enter the subnet mask (for example, 255.255.255.0) for network only.
Gateway	Specify the IP address for the default gateway.

- 5 Click **Create Static Route** to add the new static route to the table, or click another tab or sub-menu to cancel. You must reboot Logger for the changes to become effective. See [“Reboot” on page 219](#).

License & Update

Updating system software requires uploading an upgrade file provided by ArcSight Customer Support using the System Update page. The System Update page also displays the elapsed time since the appliance was last rebooted, and the version of the Logger components. The Logger version and build number is found at 'arcsight-logger'.

To upload an upgrade file:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **License & Update** from the System section.
- 3 Click **Browse** to locate the file.
- 4 Click **Upload Update**.



Note

System Update will take effect after the next reboot. To update immediately, reboot the system after performing a System Update. See [“Reboot” on page 219](#).

Process Status

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Process Status** from the System section to display a page similar to the one shown in [Figure 7-7](#).

Process Status

System	Status	Load			CPU	Memory
logger	running	[0.30]	[0.28]	[0.25]	1.4%us 0.2%sy 0.1%wa	1909144 kB [7.7%]

Process	Status	Uptime	CPU	Memory	Memory (kB)
apache	running	23h 51m	0.0%	0.0%	7300
aps	running	23h 51m	0.0%	0.7%	191736
connector	running	23h 51m	0.0%	0.0%	568
insp	running	23h 49m	0.0%	0.1%	27604
mysqld	running	23h 51m	0.0%	0.0%	21460
nullmailer	running	23h 49m	0.0%	0.0%	816
postgresql	running	23h 51m	0.0%	0.0%	9244
processors	running	23h 49m	0.0%	0.6%	149160
receivers	running	23h 49m	0.0%	0.2%	50700
reportengine	running	23h 49m	0.0%	0.3%	94856
servers	running	23h 51m	0.0%	2.3%	587224
web	running	23h 49m	1.3%	1.6%	406388

Figure 7-7 Process Status page

(In the process list, processors refers to Forwarders.)

Additional system information, specifically the system uptime and component versions, is available on the System Update page. (See [“License & Update” on page 227](#).)

Each process is a hyperlink. Clicking on an individual process displays more detail about that process, as shown in [Figure 7-7](#).

Status detail for apache

Parameter	Value
children	15
cpu_percent	0.0%
cpu_percent_total	0.0%
data_collected	Wed Sep 2 13:58:02 2009
memory_kilobytes	7300
memory_kilobytes_total	250200
memory_percent	0.0%
memory_percent_total	1.0%
monitoring_status	monitored
parent_pid	1
pid	4279
status	running
uptime	23h 53m

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

[BACK](#) [RESTART](#)

Figure 7-8 Process Status detail for apache

Support Login

Starting with this release, when Customer Support needs access to your appliance for troubleshooting and diagnostics, they work with you to assign a single-use password to the appliance. Doing so enables Support Login access to the appliance. This password is valid only for one support session and is automatically disabled after the session ends. (You can also explicitly disable Support Login access.)

ArcSight
Logger

Monitor Analyze Configuration **System Admin** admin Logout

System
Reboot
Network
License & Update
Process Status
Support Login

Logs
Audit
Error
Forwarding

Storage
CIFS
NFS
RAID Controller

Security
SSL Server Certificate
SSL Client Authentication
FIPS 140-2

Users/Groups
Authentication
Groups
Users
Change Password

Enable/Disable ArcSight Support Login

Set up a password for this appliance that ArcSight Customer Support can use to access this appliance for problem diagnosis. The password is automatically disabled after one support session. Or you can explicitly disable it.

ArcSight Support Login

Access to this appliance is currently disabled.

Request Code: PKA88

Activation Code:

Root Password:

Confirm Root Password:

[Enable Support Login](#)

When you report an issue to ArcSight Customer Support that requires them to access your appliance, they will direct you on enabling Support Login access.

The only circumstance in which you will need to explicitly disable Support Login access is if access was enabled but the support session never occurred.

To **disable** support login access:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Support Login** from the System section.
- 3 Click **Disable Support Login** in the right-side panel.

Logs (Audit and Error)

Logger audit and error logs are available for viewing.

The screenshot shows the ArcSight Logger System Admin interface. The top navigation bar includes Monitor, Analyze, Reports, Configuration, and System Admin (selected). The left sidebar lists various system components like System, Logs, Storage, Security, and Users/Groups. The main content area is titled 'Search Audit Logs' and contains several sections: 'Select Audit Type' with a dropdown for 'Logger Application Audits', 'Select Audit Section' with a dropdown for 'Manage Search', 'Select Date Range' with start and end date pickers (Oct 10, 2009 to Oct 13, 2009), and 'Select User (optional)' with a table of users. The 'View Audit Logs' button is at the bottom.

Login	First Name	Last Name	Email	Phone	Groups
<input checked="" type="checkbox"/>	admin	Default	admin@arcsight.com		Default System Admin Group, Default Logger Rights Group, Default Logger Search Group, Default Logger Report Group,

To view Audit or Error logs:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Audit** (for audit logs) or **Error** (for Error logs) from the Logs section.
- 3 Select the type of log—Application or Platform.
- 4 Select the date range for which you want to obtain the log.
- 5 Click **View Error Logs**.



To search again after clicking **View Audit Logs**, use the browser's Back button.

Audit logs, as Common Event Format (CEF) audit events, can be sent to ArcSight ESM directly for analysis and correlation because the Logger Forwarder supports ESM Manager's event protocol.

For more information about audit event forwarding, see [“Audit Forwarding” on page 230](#). For information about audit events that you can forward, see [Appendix D, Logger Audit Events, on page 281](#).

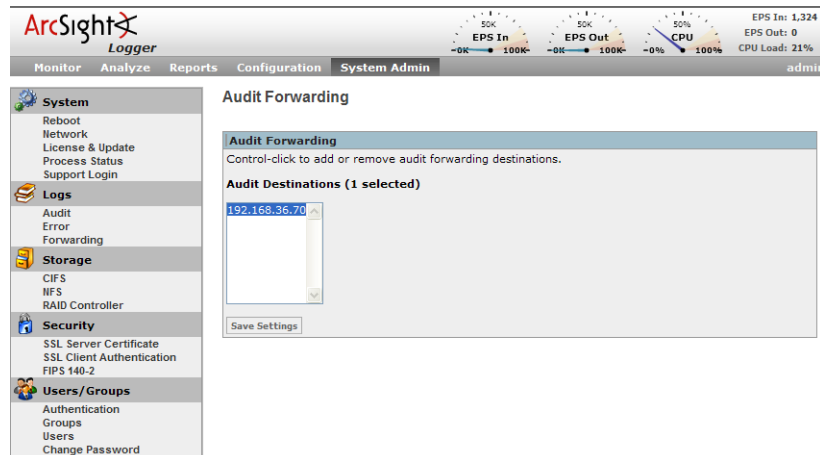
Audit Forwarding

For information about audit events that you can forward, see [Appendix D, Logger Audit Events, on page 281](#).

To forward audit events to specific destinations:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Forwarding** from the Logs section.
- 3 Select destinations from the Audit Destinations list, as shown in the following figure. Click on a destination to select a single destination, or Ctrl+click to select or de-select

multiple destinations. The destinations are ESM destinations that you configure on the ESM Destinations page (**Configuration > Event Input/Output > ESM Destinations**).



- 4 Click **Save Settings**.

Storage

Logger can mount NFS and CIFS shares. As a result, it can read log files and event data from UNIX, Linux, Windows remote hosts, and any Network Attached Storage (NAS) solutions based on these operating systems. Loggers with Storage Area Network (SAN) capability can also interface with a SAN.

The Storage tab includes the ability to configure NFS and CIFS mounts for archiving data and configure LUNs (on systems that support SAN).

In addition, this tab provides status of the hard disk array (RAID) controller and specific system processes.

CIFS Settings

Logger can mount a CIFS remote file system (Windows share) to archive data such as events, exported filters and alerts, and saved searches. A CIFS file system cannot be used as the primary storage device for Logger.

Before you mount a Windows share to a Logger, make sure

- A user account with read-write privileges to the share exists on the Windows system.
- The folder to which you are establishing the mount point is configured for sharing.

ArcSight Logger

Monitor Analyze Reports Configuration **System Admin** admin Help Options Logout

System Admin

- System
 - Reboot
 - Network
 - License & Update
 - Process Status
 - Support Login
- Logs
 - Audit
 - Error
 - Forwarding
- Storage
 - CIFS
 - NFS
 - RAID Controller
- Security
 - SSL Server Certificate
 - FIPS 140-2
- Users/Groups
 - Authentication
 - Groups
 - Users
 - Change Password

CIFS Mount Administration

Add Remote Mount Point

Name: Hurricane

File System Mount Options: rw

Remote Hostname/IP Address: 192.0.2.11

Username: admin

Password:

Share Name: CIFS1

Description: CIFS archival for Logger appliance.

Save CIFS Mount Cancel

To add a CIFS mount:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **CIFS** under the Storage section in the left panel.
- 3 Click **Add CIFS Mount** in the right panel.
- 4 Enter values for the following fields.

Parameter	Description
Name	A meaningful name for the Windows share. The name cannot contain spaces. This name is used locally on your Logger to refer to the mount point and needs to be specified when configuring archive settings for data that will be stored on the share.
File System Mount Options	Autofs options. For example, <code>ro</code> for read-only from the remote host, <code>rw</code> for read-write, or <code>hard</code> to keep retrying until the remote host responds. Note: Even if you configure <code>rw</code> permission at your mount point, <code>rw</code> permission is not granted to the remote host if the host is configured to allow read-only access.
Remote Hostname / IP Address	Host name or IP address of the host to which you are creating the CIFS mount.
Username	Name of the user account with read-write privileges to the Windows share. Make sure the username is prefixed with the domain information. For example, <code>tahoe/arcsight</code> .
Password	Password for the user name specified above.

Parameter	Description
Share Name	<p>The folder on the Windows host to which you are creating the CIFS mount. For example, <code>logger_logs</code>.</p> <p>This folder needs to be configured for sharing. (Typically, to configure a Windows folder for sharing, right click on the folder name > Properties > Sharing.)</p> <p>Note: If you cannot mount successfully, try specifying a leading slash (\) in the remote path. For example, <code>\connector_logs</code>.</p>
Description	A meaningful description of the mount point.

- 5 Click **Save CIFS Mount**.
- 6 (Optional) Click **test** in the Action column of the mount point you added to test connectivity to the Windows share.

To edit a CIFS mount:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **CIFS** under the Storage section in the left panel.
- 3 Click **edit** in the Action column for the CIFS mount that you want to edit. Change field values as needed.
- 4 Click **Save CIFS Mount**.

To delete a CIFS mount:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **CIFS** under the Storage section in the left panel.
- 3 Click **delete** in the Action column for the CIFS mount that you want to delete.
- 4 Confirm the deletion.



Caution

Deleting the CIFS mount (or detaching the SAN) used for Event Archive (see [“Archive Storage Settings” on page 168](#)) will permanently disable the Event Archive feature.

Network File System (NFS) Settings

An NFS mounted system can be used to archive data such as events, exported filters and alerts, and saved searches. Use of a Network File System (NFS) as primary storage for Logger events is not recommended.

Before you mount an NFS share of a remote system, make sure you grant Logger read and write permission on that system. The account name is 'arcsight', but use numeric ids instead: 1500 for uid, or 750 for gid.

Logger supports only NFS v3.0.



Tip

ArcSight recommends creating a Configuration Backup whenever NFS settings are changed. A current backup is useful for disaster recovery. For more information, see [“Configuration Backup and Restore” on page 206](#).

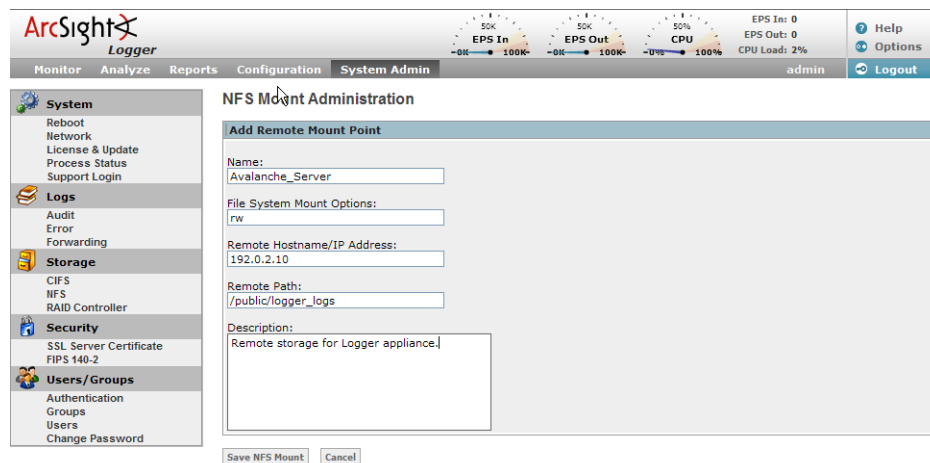


Figure 7-9 NFS Mount Administration page

To add an NFS mount:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **NFS** under the Storage section in the left panel.
- 3 Click **Add NFS Mount** in the right panel.
- 4 Enter new values for the following fields:

Parameter	Description
Name	A name for the network file system mount. The name cannot contain spaces.
File System Mount Options	Autofs options. For example, <code>ro</code> for read-only from the remote host, <code>rw</code> for read-write, or <code>hard</code> to keep retrying until the remote host responds. Note: Even if you configure <code>rw</code> permission at your mount point, <code>rw</code> permission is not granted to the remote host if the host is configured to allow read-only access.
Remote Hostname / IP Address	Host name or IP address of the host to which you are creating the NFS mount.
Remote Path	The folder on the remote host that will act as the root of the network file system mount. For example, <code>/public/logger_logs</code> .
Description	A meaningful description of the mount point.

- 5 Click **Save NFS Mount**.
- 6 (Optional) Click **test** in the Action column of the mount point you added to test the network file system connectivity.

To edit an NFS mount:

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

- 2 Click **NFS** under the Storage section in the left panel.
- 3 Click **edit** in the Action column for the NFS mount that you want to edit. Change field values as needed.
- 4 Click **Save NFS Mount** to make the changes, or click **Cancel** to quit.

To delete an NFS mount:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **NFS** under the Storage section in the left panel.
- 3 Click **delete** in the Action column for the CIFS mount that you want to delete.
- 4 Confirm the deletion.



Deleting the NFS mount (or detaching the SAN) used for Event Archive (see [“Archive Storage Settings” on page 168](#)) will permanently disable the Event Archive feature.

SAN

Some models of Logger include the ability to connect to a Storage Area Network (SAN) for various purposes. SANs contain Logical Units (LUNs), identified by their World Wide Name. As shown in [Figure 7-10](#), a LUN's Attachment Status can be 'available,' 'attached,' or 'detached. LUNs in a SAN are in one state at a time. Actions such as “attach” change from one state to another.’

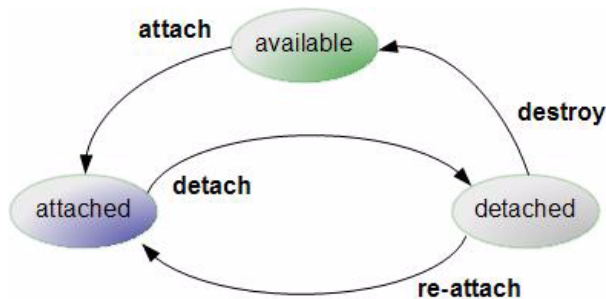


Figure 7-10 SAN Logical Unit state diagram.



Destroying a Logical Unit that has been detached puts that LUN into a state in which a subsequent attach will erase any data stored on the Logical Unit. If a LUN is accidentally destroyed, ArcSight Customer Support may be able to recover the data, provided the LUN is not attached.

The following table summarizes the states and possible actions:

Table 7-1 Logical Unit States and 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

Attachment Status	Actions	Description
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

Changes to the SAN (adding or removing LUNs, for example) will not be reflected in Logger until Logger is rebooted.

To attach a LUN:



Note

- Logger can attach to only one LUN (on SAN) at a time for primary storage. You can add more LUNs for event archival, configuration backup, and export.
- Although the HBA card on the Logger contains two physical interfaces, only a single interface can be enabled. Therefore multi-path support is not available currently.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN Storage** under the Storage section in the left panel.
- 3 Find the LUN in the SAN Logical Unit List.
- 4 In the Action column, click **attach** for that row.
- 5 The LUN's Attachment Status will change to 'attached' when the LUN is ready for use.

To detach a LUN:



Note

LUNs that are used for primary storage may not be detached.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN Storage** under the Storage section in the left panel.
- 3 In the SAN Logical Unit List, locate the LUN to be detached. In the action column, click **detach** for that row. Change field values as needed.
- 4 The LUN's Attachment Status will change to 'detached.'

To re-attach a LUN:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN Storage** under the Storage section in the left panel.
- 3 In the SAN Logical Unit List, locate the LUN to be reattached. The LUN must be in the 'detached' state. In the action column, click **re-attach** for that row.
- 4 The LUN's Attachment Status will change to 'attached.'

To release a LUN:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN Storage** under the Storage section in the left panel.
- 3 In the SAN Logical Unit List, locate the LUN to be released. The LUN must be in the 'detached' state. In the action column, click **destroy** for that row.
- 4 The LUN's Attachment Status will change to "available".

SAN Storage Administration

SAN Logical Unit List						
Name	Local Device Name	World Wide Name	Size	Attachment Status	Action	
SAN1	/dev/sdb	5006016830224f88:0000000000000000	49.95 GB	attached	detach	
san1	/dev/sdc	5006016830224f88:0001000000000000	49.95 GB	available	destroy re-attach	

Figure 7-11 SAN Storage Administration page

Restoring a SAN

To restore a SAN to either the Logger to which it was formerly attached or a new Logger (in the case of disaster recovery), follow these steps:

- 1 With Logger powered off, attach the SAN physically. Turn on Logger.
- 2 Restore the configuration to Logger. ArcSight recommends backing up the configuration regularly so that a backup file will be available for this purpose. If no backup file is available, skip this step and manually add receivers, forwarders, users, and so on, after SAN has been restored.
- 3 Enable one-time Support Login (see ["Support Login" on page 229](#)). Contact ArcSight Customer Support.
- 4 ArcSight Customer Support will login remotely, stop all Logger processes by issuing the command

```
/opt/local/monit/bin/monit stop all
```

and migrate the internal database to the SAN by creating a symbolic link with the command

```
ln -s <remote storage path> /opt/local/pgsqldata
```

When Customer Support has finished these tasks, reboot Logger. This will disable future Support Logins.

RAID Controller

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **RAID Controller** under the Storage section in the left panel to display a page similar to the one shown in [Figure 7-12](#).



Note

Logger hardware models use different RAID controllers, which display information differently.

```

Status of RAID Controller
+-----+
| General Controller Information |
+-----+
Type: RAID-5
State: Optimal

Versions:
Product Name   : PERC 6/i, Integrated
Serial No      : 1122334455667788
FW Package Build: 6.1.1-0047

Image Versions In Flash:
FW Version      : 1.21.02-0528
BIOS Version     : 2.01.00
WebBIOS Version  : 1.1-46-e_15-Rel
Ctrl-R Version   : 1.02-014B
Boot Block Version : 1.00.00.01-0011

HW Configuration:
SAS Address     : 50024e805edb8600
BBU              : Present
Alarm           : Absent
NVRAM           : Present
Serial Debugger  : Present
Memory          : Present
Flash           : Present
Memory Size     : 256MB

Device Present:
Virtual Drives   : 2
  Degraded       : 0
  Offline        : 0
Physical Devices : 7
  Disks          : 6
  Critical Disks : 0
  Failed Disks   : 0

Error Counters:
Memory Correctable Errors : 0
Memory Uncorrectable Errors : 0

Drive states:
0: Online
1: Online
2: Online

```

Figure 7-12 RAID Controller Information page

Obviously, this information is highly technical. It is not needed during normal Logger operations, but it can be helpful for diagnosing specific hardware issues. Due to the redundant nature of RAID storage, unit failure does not disable Logger. Instead, performance degrades. Use this report to determine whether a performance issue is caused by a disk failure. ArcSight Customer Support can use this information to better diagnose problems, as well.

Security

Security settings enable you to configure SSL Server certificates, enable and disable FIPS (Federal Information Processing Standards) mode on the Logger, and configure SSL client authentication for CAC support.

SSL Server Certificate

Logger uses Secure Sockets Layer (SSL) technology to communicate securely over an encrypted channel with its clients—users, SmartConnectors when using the SmartMessaging technology, and peer Loggers. To establish a typical SSL session, an SSL certificate is required on the server (Logger) side and a truststore is required on the client side. The truststore contains a list of Certificate Authorities (CA) that the client trusts.

When a client initiates communication with Logger, the Logger sends its SSL certificate to the client to authenticate itself. The client checks its truststore to validate the certificate. (In addition, the client verifies whether the hostname in the certificate matches the one with which it initiated communication, and the current time on the client machine is within the validity range specified in the certificate.) If the certificate is validated, a session key is exchanged between the client and the Logger. This key is used to encrypt and decrypt data exchanged between the Logger and the client.

Logger ships with a self-signed certificate. Although you can use this certificate, ArcSight strongly recommends using a CA-signed certificate.

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 resulting CSR should be sent to a CA, such as VeriSign, which responds with a signed certificate file.

The screenshot shows the ArcSight Logger System Admin interface. The left sidebar contains a navigation menu with categories: System (Reboot, Network, License & Update, Process Status, Support Login), Logs (Audit, Error, Forwarding), Storage (CIFS, NFS, RAID Controller), Security (SSL Server Certificate, SSL Client Authentication, FIPS 140-2), and Users/Groups (Authentication, Groups, Users, Change Password). The main content area is titled 'ArcSight SSL Settings' and includes tabs for 'Generate CSR', 'Install Cert', and 'View Results'. The 'Generate CSR' tab is active, displaying a form titled 'Generate Certificate Signing Request' with the instruction 'Please enter the Certificate Settings'. The form fields are: Country (2-letter code) [US], State/Province [California], City/Locality [Cupertino], Organization Name [ArcSight, Inc.], Organizational Unit [Support Team], Hostname [loggerA.arcsight.com], Email Address [support@arcsight.com], and Private Key Password []. A 'Generate CSR' button is at the bottom left of the form. The top of the interface shows system status metrics: EPS In (50K), EPS Out (50K), CPU (50%), EPS In: 1.324, EPS Out: 0, and CPU Load: 72%.

Figure 7-13 Certificate Signing Request page

To generate a certificate signing request:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Server Certificate** under the Security section in the left panel to display the Generate Certificate Signing Request page, as shown in [Figure 7-13 on page 239](#).
- 3 Enter new values for the following fields:

Parameter	Description
Country	A two-letter country code, such as 'US' for the United States.
State / Province	State or province name, such as 'California.'
City / Locality	City name, such as 'Cupertino.'
Organization Name	Company name, governmental entity, or similar overall organization.
Organizational Unit	Division or department within the organization.
Hostname	The host name of this Logger. This name should match 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 221 .
Email Address	The e-mail address of the administrator or contact person with regard to this CSR.

Parameter	Description
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.

- 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 7-14 Install Certificate page

To install a signed certificate:

- Click **System Admin** from the top-level menu bar.
- Click **SSL Server Certificate** under the Security section in the left panel.
- On the **Install Cert** tab (as shown in [Figure 7-14 on page 240](#)), click **Browse** to find the signed certificate file on your local file system.
- 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 238](#) to load the certificate.
- 2 Enable client certificate authentication, as described in [“Client Certificate Authentication” on page 249](#).



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 241](#).

If the client certificates used to authenticate with Logger are signed by different CAs, make sure you upload root certificates of **all** CAs.

If the client certificates are **self-signed**, upload the public portion of the client certificate.

- 4 Configure a Logger user name for each user who will be connecting to the Logger using a client certificate, as described in [“Users” on page 256](#).
- 5 (Optional) Upload a certificate revocation list (CRL), as described in [“Uploading a Certificate Revocation List” on page 242](#).
- 6 (Optional) If this Logger is configured to use **only** SSL Client Authentication, make sure this Logger's Authorization ID and Code are appropriately configured on other Loggers that peer with it. For more information, see [“Peer Loggers” on page 202](#).

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 245](#) 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 183](#) to complete configuration of this setup.
- Peer Loggers
Loggers running v4.0 automatically use FIPS 140-2 compliant algorithms. Therefore, no action is required on a peer Logger running version 4.0. A FIPS-enabled version 4.0 Logger can communicate with a non-FIPS enabled Logger running Logger v4.0. Additionally, a Logger running v3.0 SP1 Patch1 can be peered with a Logger running v4.0.
- Connector Appliance
If your Logger platform includes an integrated Connector Appliance, both products operate in FIPS mode when you enable FIPS on the Logger. However, you might need to do additional configuration on the Connector Appliance components for FIPS-mode operation. See the *Connector Appliance Administrator’s Guide* for more information.

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.

To enable or disable FIPS mode on Logger:



Note

- Your Logger needs to be set up with a signed SSL certificate before you can enable FIPS 140-2 on it. For more information, see [“SSL Server Certificate” on page 238](#).
- 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.

- 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 System Admin interface. The top navigation bar includes Monitor, Analyze, Reports, Configuration, and System Admin. The left sidebar lists various system components under System, Logs, Storage, Security, and Users/Groups. The main content area is titled 'Enable/Disable FIPS Mode' and 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 columns 'Name' and 'FIPS Enabled'.

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 245.
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 245.
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 245 .

- 1 Follow device configuration steps provided in the SmartConnector's configuration guide (available from the ArcSight Customer Support site at <https://support.arcsight.com>), then follow the installation procedure through installation of the core connector software (SmartConnector Installation step 2).

At Step 3 of the Connector setup, as shown below, click **Cancel** to exit the setup to configure the NSS DB, which is necessary for installing the connector in FIPS-compliant mode.

Step 3: When the installation of ArcSight SmartConnector core component software is finished, the following window is displayed:



- 2 Click **Cancel** to exit the configuration wizard. You will return to this wizard and resume SmartConnector configuration, after

- ◆ Enabling FIPS mode on it, and
- ◆ Importing Logger's certificate

Enable FIPS Mode on the SmartConnector

- a Create an `agent.properties` file at the following location:

`$ARCSIGHT_HOME\current\user\agent`

- b** Enter the following property, then save and close the file.

```
fips.enabled=true
```

Import Logger's Certificate on the SmartConnector

- c** In a DOS prompt window on your SmartConnector machine, from `$ARCSIGHT_HOME\current\bin`, enter the following command to turn off FIPS mode.

```
arcsight runmodutil -fips false -dbdir  
user/agent/nssdb.client
```

- d** Export the Logger certificate file and import it to the SmartConnector's NSS DB as follows:

- i** Export Logger's certificate file from the browser you use to connect to it. Refer to your browser's Help for instructions. For example, to export a Logger's certificate file on Firefox 3.0.x, click **Tools > Options > Encryption > View Certificates > Servers > Select your Logger appliance > Export**. Save the certificate file with a .crt or .cer extension.
- ii** Copy the certificate file you exported in the previous step (in this example, **loggercert.crt**) to the `$ARCSIGHT_HOME\current\bin` directory.

From `$ARCSIGHT_HOME\current\bin`, enter the following:

```
arcsight runcertutil -A -n mykey -t "CT,C,C" -d  
user/agent/nssdb.client -i bin/loggercert.crt
```

- e** Enter the following command to re-enable FIPS mode that you turned off in Step 1:

```
arcsight runmodutil -fips true -dbdir0  
user/agent/nssdb.client
```

- f** Ensure that the SmartConnector can resolve the name specified in the CN value of the Logger certificate's *Subject*: field. If the name is not resolvable, add it to SmartConnector system's Hosts file.

- g** *If you are updating your SmartConnector to be FIPS-compliant*, ensure that the connector's Logger destination host name is same as the CN value in the certificate's *Subject*: field and **exit this procedure**.

If you are installing a new SmartConnector, go to the next step.

- 3** To return to the SmartConnector configuration wizard, enter the following from `$ARCSIGHT_HOME\current\bin`:

```
arcsight connectorsetup
```

- 4** When prompted whether you want to start in Wizard Mode, click **Yes**.
- 5** The Destination selection window is again displayed; return to your SmartConnector Configuration Guide, **SmartConnector Installation step 4** to continue the connector configuration.

Note: When configuring the connector, ensure that the connector's Logger destination host name is same as the CN value in the certificate's *Subject*: field.

For the remainder of the configuration process, see the Configuration Guide for the SmartConnector you selected to install. The specific configuration guide provides information about how to configure the device for event collection, specific installation

parameters required during the configuration process, and a table of vendor-specific field mappings to ArcSight events.

Users/Groups

Authentication Settings

The Authentication settings enable you to specify settings and policies for login, password, and the authentication mechanism to use.

Login

The Authentication Settings page lets you specify the maximum number of simultaneous sessions for a single user account, which may impact system performance.

The form, shown in [Figure 7-15](#), also lets you specify how many seconds of inactivity to allow before automatically ending the current session. The default is 900 (15 minutes).

The screenshot shows the 'Global Settings' window with the 'Login' tab selected. It contains two input fields: 'Max Simultaneous Logins per User' with the value '15' and 'Session Inactivity Timeout in Seconds' with the value '900'. A 'Save Settings' button is at the bottom.

Figure 7-15 Login Settings page, changing the Session Inactivity time-out to 3 minutes.

To change login settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** under the Users/Groups section in the left panel to display the Global Settings page in the Login tab, as shown in [Figure 7-15 on page 247](#).
- 3 Enter new values for the maximum simultaneous logins per user or the session inactivity time-out.
- 4 Click **Save Settings** to make the changes, or click another tab or sub-menu to cancel. You must reboot Logger for the changes to become effective. See ["Reboot" on page 219](#).

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.

Global Settings

Login Passwords Authentication Audit Forwarding

Password Settings

Enable Password Lockout
☐ Yes ☒ No

Number of failed attempts before lockout

Maximum time between attempts (in seconds)

Lockout duration (in minutes)

Enable Password Expiration
☐ Yes ☒ No

Days until password expires

Days before expiration to notify user

Enable Password Validation
☐ Yes ☒ No

Password Length Limits

Minimum password length

Maximum password length

Minimum Requirements

Numeric characters [0-9]

Uppercase characters [A-Z]

Lowercase characters [a-z]

Non-alphanumeric characters [!\$%^*...]

Number of characters different from old password

Save Settings

Figure 7-16 Password Policy Settings page

To change password policy settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** under the Users/Groups section in the left panel to display the Password Settings page in the Password tab, as shown in [Figure 7-16 on page 248](#).
- 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.

Parameter	Description
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 .
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 219](#).

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 250](#).)

Client Certificate Authentication

Even if SSL client certificate authentication is enabled on the Logger, a user name must be defined on it for users to connect to it. See ["Users" on page 256](#) for specifics about setting up a user name for client certificate authentication.

The default 'admin' user is exempt and can log on without a certificate even if client certificate authentication is configured on a Logger.



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

To configure client certificate authentication:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** under the Users/Groups section in the left panel to display the Authentication Settings page in the Authentication tab.
- 3 Update any or all of the parameters listed in the following table:

Parameter	Description
Use client certificate	Select Yes to enable client certificate authentication. Default: No
Require additional password	Select Yes to require a password, in addition to a client certificate, for authentication. This is the password configured for a user's name on Logger. (See "Users" on page 256 for more information.) Default: No
Allow password fallback	Select Yes if a user should be allowed to log in to Logger using only the RADIUS or local password when a certificate is not available or is invalid. Default: No

- 4 Click **Save Settings** to make the changes, or click another tab to cancel.
- 5 Click **Reboot** in the left panel to reboot the appliance.

RADIUS Authentication

If RADIUS authentication is enabled, only user names that are defined as Logger users (see ["Users" on page 256](#)) 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 .
Number of retries	Number of times to retry a RADIUS request. The default is 1 .

- 4 Click **Save Settings** to make the changes, or click another tab or sub-menu to cancel.

Groups

Logger users are granted permissions by membership in a user group. A user group is a set of permissions and a set of users. User groups have types, such as Logger user groups, or Filter user groups.

User Groups

Groups are organized by type, as shown in [Figure 7-17](#). Each user group is one of the following types: System Admin, Logger Rights, Logger Search, or Logger Report.

Each type has a default user group pre-defined, and the default user group has all privileges for its type enabled. To authorize a subset of the default user group's privileges, create a new User Group (as described below) and revoke some privileges. Then move restricted users from the default user group into the newly created group.

Table 7-2 System Admin Groups

Section	Privilege
Reboot	Reboot Logger. (See "Reboot" on page 219.)
Update	Update Logger. (See "License & Update" on page 227.) Enable Maintenance Mode (See "System Maintenance" on page 208.)
System Information	Process Status. (See "Process Status" on page 228.) 3Ware RAID Controller. (See "RAID Controller" on page 237.)

Section	Privilege
SSL Certificates	<p>Generate SSL Certificate Signing Request (CSR). (See “Generating a Certificate Signing Request” on page 239.)</p> <p>Install new SSL certificates. (See “Installing a Signed Certificate” on page 240.)</p>
Platform Settings	<p>Configure DNS settings. (See “DNS Settings” on page 220.)</p> <p>Configure network settings. (See “Network” on page 221.)</p> <p>Configure time settings. (See “Time/NTP” on page 223.)</p> <p>Configure SMTP settings. (See “SMTP Settings” on page 226.)</p> <p>Configure static routes. (See “Static Routes” on page 227.)</p> <p>Configure Hosts File. (See “Hosts” on page 221.)</p> <p>Configure Security Settings. (See “Security” on page 238.)</p>
External File Systems	Configure NFS, CIFS, and SAN settings. (See “Storage” on page 231, “CIFS Settings” on page 231, and “SAN” on page 235.)
Global Settings	<p>Configure login settings. (See “Authentication Settings” on page 247.)</p> <p>Configure password settings. (See “Password” on page 248.)</p> <p>Configure password authentication. (See “Authentication” on page 249.)</p> <p>Configure audit forwarding destination. (See “Audit Forwarding” on page 230.)</p>
System Logs	<p>View Audit Logs. (See “Logs (Audit and Error)” on page 230.)</p> <p>View Error Logs. (See “Logs (Audit and Error)” on page 230.)</p>
User/Groups	<p>Manage users. (See “Users” on page 256.)</p> <p>Manage user groups. (See “User Groups” on page 251.)</p> <p>Run user entitlement reports.</p>
Console Access	<p>Allow console access. (See “Connecting to the Command Line Interface” on page 17.)</p> <p>Control support login access. (See “Support Login” on page 229.)</p>

Table 7-3 Logger Rights Groups

Section	Privilege
Monitor	<p>Monitor Logger throughput. (See “Monitor” on page 33.)</p> <p>Monitor Logger throughput on remote peers. (See “Monitor” on page 33 and “Peer Loggers” on page 202.)</p>

Section	Privilege
Application Options	View options. (See “Options” on page 33.) Edit, save, and remove options. (See “Options” on page 33.)
Filters	Use and view shared filters. (See “Filters” on page 194.) Edit, save, and remove shared filters. (See “Filters” on page 194.) Also, import and export filters.
Peers	View registered peers. (See “Peer Loggers” on page 202.) Edit, save, and remove registered peers. (See “Peer Loggers” on page 202.)
Devices and Device Groups	View devices. (See “Devices” on page 164.) Edit, save, and remove devices. (See “Devices” on page 164.) View device groups. (See “Device Groups” on page 165.) Edit, save, and remove device groups. (See “Device Groups” on page 165.)
Receivers	View receivers. (See “Receivers” on page 173.) Edit, save, and remove receivers. (See “Receivers” on page 173.)
Forwarders and Alerts	View forwarders and alerts. (See “Forwarders” on page 179 and “Alerts” on page 187.) Edit, save, and remove forwarders and alerts. (See “Forwarders” on page 179 and “Alerts” on page 187.) For alerts, this privilege enables you to import and export them.
ESM Connectors	View ESM connectors. (See “ESM Destinations” on page 183.) Edit, save, and remove ESM connectors. (See “ESM Destinations” on page 183.)
Search Filters	View search group filters (aka user group filters). (See “Search Group Filters” on page 195.) Edit, save, and remove search group filters. (See “Search Group Filters” on page 195.)
Configuration Backup	View backups. (See “Configuration Backup and Restore” on page 206.) Edit, save, and remove backups. (See “Configuration Backup and Restore” on page 206.)
Retrieve Logs	Download system logs. (See “Retrieve Logs” on page 215.)
Scheduling	View scheduled tasks. (See “Scheduled Tasks” on page 193.)
Storage Groups	View storage groups. (See “Storage Groups” on page 168.) Edit and add storage groups. (See “Storage Groups” on page 168.)

Section	Privilege
Event Archive/Restore	View event archives. (See “Event Archives” on page 166.) Edit, save, and remove event archives. (See “Event Archives” on page 166.)
Saved Search	View saved search. (See “Saved Search” on page 197.) Edit, save, and remove saved search. (See “Saved Search Jobs” on page 198.)

Table 7-4 Logger Search Groups

Section	Privilege
Search	Search for events. (See “The Need to Search Events” on page 39.) Search for events on remote peers. (See “Searching Peer Loggers (Distributed Search)” on page 60.)

Table 7-5 Logger Report Groups.

Section	Privilege
Report	Global access to all report objects and permission to change reporting configuration. (See Chapter 5, Reporting, on page 75.) 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 106.) Edit and save report style. This overrides the corresponding permission on individual report groups. (See “Applying Report Template Styles” on page 154.) View all published reports. This overrides the corresponding permission on individual report groups. (See Chapter 5, Reporting, on page 75.) View, run, and schedule all reports. This overrides the corresponding permission on individual report groups. (See “Running, Viewing, and Publishing Reports” on page 95 and “Scheduling Reports” on page 155.) Edit and save reports. This overrides the corresponding permission on individual report groups. (See Chapter 5, Reporting, on page 75.)

Each individual report group--Default Reports, Configuration Monitoring, Intrusion Monitoring, or SANS Top 5, for example--will have its own set of rights. Each report group

will have privileges for View published reports, View, run, and schedule reports, and Edit and save reports.

Groups Administration

Add User Group

System Admin Groups			
Name	Description	Number Of Members	Action
Default System Admin Group	The default group allows all system admin operations e.g., reboot, install SSL, platform configuration, and so on.	1	edit

Logger Rights Groups			
Name	Description	Number Of Members	Action
Default Logger Rights Group	The default group allows all logger operations e.g., monitor, creating and editing filters, peers, devices, receivers, forwarders, and so on.	1	edit delete

Logger Search Groups			
Name	Description	Number Of Members	Action
Default Logger Search Group	The default search group allows both local and distributed searches.	1	edit delete

Logger Report Groups			
Name	Description	Number Of Members	Action
Default Logger Report Group	The default report group allows all report operations e.g., view published report, view, run, schedule, edit, delete all reports, and so on.	1	edit delete

Figure 7-17 Groups page

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

To create a new user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Groups** under the Users/Groups section in the left panel to display the page shown in [Figure 7-17](#).
- 3 Click **Add User Group**.
- 4 Enter the definition of the new group.
 - a Define the group by choosing a type and entering a name and description.
 - b Define the group's rights and permissions.
 - c Optionally, add users to the new group.
- 5 Click **Save Group**.

To edit a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Groups** under the Users/Groups section in the left panel to display the page shown in [Figure 7-17](#).
- 3 Identify the group to be edited and click the **edit** link.
- 4 Update the user group information as necessary.
- 5 Click **Save Group**.

To delete a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Groups** under the Users/Groups section in the left panel to display the page shown in [Figure 7-17](#).

- 3 Identify the user group to be deleted and click the **delete** link.

Users

Add User

Add User Information:

Login:

First Name:

Last Name:

Password:

Confirm Password:

Email:

Phone Number:

Restrict To CLI: ☐

Authorization Groups

Select Name	Description
<input type="radio"/> Default Auth Group	This is the default Authorization Group
clear	

System Admin Groups

Select Name	Description
<input type="radio"/> Logger	Logger group
<input type="radio"/> Default System Admin Group	This is the default System Administration Group
clear	

Network Configuration Groups

Select Name	Description
<input type="radio"/> Default Net Config Mgmt Group	This is the default Network Configuration Management Group
clear	

Network Response Groups

Select Name	Description
<input type="radio"/> Default Network Response Group	This is the default Network Response Group
clear	

[Save User](#)

Figure 7-18 Add User page

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

To create a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Users** under the Users/Groups section in the left panel.
- 3 Click **Add User** in the right panel. The form shown in [Figure 7-18](#) is displayed.
- 4 Enter the following parameters.

Parameter	Description
Login	A login name for the user
First Name	User's first name.
	<p>If you enabled SSL client authentication (see “SSL Client Authentication (CAC Authentication)” on page 240), click Use Client DN to enter the Distinguished Name (Certificate Subject) information for the user. Distinguished Name should be in this format:</p> <p>ST=California, C=US, L=Cupertino, O=ArcSight, Inc., OU=Engg Team, CN=UserA D/emailAddress=email@xyz.com</p> <p>Obtain the DN information for a user from the browser that the user will use to connect to the Logger. For example, on Firefox 3.0, click Tools > Options > Encryption > View Certificates > Your Certificates > Select the certificate > View.</p>

Parameter	Description
Last Name	User's last name. This information is not required when creating a user for SSL client authentication.
Password	A password for the user.
Confirm Password	Reenter the password.
Email	An e-mail address for the user.
Phone Number	User's phone number.
Select User Groups	Select the groups to which this user belongs. This setting controls the privileges a user has on this Logger.

5 Click **Save User**.

To edit a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Users** under the Users/Groups section in the left panel.
- 3 Identify the user to be edited and click the **edit** link. Update the user information as necessary.
- 4 Click **Save User**.

To delete a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Users** under the Users/Groups section in the left panel.
- 3 Identify the user to be deleted and click the **delete** link.
- 4 Confirm the delete operation.

Changing Password

Password management is the responsibility of individual users. Users can choose their password, and they may change their password as often as desired.

Change Password For Default Admin

Enter Existing Password:

Old Password:
00 00 00 00 00 00 00 00

New Password:
00 00 00 00 00 00 00 00

Confirm:
00 00 00 00 00 00 00 00

Set Password

Figure 7-19 Change Password page

To change your password:

- 1** Click **System Admin** from the top-level menu bar.
- 2** Click **Change Password** under the Users/Groups section in the left panel to display the Change Password for <User Name> page, as shown in [Figure 7-19 on page 257](#).
- 3** Enter the old password, the new password, and enter the new password a second time to confirm.
- 4** Click **Set Password**.



Passwords are subject to the password policy specified by the Admin user.
See ["Password" on page 248](#).

Appendix A

Common Event Format

Common Event Format (CEF) is an industry standard for the interoperability of event- or log-generating devices. The myriad of formats used for event reporting, especially in the security world, greatly complicates integration. Each vendor has its own format for reporting event information, but these formats often lack key information necessary to integrate the events from their devices.

The CEF standard aims to improve the interoperability of infrastructure devices by better aligning the logging output from participating technology vendors. Vendors implementing the CEF standard log events in a format that is both useful, and more importantly, parse-able by ArcSight or any vendor following the standard. Further, this standard assures that an event and its semantics contain all necessary information.

Common Exchange Format

This specification defines a simple event format that can be readily adopted by vendors of both security and non-security devices. This format is intended to contain the most relevant information and make it easy for event consumers to parse and use events.

To simplify integration, we use syslog as a transport mechanism. This applies a common prefix to each message, containing the date and hostname:

```
Jan 18 11:07:53 zurich message
```

If an event producer is unable to write syslog messages, it is still possible to write the events to a file. In this case, omit the syslog header and start the message with the format defined below.

It is important to note that this part of the message need not be explicitly generated by the event producer. The remainder of the message is formatted using a common prefix composed of fields delimited by a bar ("|") character. The prefix is mandatory and all specified fields need to be present. Additional fields are specified in the Extension. The format is:

```
CEF:Version|Device Vendor|Device Product|Device Version|Device  
Event Class ID|Name|Severity|Extension
```

The *Extension* part of the message is a placeholder for additional fields. Those fields are documented in the Event Dictionary below and are logged as key-value pairs.

Here are definitions for the prefix fields:

Version is an integer and identifies the version of the CEF format. Event consumers use this information to determine what the following fields represent. Currently only version 0

(zero) is established in the above format. Experience may show that other fields need to be added to the "prefix" and therefore require a version number change. Adding new formats is handled through the standards body.

Device Vendor, **Device Product** and **Device Version** are strings that uniquely identify the type of sending device. No two products may use the same device-vendor and deviceproduct pair. There is no central authority managing these pairs. Event producers have to ensure that they assign unique name pairs.

Device Event Class ID is a unique identifier per event-type. This can be a string or an integer. Device Event Class ID identifies the type of event reported. In the intrusion detection system (IDS) world, each signature or rule that detects certain activity has a unique identifier assigned. This is a requirement for other types of devices as well, and helps correlation engines deal with the events.

Name is a string representing a human-readable and understandable description of the event. The event name should not contain information that is specifically mentioned in other fields. For example: "Port scan from 10.0.0.1 targeting 20.1.1.1" is not a good event name. It should be: "Port scan." The other information is redundant and can be picked up from the other fields.

Severity is an integer and reflects the importance of the event. Only numbers from 0 to 10 are allowed, where 10 indicates the most important event.

Extension is a collection of key-value pairs. The keys are part of a predefined set. The standard allows for including additional keys as outlined later. An event can contain any number of key-value pairs in any order, separated by spaces (" "). If a field contains a space, such as a file name, this is okay and can be logged in exactly that manner. For example: fileName=c:\Program Files\ArcSight is a valid token.

Here is a sample message to illustrate appearance:

```
Sep 19 08:26:10 zurich CEF:0|security|threatmanager|1.0|100|worm
successfully stopped|10|src=10.0.0.1 dst=2.1.2.2 spt=1232
```

Here are further details about character encoding:

The entire message has to be **UTF-8** encoded.

If a pipe (|) is used in the prefix, it has to be escaped with a backslash (\). But note that pipes in the extension do not need escaping. Here is an example message:

```
Sep 19 08:26:10 zurich
CEF:0|security|threatmanager|1.0|100|detected a \ in
message|10|src=10.0.0.1 act=blocked a | dst=1.1.1.1
```

If a backslash (\) is used in the prefix, it has to be escaped with another backslash (\). Again, note that backslashes in the extension do not need escaping. Here is an example:

```
Sep 19 08:26:10 zurich
CEF:0|security|threatmanager|1.0|100|detected a \ in
packet|10|src=10.0.0.1 action=blocked a \ dst=1.1.1.1
```

If an equal sign (=) is used in the extensions, it has to be escaped with a backslash (\). Equal signs in the prefix need no escaping. For example:

```
Sep 19 08:26:10 zurich
CEF:0|security|threatmanager|1.0|100|detected a = in
message|10|src=10.0.0.1 action=blocked a \= dst=1.1.1.1
```

Multi-line fields can be sent by Common Event Format (CEF) by encoding the newline character as \n or \r. Note that multiple lines are only allowed in the value part of the extensions. See this example:

```
Sep 19 08:26:10 zurich
CEF:0|security|threatmanager|1.0|100|Detected a threat. No action
needed.|10|src=10.0.0.1 message=Detected a threat.\nNo action
needed.
```

Common Extension Dictionary

The following table contains predefined keys that establish usages for both event producers and consumers. The standard allows for defining additional keys, with the understanding that those fields may not be interpreted by other event consumers.

The table below contains key names as well as the full name for each key. The key name is the one that is required in events.

Key Name	Full Name	Data Type	Meaning
act	deviceAction	String	Action mentioned in the event.
app	applicationProtocol	String	Application level protocol. Example values include: HTTP, HTTPS, SSHv2, Telnet, POP, IMAP, IMAPS.
in	bytesIn	Integer	Number of bytes transferred inbound. Inbound relative to the source to destination relationship, meaning that data was flowing from source to destination.
out	bytesOut	Integer	Number of bytes transferred outbound. Outbound relative to the source to destination relationship, meaning that data was flowing from destination to source.
dst	destinationAddress	IPv4 Address	Identifies destination that the event refers to in an IP network. The format is an IPv4 address, such as "192.168.10.1".
dhost	destinationHostName	FQDN	Identifies the destination that the event refers to in an IP network. The format is a fully qualified domain name (FQDN) associated with the destination node, such as "zurich.domain.com".
dmac	destinationMacAddress	String	Six colon-separated hexadecimal numbers, such as "00:0D:60:AF:1B:61"
dntdom	destinationNtDomain	String	The Windows domain name of the destination address.

Key Name	Full Name	Data Type	Meaning
dpt	destinationPort	Integer	The valid port numbers are between 0 and 65535.
dproc	destination ProcessName	String	The name of the process which is the event's destination, such as "telnetd" or "sshd".
duid	destinationUserId	String	Identifies the destination user by ID. For example, in Unix, the root user is generally associated with ID 0.
dpriv	destination UserPrivileges	String	The allowed values are: "Administrator", "User", and "Guest". This identifies the destination user's privileges. In Unix, for example, activity executed on the root user would be identified with destinationUserPrivileges of "Administrator".
duser	destinationUserName	String	Identifies the destination user by name. This is the user associated with the event's destination. E-mail addresses are also mapped into the UserName fields. The recipient is a candidate to put into destinationUserName.
end	endTime	TimeStamp	The time at which the activity related to the event ended. The format is MMM dd yyyy HH:mm:ss or milliseconds since epoch (January 1, 1970). An example would be reporting the end of a session.
fname	fileName	String	Name of the file.
fsize	fileSize	Integer	Size of the file.
msg	message	String	An arbitrary message giving more details about the event. Multi-line entries can be produced by using '\n' as the newline separator.
rt	receiptTime	TimeStamp	The time at which the event related to the activity was received. The format is MMM dd yyyy HH:mm:ss or milliseconds since epoch (January 1, 1970).
request	requestURL	String	In the case of an HTTP request, this field contains the URL accessed. The URL should contain the protocol as well, such as "http://www.security.com".
src	sourceAddress	IPv4 Address	Identifies the source that the event refers to in an IP network. The format is an IPv4 address, such as "192.168.10.1".

Key Name	Full Name	Data Type	Meaning
shost	sourceHostName	FQDN	Identifies the source that the event refers to in an IP network. The format is a fully qualified domain name (FQDN) associated with the source node, such as "zurich.domain.com".
smac	sourceMacAddress	String	Six colon-separated hexadecimal numbers, such as "00:0D:60:AF:1B:61"
sntdom	sourceNtDomain	String	The Windows domain name of the source address.
spt	sourcePort	Integer	The valid port numbers are between 0 and 65535.
suid	sourceUserId	String	Identifies the source user by ID. This is the user associated with the source of the event. For example, in Unix, the root user is generally associated with ID 0.
spriv	sourceUserPrivileges	String	<p>The allowed values are: "Administrator", "User", and "Guest". This identifies the source user's privileges. In Unix, for example, activity executed on the root user would be identified with sourceUserPrivileges of "Administrator".</p> <p>This is an idealized and simplified view of privileges and can be extended in the future.</p>
suser	sourceUserName	String	Identifies the source user by name. E-mail addresses are also mapped into the UserName fields. The sender is a candidate to put into sourceUserName.
start	startTime	TimeStamp	The time when the activity the event referred to started. The format is MMM dd yyyy HH:mm:ss or milliseconds since epoch (January 1, 1970).
proto	transportProtocol	String	Identifies the Layer-4 protocol used. The possible values are protocol names such as TCP or UDP.

Appendix B

Regular Expressions

Regular String Search expressions (Perl Regex) are used to compose Logger filters. The following describes the syntax of regular expressions in Perl.

Regex Overview

A regular expression is a string of characters which tells the searcher which string (or strings) you are looking for. The following explains the format of regular expressions in detail. If you are familiar with Perl, you already know the syntax. If you are familiar with Unix, you should know that there are subtle differences between Perl's regular expressions and Unix' regular expressions.

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.+ze`

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 numbers separated by a comma (,) and surrounded by braces ({}). It is like the star metacharacter, except the length of the string it matches must be within the minimum and maximum length specified by the two numbers in braces. Thus,

```
ab{3,5}c
```

will match "abbbc", "abbbbc" or "abbbbbc". No other string is matched. Likewise,

```
.{3,5}pentane
```

will match "cyclopentane", "isopentane" or "neopentane", but not "n-pentane", since "n-" is only two characters long.

The alternative metacharacter is represented by a vertical bar (|). It indicates an either/or behavior by separating two or more possible choices. For example:

```
isopentane|cyclopentane
```

will match any event containing the strings "isopentane" or "cyclopentane" or both. However, it will not match "pentane" or "n-pentane" or "neopentane." The last metacharacter is the **brackets** metacharacter. The bracket metacharacter matches one occurrence of any character inside the brackets ([]). For example,

```
\s[cmt]an\s
```

will match "can", "man" and "tan", but not "ban", "fan" or "pan". Similarly,

```
2,[23]-dimethylbutane
```

will match "2,2-dimethylbutane" or "2,3-dimethylbutane", but not "2,4-dimethylbutane", "2,23-dimethylbutane" or "2,-dimethylbutane". Ranges of characters can be used by using the dash (-) within the brackets. For example,

```
a[a-d]z
```

will match "aaz", "abz", "acz" or "adz", and nothing else. Likewise,

```
textfile0[3-5]
```

will match "textfile03", "textfile04", or "textfile05" and nothing else.

If you wish to include a dash within brackets as one of the characters to match, instead of to denote a range, put the dash immediately before the right bracket. Thus:

```
a[1234-]z
```

and

```
a[1-4-]z
```

both do the same thing. They both match "a1z", "a2z", "a3z", "a4z" or "a-z", and nothing else.

The bracket metacharacter can also be inverted by placing a caret (^) immediately after the left bracket. Thus,

```
textfile0[^02468]
```

matches any ten-character string starting with "textfile0" and ending with anything except an even number. Inversion and ranges can be combined, so that

```
\W[^f-h]ood\w
```

matches any four letter word ending in "ood" **except** for "food", "good" or "hood". (Thus "mood" and "wood" would both be matched.)

Note that within brackets, ordinary quoting rules do not apply and other metacharacters are not available. The only characters that can be quoted in brackets are "[", "]", and "\". Thus,

```
[\\[\]]abc
```

matches any four letter string ending with "abc" and starting with "[", "]", or "\".

Forbidden Characters

Because of the way the searcher works, the following metacharacters should not be used, even though they are valid Perl metacharacters. They are:

- \$ (allowed within brackets)
- \n
- \r
- \t
- \f
- \b
- () (allowed within brackets. Note that if you wish to search for parentheses within text outside of brackets, you should quote the parentheses.)
- \1, \2 ... \9
- \B
- :
- !

Things To Remember

Here are some other things you should know about regular expressions.

Because regular expressions can be complex, it can be more work mastering a search than just sifting through a long list of matches (unless you're already familiar with regular expressions).

The search is case insensitive; thus

```
mopac
```

and

```
Mopac
```

and

```
MOPAC
```

all search for the same set of strings. Each will match "mopac", "MOPAC", "Mopac", "mopaC", "MoPaC", "mOpAc" and so forth. Thus you need not worry about capitalization. (Note, however, that metacharacters must still have the proper case. This is especially important for metacharacters whose case determines whether their meaning is reversed or not.)

Outside of the brackets metacharacter, you must quote parentheses, brackets and braces to get the searcher to take them literally.

Appendix C

Restoring Factory Settings

ArcSight Logger can be restored to its original factory settings using built-in Acronis True Image software.



Restoring Logger to factory settings will irrevocably delete all event data and configuration settings.



The screens shown here are examples only. Your Logger partitions might vary, and the overall capacity might be different.

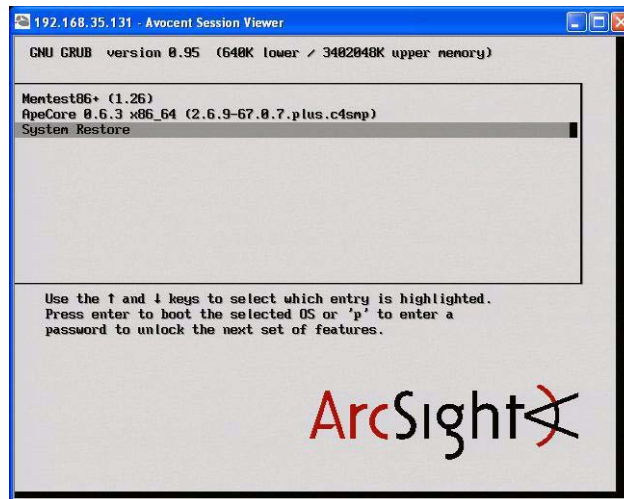
To restore factory settings

To restore Logger to its original factory settings, perform these steps:

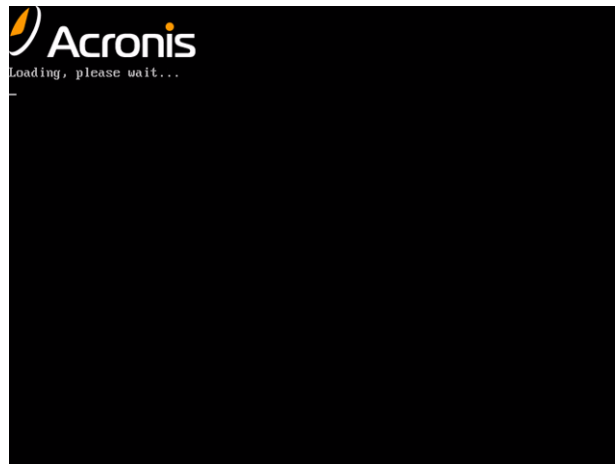
- 1 Attach a keyboard, monitor, and mouse directly to the ArcSight Logger appliance.
- 2 Reboot Logger from the web interface by clicking the **System Admin** tab, the **System Reboot** command in the sub-menu, and the **Start Reboot Now** button.
- 3 Once the following screen is displayed, press any key on your keyboard. This screen is displayed for a very short time, therefore, make sure you press a key on your keyboard quickly; otherwise, the Logger will continue to boot normally.



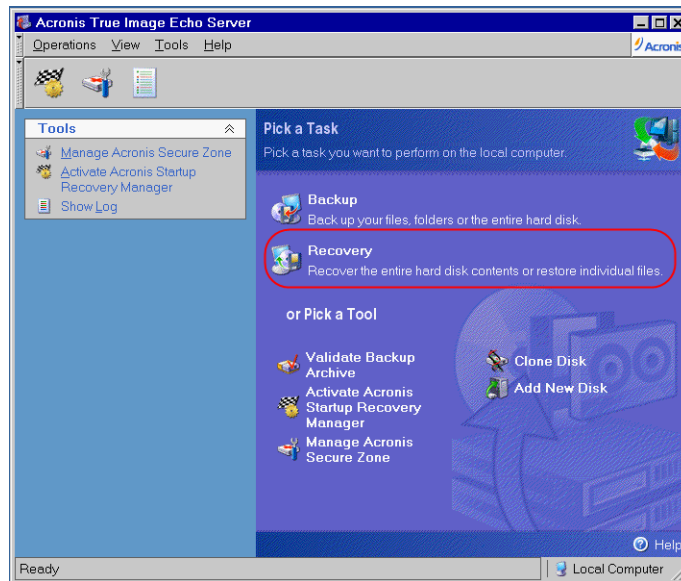
- 4 A screen similar to the following is displayed on the attached monitor. Use the mouse or arrow keys to select System Restore and press the **Enter** key on your keyboard.



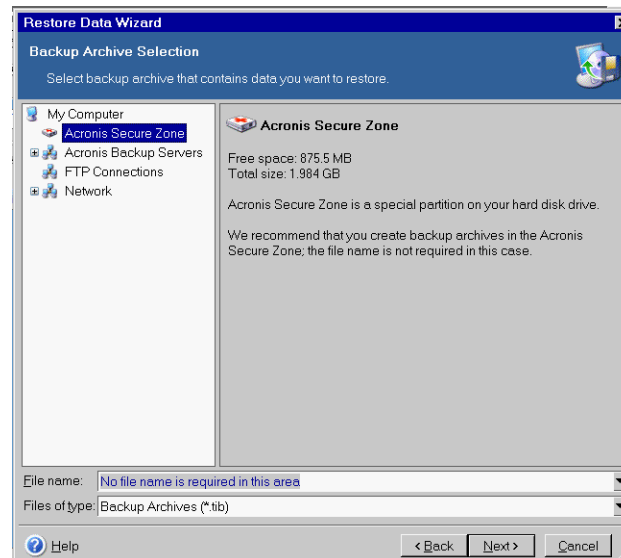
- 5 Click **Acronis True Image Server** to continue.



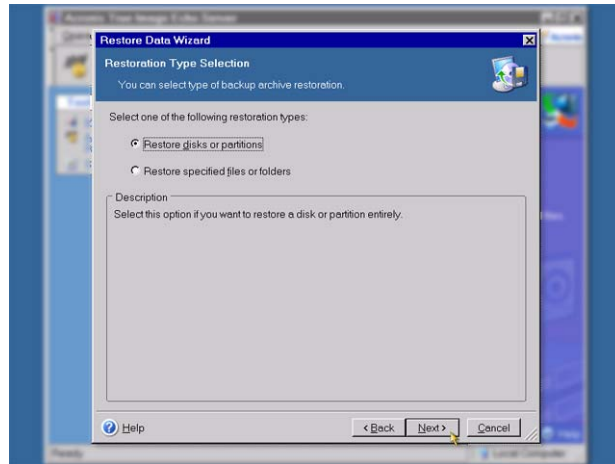
- 6 On the Pick a Task list, as shown in the following figure, choose **Recovery**. On the next page (Welcome to the Restore Data Wizard), click **Next** to continue.



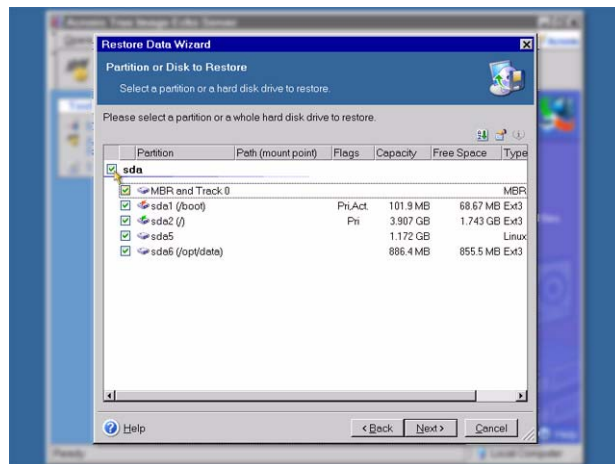
- 7 Select the Acronis Secure Zone, as shown in the following figure, and click **Next**. You will have a chance to review the choices you make on this page and the wizard pages that follow before initiating the restore process.



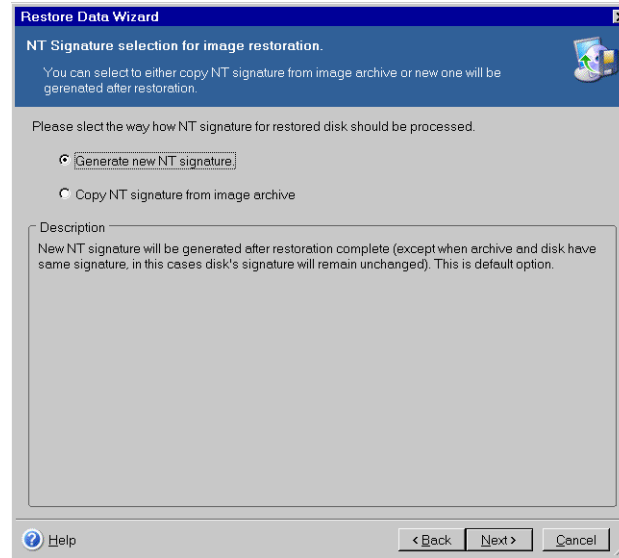
- 8 Select **Restore disks or partitions** and click **Next**. Only choose other options if specifically directed to do so by ArcSight Customer Support.



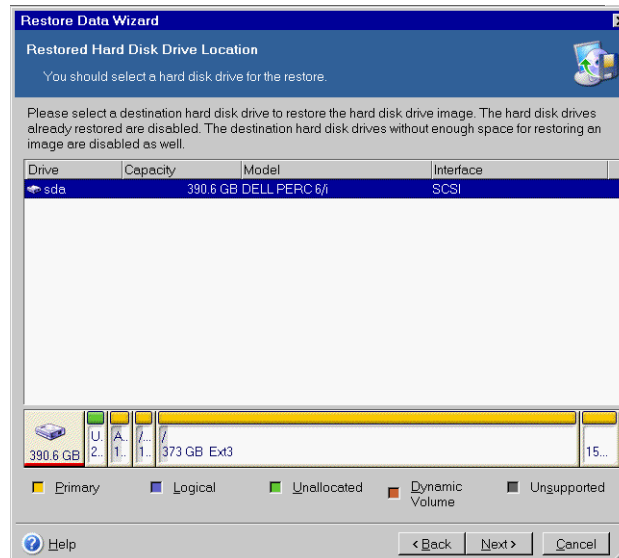
- 9 Select the entire drive, labeled 'sda' in the following figure. Click **Next** to continue.



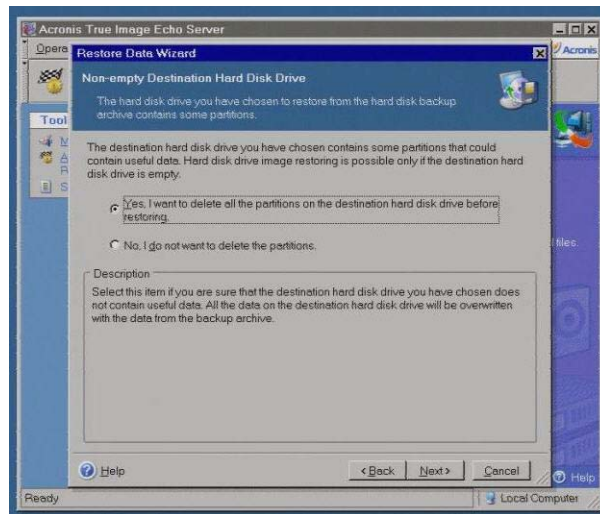
- 10 Select the way in which the NT signature for the restored disk should be processed and click **Next**.



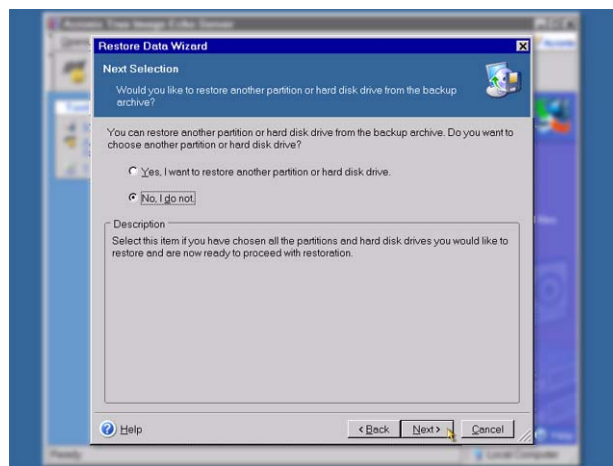
- 11 Choose the drive to restore ('sda') and click **Next**.



- 12 Select, "Yes, I want to delete all partitions on the destination hard disk drive before restoring", as shown in the following figure.

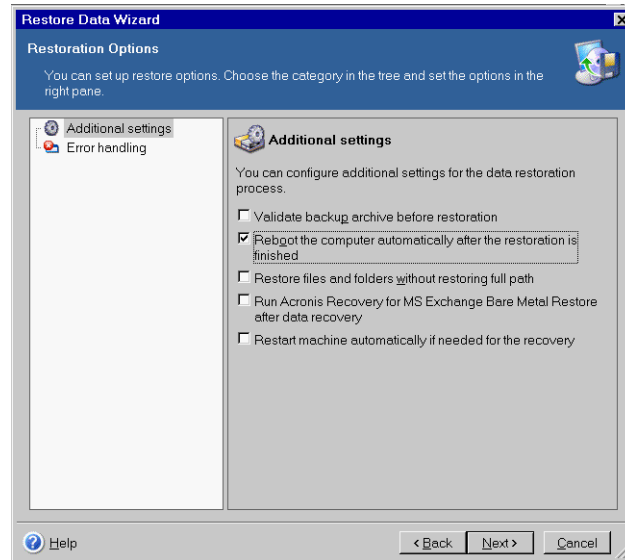


- 13 Because there are no other partitions or disks to restore, choose "No, I do not," on the Next Selection page of the wizard. Click **Next**.

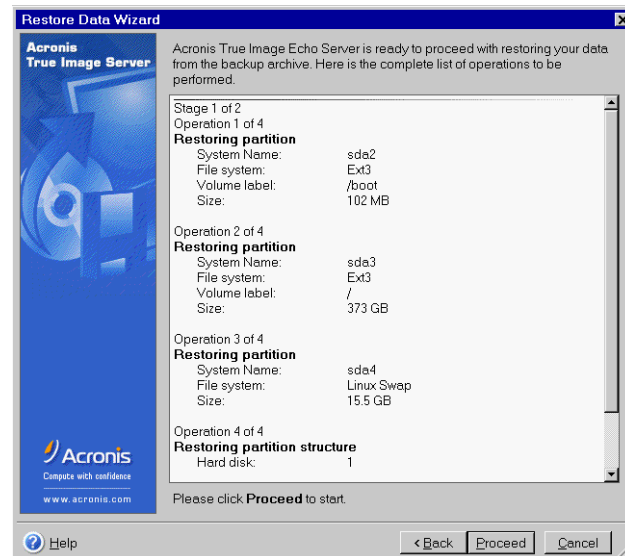


- 14 Validating the archive before restoring is optional. Check the box to validate the archive or leave it unchecked to skip this step. Check the box labeled "Reboot the

computer automatically after the restoration is finished” to automatically reboot. Click **Next**.



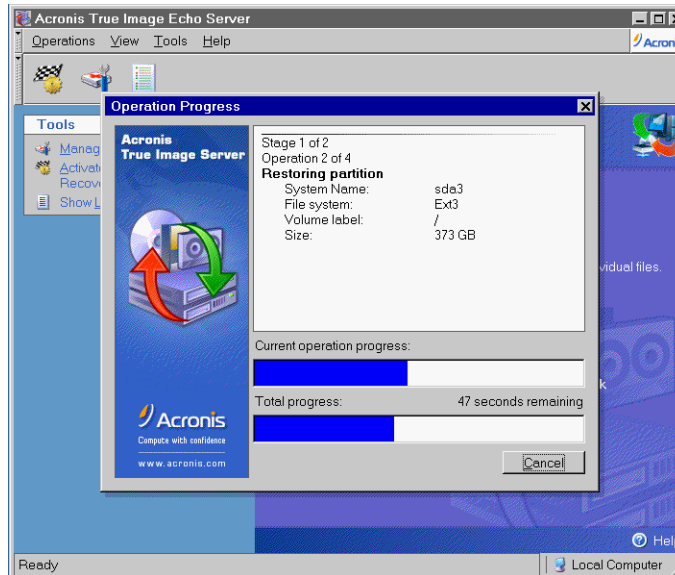
- 15** Review the checklist of operations to be performed, as shown in the following figure, and click **Proceed** to begin the restore process, or click **Back** to revisit previous wizard pages.



Caution

Do not interrupt or power-down Logger during the restore process. Interrupting the restore process may force the system into a state from which it cannot be recovered.

- 16** The progress bars (shown in the following figure) display the status of the current and total operations. When the restoration is complete, an alert is displayed that says "Data was successfully restored." Click **OK**.



If you specified automatic reboot in Step 13, Logger will reboot when the restore is complete. Otherwise, reboot manually.

Appendix D

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 [“Audit Forwarding” on page 230](#)) to forward the events. For a detailed understanding of the format of CEF events, see [Appendix A, Common Event Format, on page 259](#).

The following events are logged and available for Audit Forwarding to ArcSight ESM.

[“Platform Events” on page 282](#)

[“Logger Application Events” on page 284](#)

Types of Audit Events

Two types of audit events are generated on Logger:

- Platform Events—related to the Logger hardware/system
- Logger Application Events—related to Logger functions and configuration changes on it

In addition to these events, a Logger appliance that has an ArcSight Connector Appliance installed on it generates Connector Appliance audit events. For a list of Connector Appliance audit events, see the *Administrator's Guide for Connector Appliance version 5.5*.



Note

Platform audit events are not stored on the Logger appliance; therefore, you cannot search for them using the search facility available on the appliance. However, you can search on Alerts that are configured for these audit events.

Logger application events are searchable as usual.

Information in an Audit Event

A Logger audit event (in CEF format) contains information about the following prefix fields:

- Device Event Class ID
- Device Severity
- Name
- 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.

Device Event Class ID	Severity	Name	Device Event Category (cat)
platform: 200	7	Failed password change	/Platform/Authentication/Failure/Password
platform: 201	7	Failed login attempt	/Platform/Authentication/Failure/Login
platform: 202	5	Password changed	/Platform/Authentication/Password
platform: 205	5	Access enabled for support personnel	/Platform/Support/Enable
platform: 206	1	Access disabled for support personnel	/Platform/Support/Disable
platform: 210	3	Global login settings modified	/Platform/Configuration/Global/Login
platform: 211	3	Password policy modified	/Platform/Configuration/Global/Policy
platform: 212	5	Authentication settings modified	/Platform/Configuration/Global/RADIUS
platform: 213	7	Audit forwarding modified	/Platform/Configuration/Global/AuditEvents
platform: 220	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: 225	7	Uploaded file damaged or corrupt	/Platform/Update/Failure/Upload
platform: 226	7	Uploaded package damaged or corrupt	/Platform/Update/Failure/Installation
platform: 227	5	Applied appliance update	/Platform/Update/Applied
platform: 228	5	Failed to install package	/Platform/Update/Failure/Installation
platform: 230	5	Successful login	/Platform/Authentication/Login

Device Event Class ID	Severity	Name	Device Event Category (cat)
platform:231	5	Successful login (RADIUS)	/Platform/Authentication/Login/RADIUS
platform:232	7	Failed login attempt (BADUSER)	/Platform/Authentication/Failure/BADUSER
platform:233	7	Failed login attempt (BADPASS)	/Platform/Authentication/Failure/BADPASS
platform:234	7	Failed login attempt (LOCKED)	/Platform/Authentication/Failure/LOCKED
platform:235	7	Failed login attempt (INTERNAL)	/Platform/Authentication/Failure/INTERNAL
platform:236	7	Failed login attempt (EBADAUTH)	/Platform/Authentication/Failure/EBADAUTH
platform:237	7	Failed login attempt (ETIMEOUT)	/Platform/Authentication/Failure/ETIMEOUT
platform:238	7	Failed login attempt (NOACCESS)	/Platform/Authentication/Failure/NOACCESS
platform:239	1	User logout	/Platform/Authentication/Logout
platform:240	3	Added user group	/Platform/Groups/Add
platform:241	3	Updated user group	/Platform/Groups/Update
platform:242	3	Removed all members from group	/Platform/Groups/Membership/Remove
platform:243	3	Modified user group membership	/Platform/Groups/Membership/Update
platform:244	3	Deleted user group	/Platform/Groups/Remove
platform:245	3	Added user	/Platform/Users/Add
platform:246	3	Updated user	/Platform/Users/Update
platform:247	3	Deleted user	/Platform/Users/Delete
platform:250	5	Added remote mount point	/Platform/Storage/NFS/Add
platform:251	5	Edited remote mount point	/Platform/Storage/NFS/Edit
platform:252	7	Failed to create remote mount point	/Platform/Storage/NFS/Failure
platform:253	5	Removed remote mount point	/Platform/Storage/NFS/Remove
platform:254	5	Destroyed SAN Logical Unit	/Platform/Storage/SAN/Destroy
platform:255	5	Attached SAN Logical Unit	/Platform/Storage/SAN/Attach
platform:256	7	Detached SAN Logical Unit	/Platform/Storage/SAN/Detach

Device Event Class ID	Severity	Name	Device Event Category (cat)
platform: 257	5	Removed SAN Logical Unit	/Platform/Storage/SAN/Remove
platform: 259	5	Reattached SAN Logical Units	/Platform/Storage/SAN/Reattach
platform: 260	5	Staticroute modified	/Platform/Configuration/Network/Route/Update
platform: 261	5	Staticroute deleted	/Platform/Configuration/Network/Rou/Remove
platform: 262	5	Appliance time modified	/Platform/Configuration/Time
platform: 263	5	Network settings modified	/Platform/Configuration/Network
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	SMTPsettingsmodified	/Platform/Configuration/Network/SMTP
platform: 268	5	Static route added	/Platform/Configuration/Network/Route/Add
platform: 270	9	Stopped process '<process>'	/Platform/Process/Control/Stop
platform: 271	7	Restarted process '<process>'	/Platform/Process/Control/Restart
platform: 272	5	Started process '<process>'	/Platform/Process/Control/Start
platform: 280	7	Appliance reboot initiated	/Appliance/State/Reboot/Initiate
platform: 281	3	Appliance reboot canceled	/Appliance/State/Reboot/Cancel
platform: 282	9	Appliance poweroff initiated	/Appliance/State/Shutdown

Logger Application Events

The following table lists the information contained in audit events related to various Logger functions and configuration changes on it. The Severity for all Logger application events is 2.

Device Event Class ID	Name	Device Event Category (cat)
Filters		
logger: 500	Filter added	/Logger/Resource/Filter/Configuration/Add
logger: 501	Filter deleted	/Logger/Resource/Filter/Configuration/Delete

Device Event Class ID	Name	Device Event Category (cat)
logger:502	Filter updated	/Logger/Resource/Filter/Configuration/Update
Devices		
logger:510	Device added	/Logger/Resource/Device/Configuration/Add
logger:511	Device deleted	/Logger/Resource/Device/Configuration/Delete
logger:512	Device updated	/Logger/Resource/Device/Configuration/Update
Groups		
logger:513	Group added	/Logger/Resource/Group/Configuration/Add
logger:514	Group deleted	/Logger/Resource/Group/Configuration/Delete
logger:515	Group updated	/Logger/Resource/Group/Configuration/Update
Archives		
logger:520	Archive added	/Logger/Resource/Archive/Configuration/Add
logger:521	Archive deleted	/Logger/Resource/Archive/Configuration/Delete
logger:523	Archive loaded	/Logger/Resource/Archive/Configuration/Load
logger:524	Archive unloaded	/Logger/Resource/Archive/Configuration/Unload
logger:525	Archive archived	/Logger/Resource/Archive/Configuration/Archive
logger:526	Event archive settings added	/Logger/Resource/Archive/Add
logger:527	Daily archive task settings updated	/Logger/Resource/Archive/Update
Storage Groups		
logger:530	Storage group added	/Logger/Resource/StorageGroup/Configuration/Add
logger:532	Storage group updated	/Logger/Resource/StorageGroup/Configuration/Update
Storage Rule		
logger:533	Storage rule added	/Logger/Resource/StorageRule/Configuration/Add
logger:535	Storage rule updated	/Logger/Resource/StorageRule/Configuration/Update
Storage Volume		
logger:536	Storage volume added	/Logger/Resource/StorageVolume/Configuration/Add

Device Event Class ID	Name	Device Event Category (cat)
Saved Search		
logger:540	Saved search added	/Logger/Resource/SavedSearch/Configuration/Add
logger:541	Saved search deleted	/Logger/Resource/SavedSearch/Configuration/Delete
logger:542	Saved search updated	/Logger/Resource/SavedSearch/Configuration/Update
Peer Loggers		
logger:550	Peer Logger added	/Logger/Resource/PeerLogger/Configuration/Add
logger:551	Peer Logger deleted	/Logger/Resource/PeerLogger/Configuration/Delete
logger:570	Peer Logger authorization added	/Logger/Resource/PeerLogger/Authorizations/Configuration/Add
logger:571	Peer Logger authorization deleted	/Logger/Resource/PeerLogger/Authorizations/Configuration/Delete
Event Input/Output		
logger:600	Receiver added	/Logger/Component/Receiver/Configuration/Add
logger:601	Receiver deleted	/Logger/Component/Receiver/Configuration/Delete
logger:602	Receiver updated	/Logger/Component/Receiver/Configuration/Update
logger:603	Receiver enabled	/Logger/Component/Receiver/Configuration/Enable
logger:604	Receiver disabled	/Logger/Component/Receiver/Configuration/Disable
logger:605	Forwarder added	/Logger/Component/Forwarder/Configuration/Add
logger:606	Forwarder deleted	/Logger/Component/Forwarder/Configuration/Delete
logger:607	Forwarder updated	/Logger/Component/Forwarder/Configuration/Update
logger:608	Forwarder enabled	/Logger/Component/Forwarder/Configuration/Enable
logger:609	Forwarder disabled	/Logger/Component/Forwarder/Configuration/Disable
logger:663	Forwarder paused	/Logger/Component/Forwarder/Configuration/Pause
logger:664	Forwarder resumed	/Logger/Component/Forwarder/Configuration/Resume
logger:640	ESM destination added	/Logger/Component/EsmDestination/Configuration/Add

Device Event Class ID	Name	Device Event Category (cat)
logger: 641	ESM destination deleted	/Logger/Component/EsmDestination/Configuration/Delete
logger: 643	Certificate added	/Logger/Component/Certificate/Configuration/Add
logger: 650	Certificate deleted	/Logger/Component/Certificate/Configuration/Delete
logger: 651	Certificate updated	/Logger/Component/Certificate/Configuration/Update
logger: 644	SNMP destination added	/Logger/Component/SnmpDestination/Configuration/Add
logger: 645	SNMP destination deleted	/Logger/Component/SnmpDestination/Configuration/Delete
logger: 647	Syslog destination added	/Logger/Resource/SyslogDestination/Configuration/Add
logger: 648	Syslog destination deleted	/Logger/Component/SyslogDestination/Configuration/Delete
logger: 649	Syslog destination updated	/Logger/Component/SyslogDestination/Configuration/Update
Alerts		
logger: 610	Alert added	/Logger/Component/Alert/Configuration/Add
logger: 611	Alert deleted	/Logger/Component/Alert/Configuration/Delete
logger: 612	Alert updated	/Logger/Component/Alert/Configuration/Update
logger: 613	Alert enabled	/Logger/Component/Alert/Configuration/Enable
logger: 614	Alert disabled	/Logger/Component/Alert/Configuration/Disable
logger: 615	Alert sent	/Logger/Component/Alert/Configuration/Sent
Configuration Backup		
logger: 660	Configuration backup updated	/Logger/Component/ConfigBackup/Update
logger: 661	Configuration backup enabled	/Logger/Component/ConfigBackup/Enable
logger: 662	Configuration backup disabled	/Logger/Component/ConfigBackup/Disable
Search		
logger: 680	Search indices added	/Logger/Search/Index/Update
logger: 690	Search options updated	/Logger/Search/Options/Update

Device Event Class ID	Name	Device Event Category (cat)
Maintenance Mode		
logger: 700	Maintenance mode entered	/Logger/Server/MaintenanceMode/Enter

Appendix E

Event Field Name Mappings

The nomenclature for field names depends on the function area of the Logger. The following table provides the mapping between 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_deviceSeverity	deviceSeverity	deviceSeverity	Device Severity
arc_deviceZoneURI	deviceZoneURI	deviceZoneURI	Device Zone URI
arc_endTime	endTime	rt	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
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

Database Name	Search Results	Actual Events (CEF Key)	Reports
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
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

Database Name	Search Results	Actual Events (CEF Key)	Reports
arc_sourceZoneResource	sourceZoneResource	sourceZoneResource	Source Zone Resource

Connector Appliance Documentation

This information is applicable only to Logger platforms with integrated Connector Appliance.

Connector Appliance documentation is available as follows:

- Through the Help icon (?) on any user interface page, when you are in the Connector Appliance context. When you click this icon, a PDF of the Connector Appliance Administrator's Guide is displayed. **All information in this guide except system administration is applicable to your product.**
- Through the ArcSight Customer Support site at <https://support.arcsight.com>.

Logger Search From An ESM Console

If you have ArcSight Logger and ArcSight ESM deployed in your network infrastructure, you can perform a Logger search operation directly from your ESM Console.

Understanding the Integrated Search Functionality

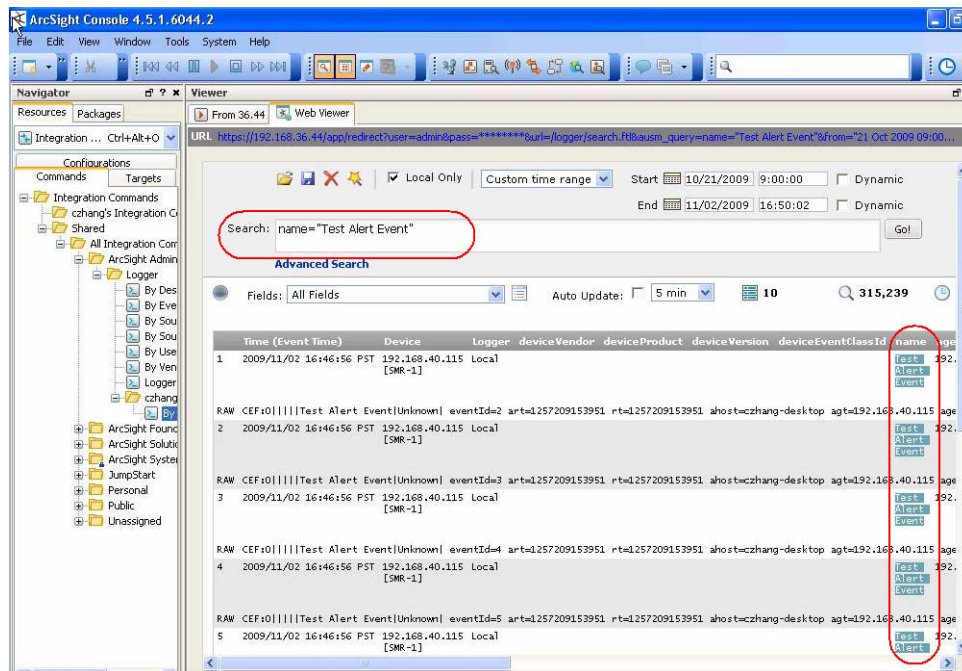
There are two ways to perform a search operation on Logger from an ESM Console:

- Search—a regular search operation in which you can specify search options
- Quick search—a search operation based on field and value you select in an ESM Console active channel; you are not prompted for any search options.

To run a Logger search, you right click on an event in an active channel of the ESM Console to display a menu. You select the search method—Logger Search or Logger Quick Search—from the menu.

If you select Logger Search, you need to select search options such as Event Name, Destination, Source, and so on, and the Logger appliances on which the search should be run (if there are multiple Logger appliances). If you select Logger Quick Search, you are not prompted to specify any search options because the search is run on the field you had clicked on.

The search results are displayed in the ESM Console, as shown in the following figure:



Prerequisites

The Logger on which search will be performed must be running Logger v4.0. The ESM Manager from which search will be initiated must be running ESM v4.5 SP1 Patch2.

Setup and Configuration

ESM

Follow these instructions to set up and configure ArcSight ESM Manager to run integrated search operations:

- 1 Ensure that the ESM Manager is running v4.5 SP1 Patch2.
- 2 Follow instructions in the *ArcSight ESM v4.5 SP1 Patch2 Release Notes* to set up ESM Console for integrated searches on Logger.

The ESM release notes are available from the ArcSight Customer Support web site at <https://support.arcsight.com>.

Logger

Make sure:

- 1 Your Logger is running v4.0.
- 2 A Logger user name that you specified when creating an integration parameter on ESM Console (Step 2 of ESM in "Setup and Configuration" on page 298) 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

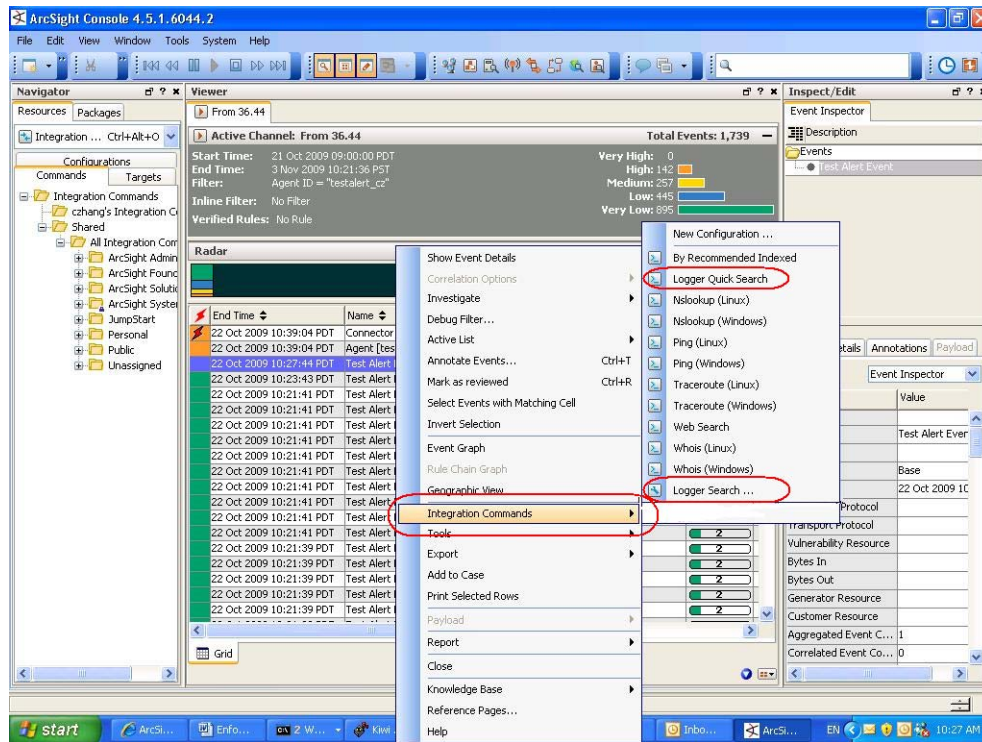
- A field-based search query is used to perform search on the Logger.
- A search operation only from an active channel of an ESM Console is supported; search operation from other ESM resources is not supported.
- Multiple search options (see [“Supported Search Options” on page 299](#)) cannot be specified for one search operation. That is, you cannot select by Event Name and By Destination for one search operation.
- You can run the search operation on only one Logger at a time. That is, you cannot select multiple Loggers from the menu list on ESM Console.

Additionally, even if a Logger peers with other Loggers, integrated search runs only on the local Logger that you select from the menu list on ESM Console.

Searching on Logger From ESM Console

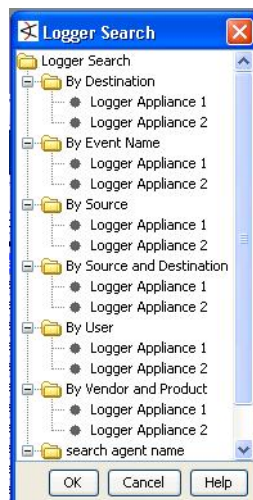
- To run a **Quick Search** on Logger (as described in [“Understanding the Integrated Search Functionality” on page 297](#)):
 - a Right click on the event field in an active channel of the ESM Console.

- b From the menu list, select **Integration Commands > Logger Quick Search**, as shown in the following figure.



OR

- To run a regular **Search** (in which you specify search options):
 - a Right click on any field of an event in an active channel of the ESM Console.
 - b From the menu list, select **Integration Commands > Logger Search > Select Search Options**, as shown in the following figure.



- c Click **OK** to run the search or **Cancel** to quit.

The search results are displayed in the ESM Console Web Viewer.

Index

A

- accounts, user. See User.
- Acronis True Image Server 273
- agents. See SmartConnectors
- Alerts
 - about 187
 - adding 189
 - disable 190
 - enable 190
 - remove 190
- apache status 228
- archive, scheduled 167
- ArcSight ESM 2, 6, 7, 26, 28, 179, 186
- ArcSight Logger Streaming SmartConnector 179
- ArcSight Manager 7, 26, 28, 179, 184
- ArcSight SmartConnectors 7, 27
- Audit forwarding 281
- Audit Log 230, 281
- authentication, RADIUS 249
- auto update feature 71
- automatic timeout 247

B

- backup, configuration 206
- browser requirements 31

C

- CAC support 13, 240
- CACERTS for ESM Destination 185
- canonical equality check 33
- case-sensitive search 33
- CEF 69, 259
- CEF event filters 69
- certificate revocation list 242
- Certificate Signing Request 239
- Certificate, installation 240
- character encoding 176
- CLI 17
 - command table 18
- command line interface (CLI) 17
- Common Event Format (CEF) 69, 259
- Common Extension Dictionary 261
- Configuration Backup 206
- configuration monitoring, reports for 79
- Configuration tab 163, 219
- Connector Forwarder 179
- connectors. See SmartConnectors
- constraints, search 40
- CSR

- generating a certificate signing request 239
- current time, changing 224

D

- dashboard
 - reports 82
 - reports, preference for display 94
- date/time format 178
- default gateway 222
- Default Storage Group 2, 24
- deploying
 - report package 159
- Device 163, 164
 - delete 164
 - edit 164
 - pre-defining 164
- Device Group 165
 - creating 165
 - deleting 166
 - editing 165
- device group
 - maximum number 165
- devices
 - maximum number 164
- DNS Settings 220
 - changing 220, 221
- dynamic search 44

E

- e-mailing
 - reports 104
- encoding 176
- Error Log 230
- ESM (ArcSight Enterprise Security Manager) 2, 6, 7, 26, 28, 179
- ESM Destination 183
 - creating 184
 - deleting 185
 - updating CACERTS 185
- ESM SmartConnector status 228
- etc/hosts.txt 221
- Event Archive 166
 - adding 166
 - deleting 167
 - loading 167
 - settings 168
 - unloading 167
- event archive, scheduled 167
- Event Input/Output 173
- event storage, remote 172

- events
 - search 40
- export
 - search results 62

F

- factory settings, restoring 273
- field query
 - indexing fields 64
- field set, search 40
- fields, indexing 64
- File events to ESM 186
- Filter 194
 - copying 195
 - creating 194
 - deleting 195
 - editing 195
- Filter, Report Category 162
- filter, search 40
- filters, system 68
- find, events 40
- FIPS 140-2
 - enabling on Connector Appliance 12, 242
- Firefox (web browser) 31
- Forwarder 179
 - creating 180
 - deleting 182
 - editing 182
- Forwarder status 228
- forwarding file events to ESM 186
- function tabs 32

G

- gateway, default 222
- gauge range 33
- gauges 32
- gid 233
- Global Settings 247

H

- health, system 69
- Help
 - how to use Console Online Help xi
- help 32
- Hosts file 221

I

- i18n options 33
- indexing fields 64
- initialization 16
- insp status 228
- interface homing 222
- Internal Storage Group 2, 169
- Internet Explorer (web browser) 31
- intrusion monitoring, reports for 78
- IP addresses
 - assigning 16
 - changing 222

L

- list selection 61

- localhost 221
- Logger
 - rebooting 220
- login 31
- Login Settings 247
 - changing 247
- logout 32, 33
- logout, automatic 33
- Logs 230
- logs, internal 215
 - retrieving 215

M

- maintenance mode 209
- Manager 7, 26, 28, 179, 184
- Monitor tab 33
- multi-homing 222

N

- navigation 32
- Network Settings 221
 - changing 222
- network speed 222
- NTP Server 225
- NTP setting 225, 226

O

- Online Help
 - see *Help*
- online help 32
- options 32, 33

P

- package contents 15
- parameter value groups, in reports 152
- parameters
 - in report queries 145
 - quick run report 99
 - run reports 101
- Password policy
 - changing 248
- Password, changing 257, 258
- PCI Storage Group 2, 169
- Peer Logger 202
 - adding 204
 - deleting 206
- peer Logger, searching 60
- postgresql status 228
- predefined filters 68, 69
- Process Status 228

Q

- queries
 - in reports 122
- query
 - events 40
- query controls 32

R

- RADIUS authentication 249

- RAID controller status 228, 237
- range, gauge 33
- rebooting Logger 220
- Receiver 173
 - creating 174
 - deleting 175
 - editing 175
 - types 173, 175
- Receiver status 228
- regular expressions (regex)
 - predefined 69
 - tutorial 265
- Remote Authentication Dial-In User Service (RADIUS) 249
- remote event storage 172
- remote file system mount
 - adding 234, 236
 - deleting 233, 235
 - editing 233, 234, 236
- Report Category Filter 162
- reports 198
 - access rights 122
 - administration 160
 - categories 76
 - configuration monitoring 79
 - creating new 107
 - dashboard 82
 - delivery options 104
 - designing 106
 - editing 120
 - e-mailing 104
 - exporting 105
 - file formats 102
 - foundation 77
 - groups 76
 - intrusion monitoring 78
 - navigating to 75
 - parameter value groups 152
 - PCI solution add-on 79
 - publishing 103
 - query parameters 145
 - quick run parameters 99
 - remove scheduled 156
 - run parameters 101
 - running 95
 - SANS Top 5 77
 - saving 105
 - scheduling 155
 - solution add-ons 79
 - template styles 154
 - user-created 80
 - viewing published 106
 - viewing, editing schedules 155
- reset to factory settings 273
- restore to factory settings 273
- restoring a SAN 237
- Retrieve Logs 215
 - filter 67
 - search 67
- Saved Search 197
 - adding 197
 - deleting 198
 - editing 197
- Saved Search Files 200
- Saved Search Job 198
 - adding 198
 - deleting 200
 - editing 199
- scheduled event archive 167
- Scheduled Task 193
 - currently running 194
 - finished 194
- scheduling
 - export of search results 60
 - reports 155
- SCP file receiver 175
- search
 - constraints 40
 - events 40
 - exporting results 62
 - field set 40
 - filter 40, 67
 - peer Loggers 60
 - query, defining a 40
 - results, scheduling export of 60
 - saved 67
 - system filters 68
 - time range 40
- Search Group Filter
 - associating with user group 196
 - report category filter 162
- Search Group Filters 195
- Search Results tab 61
- selection in lists 61
- servers status 228
- SFTP file receiver 175
- SmartConnectors 7, 27
- SmartMessage 173
- solutions
 - reports 79
- speed, network 222
- SSL
 - Certificate Signing Request 239
- SSL Settings 238
- Static Route 227
 - adding 227
- statistics 32
- status
 - 3Ware RAID Controller 228, 237
 - process 228
- Storage 168
- Storage Area Network 235
- Storage Group 168
 - adding 169
 - Default 24
 - editing 169
- Storage Group, default 2, 24
- Storage Group, internal 2, 169
- Storage Group, PCI 2, 169
- Storage Rule 24, 170
 - adding 171
 - deleting 172

S

- safety precautions 16
- SAN Storage 235
- SAN, restore 237
- SANS Top 5, reports for 77
- saved

- editing 171
- Storage Settings 172
- Storage Volume 172
 - settings 172
- streaming SmartConnector 179
- subnet mask 222
- System Admin tab 219
- system filters 68, 69
- system health, monitoring 69
- System Information 231
- System Reboot 219
- System Update 227

T

- template styles
 - for reports 154
- time configuration 225, 226
- time range, dynamic 44
- time range, search 40
- Time Settings 223
 - changing 225, 226
- time, changing 224
- timeout, automatic 247
- timezone 225

U

- uid 233
- Unicode options 33

- US-ASCII encoding 176
- User 251, 256
 - changing password 258
 - creating 256
 - deleting 257
 - editing 257
- User Group 251
 - associating with Search Group Filter 196
 - creating 255
 - deleting 255
 - editing 255
- user interface 32
 - Search Results tab 61
- User password, changing 258
- UTF-8 encoding 176

V

- version, component 227

W

- Web
 - viewing Online Help in Web browser xi
- web browser requirements 31
- web status 228
- What's New 8
- widgets
 - in report dashboards 91