

HP Operations Log Intelligence

for the Linux operating system

Software Version: 1.0.0

Administrator's Guide

Document Release Date: January 2014

Software Release Date: January 2014



Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notices

© Copyright 2005 - 2014 Hewlett-Packard Development Company, L.P.

Documentation Updates

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

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

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

<http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign-in. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

Visit the HP Software Support web site at:

<http://www.hp.com/go/hpsoftwaresupport>

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Contents

Chapter 1 Overview	15
Introduction	15
OLI Features	16
Storage Configuration	16
Receiver Configuration	17
Analyzing Events	18
Grouping Events	19
Exporting	19
Forwarder Configuration	20
User Management	20
Other Setup and Maintenance	20
Deployment Scenarios	21
Chapter 2 Installation and Initialization	23
Deployment Planning	23
Storage Strategy	23
Retention Policy	23
Initial Configuration	24
Storage Volume	25
Storage Groups	25
Indexed Fields and Full-text Indexing	25
Receivers	25
Licensing	26
Installing an OLI	26
Supported Platforms	27
Downloading the OLI	27
Acquiring a License for OLI	27
How Licensing Works on the OLI	27

Prerequisites for Installation	29
Installation Modes	30
Installation Steps	31
Using the GUI Mode to Install OLI	31
Using the Console Mode to Install OLI	33
Using the Silent Mode to Install OLI	34
Starting and Stopping the OLI	36
Uninstalling the OLI	37
Connecting to OLI for the first time	37
Configuring OLI	38
Receivers	38
Enabling the Pre-configured Receivers	39
Devices	40
Device Groups	40
Storage Rules	40
Using SmartConnectors to Collect Events	41
SmartMessage	41
Downloading SmartConnectors	41
Configuring a SmartConnector to Send Events to OLI	42
Configuring SmartConnectors to Send Events to Both OLI and an ArcSight Manager	42
Configuring SmartConnectors for Failover Destinations	43
Chapter 3 User Interface and Dashboards	45
Connecting to the OLI User Interface	45
Logging In	45
Navigating the User Interface	46
Help	47
Options	47
Logout	49
Summary	49
Dashboards	51
Creating and Managing Dashboards	56
Adding and Managing Panels in a Dashboard	58
The Default Monitor Dashboard	62
Platform	63

Network	64
Operations Log Intelligence	64
Receivers	65
Forwarders	65
Storage	65
Chapter 4 Searching and Analyzing Events	67
The Need to Search Events	67
The Process of Searching Events	68
Elements of a Search Query	68
Query Expression	69
Indexed Search	69
Search Operators	73
Time Range	73
Field Set	75
Constraints	79
Syntax Reference for Query Expression	81
Using the Search Builder Tool	86
Accessing Search Builder	86
Nested Conditions	89
Alternate Views for Query Building in Search Builder	89
Search Analyzer	90
Performance Optimizations for Indexed Fields in Search Queries	91
Regex Helper Tool	91
Search Helper	93
Autocomplete Search	94
Search History	95
Search Operator History	95
Examples	95
Usage	95
Suggested Next Operators	95
Help	96
Searching for Events on OLI	96
Advanced Search Options	97
Searching Peer OLIs (Distributed Search)	97
Tuning Search Performance	98

Understanding the Search Results Display	99
User-defined Fields in Search Results	101
Viewing Search Results using Field Sets	101
Using the Histogram	102
Multi-line Data Display	103
Auto Updating Search Results	104
Chart Drill Down	104
Understanding Field Summary	105
Exporting Search Results	108
Scheduling an Export Operation	111
Indexing	111
How indexing works	111
Full-text Indexing (Keyword Indexing)	112
Field-based Indexing	112
Guidelines for Field-based Indexing	114
Enabling Indexing	115
Adding Fields to Field-based Index	115
Saving Queries (Saved Filters and Searches)	116
Saving a Query	116
Using a Saved Filter or a Saved Search	118
System Filters/Predefined Filters	119
Using a System Filter	121
Alerts	122
Viewing Alerts	122
Receiving Alerts for Events	122
Base Event Fields	123
Go, Export, and Auto Update Options	123
Live Event Viewer	123
Chapter 5 Configuration	127
Devices	127
Devices	127
Device Groups	129
Event Archives	131
Guidelines for Archiving Events	132
Archiving Events	133

Scheduled Event Archive	134
Archive Storage Settings	135
Loading and Unloading Archives	136
Storage	136
Storage Groups	136
Storage Rules	138
Storage Volume	140
Event Input	140
Receivers	140
File Based Receivers	141
Working with Receivers	143
Source Types	155
Working with Source Types	156
Parsers	159
Using Parsers with Source Types	159
Using the Parse Command	160
Working with Parsers	160
Example: Creating an Extract Parser	163
Event Output	165
Forwarders	165
Alerts	172
Alert Triggers and Notifications	174
When are Alert events triggered?	174
Receiving Alert Notifications	174
Sending Notifications to E-mail Destinations	175
Sending Notifications to Syslog and SNMP Destinations	175
Configuring and Managing Real Time Alerts	176
Creating a Real Time Alert	176
Creating and Managing Saved Search Alerts	179
Creating a Saved Search Alert	179
Sending Notifications to SNMP Destinations	185
Sending Notifications to Syslog Destinations	186
Scheduled Tasks	187
Scheduled Tasks	187
Currently Running Tasks	188
Finished Tasks	188

Filters	189
Filters	189
Search Group Filters	191
Saved Searches	192
Saved Searches	192
Scheduled Saved Search	193
Saved Search Files	197
Search	197
Adding Search Indexes	197
Tuning Advanced Search Options	197
Viewing and Deleting Field Sets	200
Viewing Default Fields	201
Viewing Custom Fields	201
Running Search Tasks	202
Ending Currently Running Tasks	202
View and Add Parsers for Specific Log Types	203
Peer OLIs	203
Guidelines	204
Authorizing Peers	207
Configuration Backup and Restore	207
Running a Configuration Backup (Ad-hoc or Scheduled)	208
Restoring from a Configuration Backup	209
Editing Configuration Backup Settings	210
System Maintenance	210
Entering Maintenance Mode	211
Exiting Maintenance Mode	211
Checking Status of a Maintenance Operation	211
Database Defragmentation	211
Guidelines for Database Defragmentation	212
Defragmenting an OLI	212
Global Summary Persistence Defragmentation	216
Guidelines for Global Summary Persistence Defragmentation	216
Storage Volume Size Increase	217
Adding Storage Groups	219
Adding or Importing Schema Fields	221
Importing Schema Fields from Peers	223

License Information	227
Data Volume Restrictions	228
Retrieve Logs	229
Content Management	229
Importing Content	230
Importing Guidelines	230
Exporting Content	231
Exporting Guidelines	231
Chapter 6 System Admin	235
System	235
System Locale	235
Impact of Daylight Savings Time Change on Logger Operations	236
SMTP	236
License & Update	236
Updating the License File	236
Process Status	237
System Settings	238
Logs	238
Audit Logs	238
Security	239
SSL Server Certificate	239
Generating a Self-Signed Certificate	240
Generating a Certificate Signing Request (CSR)	242
Importing a Certificate	243
SSL Client Authentication	244
Configuring OLI to Support SSL Client Authentication	244
Uploading Trusted Certificates	245
Uploading a Certificate Revocation List	246
Enabling Client Certificate Authentication	247
FIPS 140-2	247
Things to be Aware of When Enabling FIPS Mode on OLI	248
Users/Groups	251
Authentication	251
Sessions	251
Local Password	252

Users Exempted From Password Expiration	254
External Authentication	255
Login Banner	260
User Management	261
Users	261
Groups	264
Change Password	266
Other System Administration Information	266
Monitoring System Health	266
System Health Events	267
Appendix A Search Operators	269
chart	269
Aggregation Functions	271
Multi-Series Charts	271
The span function	272
dedup	276
eval	277
extract	277
fields	279
head	280
keys	280
parse	281
rare	282
regex	283
rename	283
replace	284
rex	286
sort	288
tail	289
top	289
transaction	290
where	292

Appendix B Audit Events	293
Types of Audit Events	293
Information in an Audit Event	293
Platform Events	293
OLI Application Events	301
Appendix C Examples of System Health Events	317
Appendix D Using the Rex Operator	325
Syntax of the rex Operator	325
Understanding the rex Operator Syntax	325
Ways to Create a rex Expression	326
Creating a rex Expression Manually	327
Samples of rex Expressions	327
Index	331

Chapter 1: Overview

This chapter provides an overview of Operations Log Intelligence (OLI), with references to other parts of this document for more detail.

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

[“Introduction” on page 15](#)

[“OLI Features” on page 16](#)

[“Deployment Scenarios” on page 21](#)

Introduction

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

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

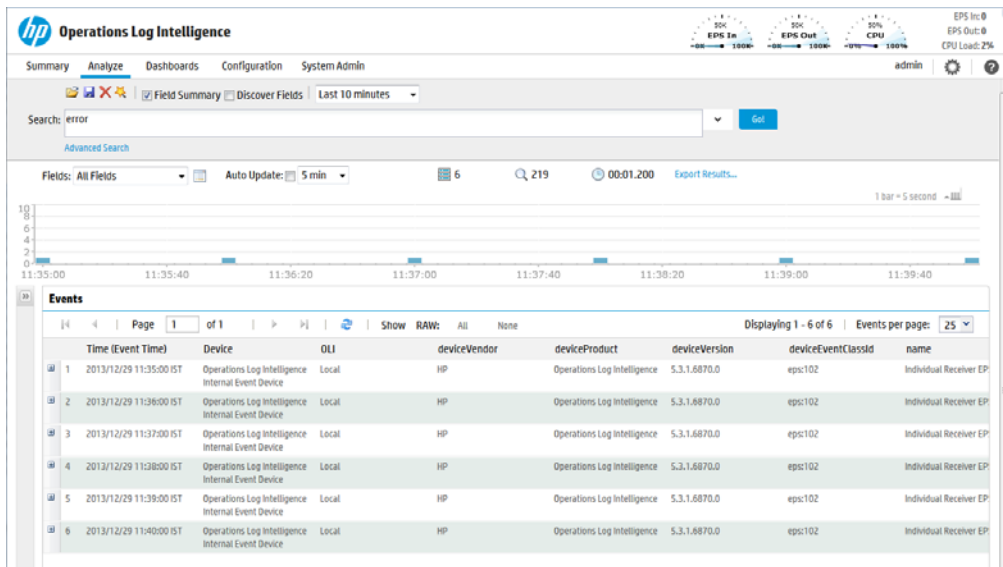


Figure 1-1 OLI web interface, Analyze tab

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

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

OLI is available as a software solution. The software solution enables you to install OLI on a supported platform of your choice.

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

OLI Features

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

Storage Configuration

For OLI, you need to have at least the minimum disk space described in the Release Notes to store events. The disk space needs to be on the partition where the `<install_dir>` directory exists. Specifically, most of this space should be available for the `<install_dir>/data/logger` directory. Using NFS as primary storage for events on the OLI is not recommended.

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

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

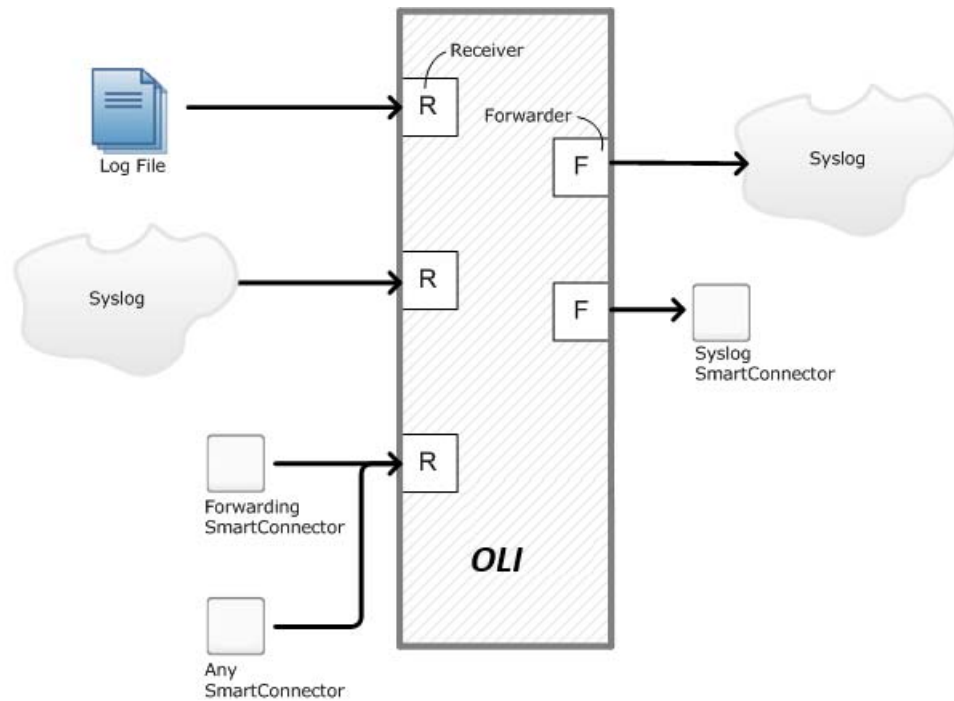
The storage volume, either external or local, can be divided into multiple storage groups, each with a separate retention policy. Two storage groups are created when OLI is first configured. New storage groups can be added later, and a storage group's size can be increased or decreased, and the retention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

- [“Deployment Planning” on page 23](#)
- [“Storage” on page 136](#)

Receiver Configuration

OLI 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 OLI can receive syslog and CEF messages using the more reliable Transmission Control Protocol (TCP) as well.

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



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

- [“Receivers” on page 140](#)
- [“Using SmartConnectors to Collect Events” on page 41](#)

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 terms in the event table. Queries can be based on plain English keywords (full-text search), predefined fields, or specified as regular expressions. OLI supports a flow-based search language that allows you to specify multiple search commands in a pipeline format.

By default, an OLI queries only its primary data store even if peer OLIs are configured. However, you can configure it to distribute a query across peer OLIs 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 OLI” on page 96](#)
- [“Saving Queries \(Saved Filters and Searches\)” on page 116](#)
- [“Filters” on page 189](#)

- [“Saved Searches” on page 192](#)
- [“View and Add Parsers for Specific Log Types” on page 203](#)

Grouping Events

The combination of a source IP address and an OLI 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 OLIs 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 127](#)
- [“Storage Rules” on page 138](#)
- [“Searching Peer OLIs \(Distributed Search\)” on page 97](#)

Exporting

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

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

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

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

- [“Exporting Search Results” on page 108](#)
- [“Scheduled Saved Search” on page 193](#)

Forwarder Configuration

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

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

For more information about Forwarders, refer to ["Forwarders" on page 165](#).

User Management

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

- ["Users/Groups" on page 251](#)
- ["Change Password" on page 266](#)
- ["Search Group Filters" on page 191](#)

Other Setup and Maintenance

OLI 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 OLI activity can be downloaded through the browser on demand, for debugging or other reasons. Other system information is available for viewing. The OLI can be rebooted by restarting OLI service and related processes.

Follow instructions in ["Starting and Stopping the OLI" on page 36](#) to start, stop, or restart OLI service on an OLI.

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

- ["Configuration Backup and Restore" on page 207](#)
- ["Retrieve Logs" on page 229](#)
- ["System Locale" on page 235](#)
- ["License & Update" on page 236](#)

Deployment Scenarios

Typically, OLI is deployed inside the perimeter firewall with a high degree of physical security to prevent tampering with the collected event information. OLI receives and forwards syslog and log file events created by a wide variety of hardware and software network products.

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




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

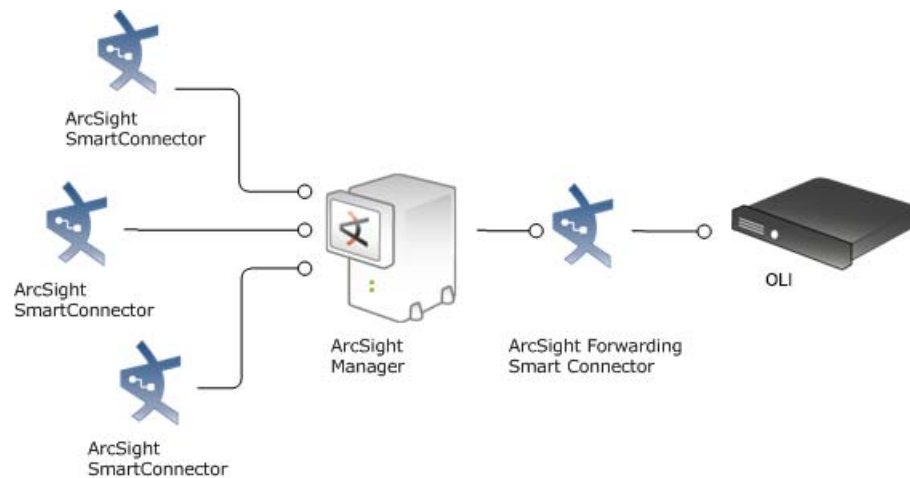


Figure 1-3 OLI can store events sent by ArcSight Manager




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

Chapter 2: Installation and Initialization

This chapter includes deployment and configuration information.

This chapter includes information on the following topics.

[“Deployment Planning” on page 23](#)

[“Initial Configuration” on page 24](#)

[“Installing an OLI” on page 26](#)

[“Installation Steps” on page 31](#)

[“Connecting to OLI for the first time” on page 37](#)

[“Configuring OLI” on page 38](#)

[“Using SmartConnectors to Collect Events” on page 41](#)

Deployment Planning

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

Storage Strategy

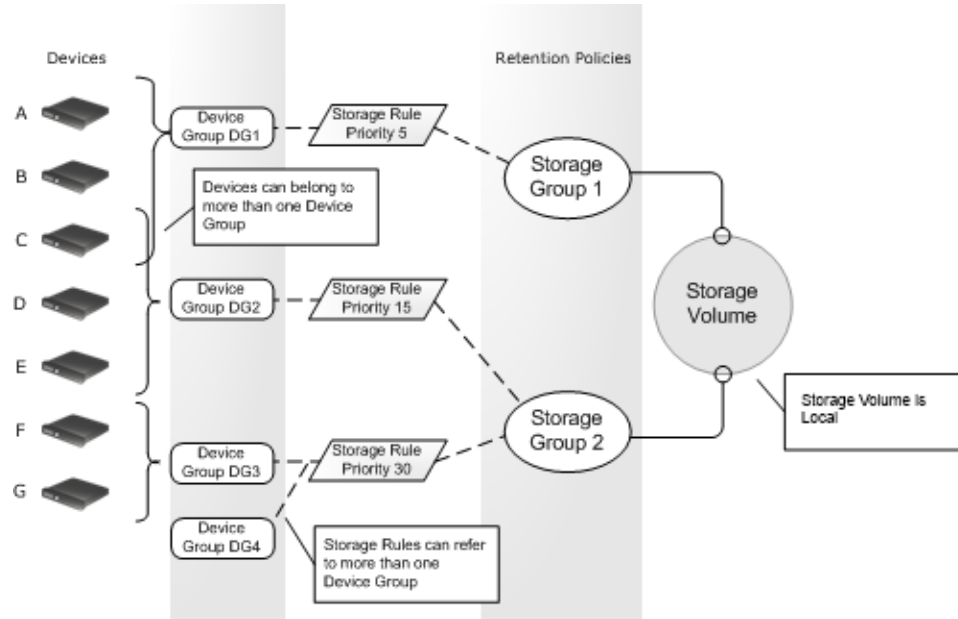
OLI events are stored locally.

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

Retention Policy

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

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



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



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

Initial Configuration

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

- [“Storage Groups” on page 25](#)
- [“Indexed Fields and Full-text Indexing” on page 25](#)
- [“Receivers” on page 25](#)

After the initial configuration, you can do additional configuration on OLI to implement your retention policies. See [“Configuring OLI” on page 38](#) for information on devices, storage groups, and storage rules.

Storage Volume

The initialization process sets up the storage volume of the non-enterprise OLI with 7 GB.

Storage Groups

Two storage groups, the Default Storage Group and the Internal Event Storage Group, are created automatically during OLI initialization.

These storage groups come pre-configured with the following settings:

Table 2-1 Pre-configured Default Storage Group Settings

Attribute	OLI
Size	Storage Volume/2
Retention Period	180 days (Usually not specified in license.)

Table 2-2 Pre-configured Internal Storage Group Settings

Attribute	OLI
Size	3 GB
Retention Period	365 days (Usually not specified in license.)

OLI can have a maximum of six storage groups; therefore, you can create an additional four storage groups after your OLI has been initialized. Each storage group can have different settings. You can change the retention policy and size for all storage groups, but you can only change the name of the user-defined storage groups. See [“Storage Groups” on page 136](#) for the details of adding and resizing storage groups, and changing their retention policies.

Indexed Fields and Full-text Indexing

Frequently used fields are indexed during initialization. You can add additional fields to the index, but once a field has been added, you cannot unindex it. See [“Indexing” on page 111](#) for more information. OLI comes prepared for full-text indexing.

Receivers

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

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

- A UDP receiver—Enabled by default. If you are installing OLI as root, the UDP receiver is on port 514. For non-root installs, it is on port 8514. If this port is already occupied, the initialization process selects the next higher

unoccupied port. This port should be allowed through any firewall rules you have configured.

- A TCP receiver—Enabled by default. If you are installing OLI as root, the TCP receiver is on port 515. For non-root installs, it is on port 8515. If this port is already occupied, the initialization process selects the next higher unoccupied port. This port should be allowed through any firewall rules you have configured.
- A SmartMessage receiver—Enabled by default. To receive events from a SmartConnector, download the SmartConnector and set the **Receiver Name** to be “SmartMessage Receiver” when configuring the destination.

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



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

The preconfigured folder follower receivers include:

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



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

For more information about how to enable these receivers, see [“Receivers” on page 25](#). For information about receivers in general, see [“Receivers” on page 140](#).

Licensing

After installing OLI, you can view the specific details of the current license on the **Configuration > License Information** page and the **System Administration > License and Update page**. For more information, see [“License Information” on page 227](#), and [“License & Update” on page 236](#).

Installing an OLI

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

Supported Platforms

For information about the platforms on which you can install and use OLI, refer to the Release Notes for your version.

Downloading the OLI

The OLI is available in these types: Downloadable Version and the Enterprise Version. The Downloadable Version is free.

Use the following table to determine where you can download the software.

OLI type...	Download from...
Downloadable Version (Free!)	HP Software Depot home: https://h20575.www2.hp.com/ecommerce/efulfillment/downloadpage.do
Enterprise Version	Follow the URL included in the Electronic Delivery Receipt you receive from HP in an email after placing the order.

You need to have a server with supported operating system and storage available to install the OLI, as described in the Release Notes.

Acquiring a License for OLI

OLI includes a trial license that you can use for a limited period of time for test and evaluation purposes. The trial license provides a user with a fully functional OLI and permits a minimal amount of device event sources and a minimal data volume.

- The Enterprise version of OLI requires a license file. You can apply the license file when you install OLI or apply one later. You can get a trial license for the Enterprise OLI version.
- The Downloadable version uses the trial license. You can upgrade by purchasing the Enterprise version and applying the license file.

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



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

Note

How Licensing Works on the OLI

The license for OLI defines its device event source limit, data volume limit, and the aggregated storage limit.

Device event source limit—The licensed maximum number of device event sources that OLI can configure.

Daily data volume limit—A per day limit on the amount of incoming data. For example, the limit might be 20 GB per day. The sum of the size of the events is used to determine this value.

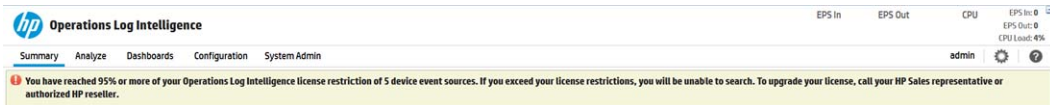
Aggregated storage limit—The licensed maximum storage, for example, 80 GB.

Even if this limit is exceeded, the OLI continues to collect and store events; therefore, no events are lost.

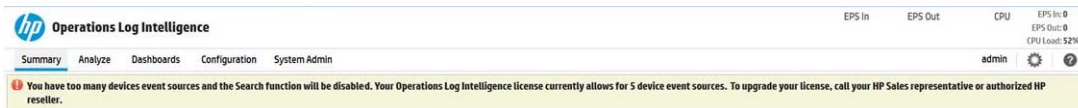


If you are using ArcSight Connectors to send events to the OLI, make sure you are running connector version 5.1.3.5870.0 or later on your connectors to ensure that event size is accurately accounted on the OLI.


When the number of device event sources is close to the maximum allowed by the license, the header in all the OLI pages displays a warning, as shown in the following figure.



When the number of device event sources exceeds the maximum allowed by the license, the header in all the OLI pages displays an error, as shown in the following figure and the search function is disabled.



If you have at least one data violation in the past 30 days, the header in all the OLI pages displays the following warning message. After 30 days, the warning is no longer displayed.



Only one error/warning message appears on the header in all the OLI pages. Priority of the messages:

- 1 Device event source limit exceeded error
- 2 Devices event source limit is going to be exceeded warning
- 3 Data volume exceeded warning

You can also view the data limit violation information for the last 30 days on the License Information page (**Configuration > License Information**).

The License Information page lists the data stored on your OLI on a day-by-day basis for the last 30 days. It also indicates the days on which data limits were exceeded as shown in the following figure.

Date	Data Stored (MB)	Limit Exceeded
11/22/13	0	false
11/23/13	0	false
11/24/13	0	false
11/25/13	0	false
11/26/13	0	false
11/27/13	0	false
11/28/13	0	false
11/29/13	0	false
11/30/13	0	false
12/1/13	0	false
12/2/13	0	false
12/3/13	0	false
12/4/13	0	false
12/5/13	0	false
12/6/13	0	false
12/7/13	0	false
12/8/13	33844	true

You can view the number of device event sources configured by OLI on the Devices page (**Configuration > Devices**).

Name	IP Address	Receiver	Creator	Last Editor
labm3arrendb05.devlab.ad [SmartMessage Receiver]	16.60.134.202	SmartMessage Receiver	System	
labm3rden02.devlab.ad [SmartMessage Receiver]	16.55.245.65	SmartMessage Receiver	System	
labm3rden03.devlab.ad [SmartMessage Receiver]	16.59.60.216	SmartMessage Receiver	System	

If you exceed the data limit frequently, you should consider purchasing a license that suits your needs. Contact your HP sales representative or authorized HP reseller to purchase a new license. Once you obtain the new license, follow the instructions in this Guide to apply it on your OLI.

Prerequisites for Installation

Make sure these prerequisites are met before you install an OLI:

- Before deploying in a production environment, get the valid license file. If you do not have a license file, see [“Acquiring a License for OLI” on page 27](#).



OLI includes a limited trial license for test and evaluation purposes.

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

Installation Modes

The OLI can be installed in the following three modes:

- GUI—In this mode, a wizard steps you through the installation and configuration of OLI. For instructions, see [“Using the GUI Mode to Install OLI” on page 31](#).
- Console—In this mode, a command-line process steps you through the installation and configuration of OLI. For instructions, see [“Using the Console Mode to Install OLI” on page 33](#).
- Silent—In this mode, you provide the input required for installation and configuration through a file. Therefore, you do not need to interact with the

installer to complete the installation and configuration. However, before you can use this mode, you must run the installation and configuration using one of the other modes to record the input in a file. For instructions, see [“Using the Silent Mode to Install OLI” on page 34](#).

Installation Steps

This section describes all three modes of OLI installation.

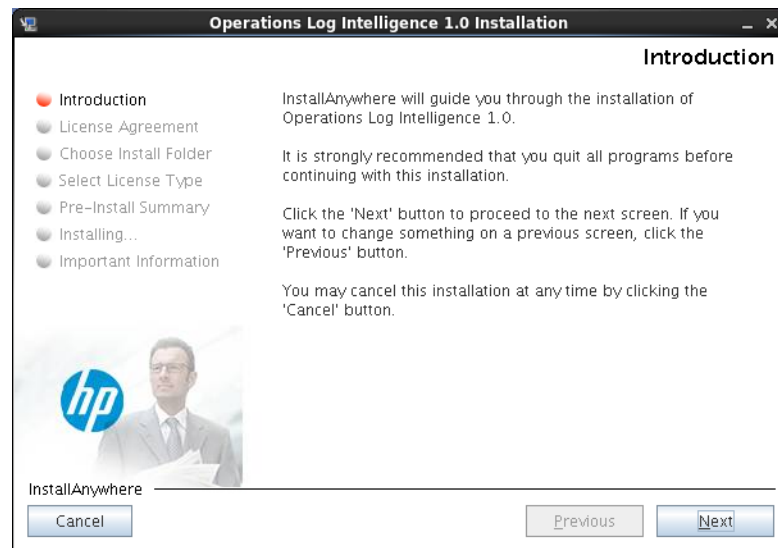
Using the GUI Mode to Install OLI

You can install OLI as a root user or as a non-root user. See [“Prerequisites for Installation” on page 29](#) for details and restrictions.

To install the OLI using the GUI mode:

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


```
chmod +x HP-Operations-Log-Intelligence-1.0.0.XXXX.bin
./HP-Operations-Log-Intelligence-1.0.0.XXXX.bin
```
- 2 The installation wizard launches, as shown in the following figure. Click **Next**.



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



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

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

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



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

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

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

Field	Notes
Configure Logger as a service	<p>Indicate whether to configure OLI to run as a service.</p> <p>Select this option to create a service called <code>arcsight_logger</code>, and enable it to run at levels 2, 3, 4, and 5.</p> <p>If you do not enable OLI to start as service during the installation process, you still do so later. For instructions on how to enable OLI to start as a service, see “System Settings” on page 238.</p>

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

11 Click **Next** to initialize OLI components.

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

12 Click **Next** to configure storage groups and storage volume and restart OLI.

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

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

Now that you are done installing and initializing your OLI, you can connect, log in, and start configuring your OLI to receive events. For instructions and information, see [“Connecting to OLI for the first time” on page 37](#) and [“Configuring OLI” on page 38](#).

Using the Console Mode to Install OLI

Make sure the machine on which you will be installing the OLI complies with the platform requirements listed in the Release Notes, and that the prerequisites listed in [“Prerequisites for Installation” on page 29](#) are met.

You can install OLI as a root user or as a non-root user. See [“Prerequisites for Installation” on page 29](#) for details and restrictions.

To install the OLI using the Console mode:

1 Run these commands from the directory where you copied the OLI:

```
chmod +x HP-Operations-Log-Intelligence-1.0.0.XXXX.bin
./HP-Operations-Log-Intelligence-1.0.0.XXXX.bin
```

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

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

```
InstallAnywhere will guide you through the installation of
HP_Operations Log Intelligence 1.0.0
```

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

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

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

PRESS <ENTER> TO CONTINUE:

- 3 The next screens display license information. Installation and use of OLI 1.0.0 requires acceptance of the license agreement. Press Enter to display each part of the license agreement, until you reach the following prompt:

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

- 4 Type Y and press Enter to accept the terms of the License Agreement.
- 5 The subsequent prompts are exactly similar to the ones described for the GUI mode install in [“Using the GUI Mode to Install OLI” on page 31](#). Follow the instructions provided for the GUI mode install to complete the installation.

Using the Silent Mode to Install OLI

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

About Licenses for Silent Mode Installations

As for any OLI installation, each silent mode installation requires a unique license file. You must obtain licenses as described in [“Acquiring a License for OLI” on page 27](#) and place them on the machines on which you will be installing OLI in silent mode, or ensure that the location where the licenses are placed is accessible from those machines.

Generating the Silent Install Properties File

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

- 1 Log in to the machine on which you can install OLI to generate an installation properties file.
- 2 Run these commands:

```
chmod +x HP-Operations-Log-Intelligence-1.0.0.XXXX.bin
./HP-Operations-Log-Intelligence-1.0.0.XXXX.bin -r
<directory_location>
```

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

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

- 3 Install OLI in GUI mode, as described in [“Using the GUI Mode to Install OLI” on page 31](#).

- 4 Once the installation completes, navigate to the directory location you specified for the `installer.properties` file earlier.

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

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

#Choose Install Folder
#-----
USER_INSTALL_DIR=/opt/Operations Log Intelligence/53

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

Installing OLI in Silent Mode

Make sure the machine on which you will be installing the OLI complies with the platform requirements listed in the Release Notes for your version, and that the prerequisites listed in [“Prerequisites for Installation” on page 29](#) are met.



To install OLI in silent mode, you must log in as a non-root user.

To install the OLI using the silent mode:

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

OR

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

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

```
chmod +x HP-Operations-Log-Intelligence-1.0.0.XXXX.bin

./HP-Operations-Log-Intelligence-1.0.0.XXXX.bin -i SILENT -f
<path to installer.properties>
```

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

Starting and Stopping the OLI

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

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

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

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

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

The screenshot displays the HP Operations Log Intelligence (OLI) System Admin interface. The 'Process Status' section is active, showing a 'System' section and a 'Processes' section. The 'System' section includes a table with columns for System, Status, Load, CPU Usage, Memory Usage, and Data Collected. The 'Processes' section includes a table with columns for Process, Status, Uptime, CPU Usage, and Memory Usage. Red arrows point to the 'System Section' and 'Processes Section' labels. The 'System' table shows 'lenaDLDEVLABAD' with status 'running'. The 'Processes' table shows 'apache' with status 'running' and other processes like 'aps', 'connector', 'mysqld', and 'postgres'.

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

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

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

Uninstalling the OLI

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

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

```
./UninstallerData/Uninstall HP-Operations-Log-Intelligence-1.0.0
```

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

Connecting to OLI for the first time

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

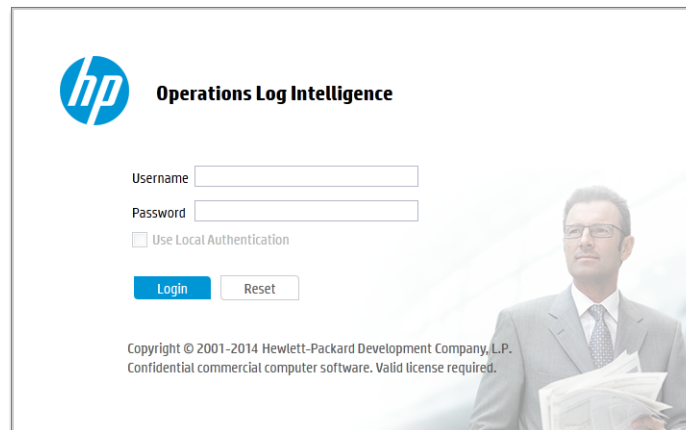
To connect and log into OLI:

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

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

Once you connect, the following Login screen is displayed.

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



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

- 3 Username: `admin`
Password: `password`
- 4 Once you have successfully logged in, go to the section, [“Configuring OLI” on page 38](#) for information on how to set up your OLI to start receiving events.



For security reasons, be sure to change the default credentials as soon as possible after connecting to OLI for the first time. Refer to [“Change Password” on page 266](#) for instructions.

For more information about the log in screen and connecting to OLI, see [“Connecting to the OLI User Interface” on page 45](#).

Configuring OLI

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

Receivers

Now that you have finished installing OLI, you can set up receivers to listen for events. You can use the preconfigured receivers or add your own. OLI comes with preconfigured with a TCP receiver, a UDP Receiver, and a SmartMessage receiver already enabled and ready to receive events. OLI also comes

pre-configured with folder follower receivers for OLI's Apache Access Error Log, the system Messages Log, and the system Messages Audit Log (if auditing is enabled on your Linux OS). You must enable these receivers in order to use them. Receivers can be disabled and re-enabled later. You can add, change, and delete them as needed.

For more information on receivers, see ["Receivers" on page 140](#).

Enabling the Pre-configured Receivers

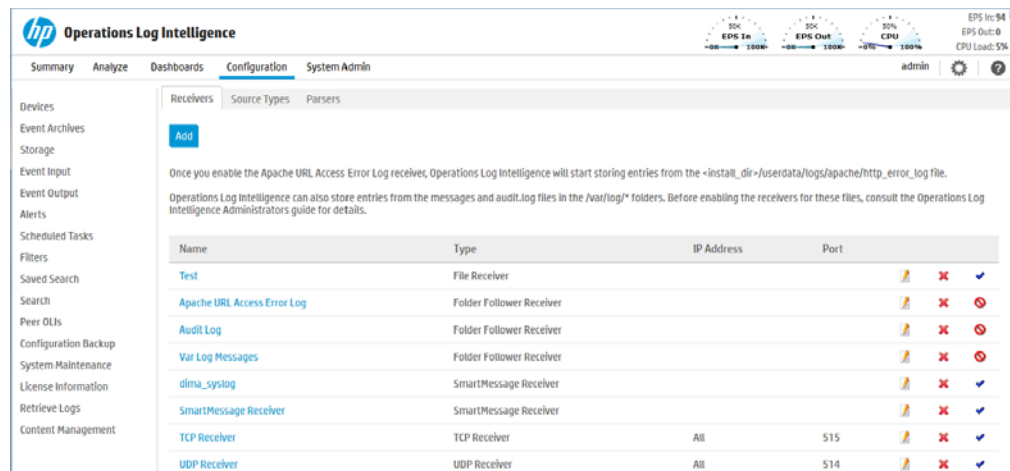
When you first log in by using the URL you configured, OLI will display a banner like the one below, telling you about the disabled receivers.



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



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




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

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

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

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Receivers** tab (right panel).

- 3 Click the disabled icon () at the end of the row.

For information on how to use the preconfigured SmartMessage receiver, see [“Using SmartConnectors to Collect Events” on page 41](#). For more information on enabling and disabling receivers, see [“Working with Receivers” on page 143](#).

Devices

OLI begins storing events when an enabled receiver receives data or, in the case of a file receivers, when the files become available. Using a process called autodiscovery, OLI 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 OLI received events.

You can also create devices preemptively, by entering the IP addresses or hostnames of data sources that you expect to be sending events to OLI. You might do this if you do not want to wait for autodiscovery, or if you want to control the initial naming of each device. Discovered devices are named for their host, or if the DNS lookup fails, for their IP address, and their receiver. For information about creating devices, see [“Devices” on page 127](#).

Device Groups

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

You can change and delete device groups freely as your needs change. Setting up device groups initially is not critical; incoming events that are not assigned to a device group are automatically sent to the Default Storage Group. For the details of setting up device groups, see [“Device Groups” on page 129](#).

Storage Rules

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

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

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

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

device group, and send them to a particular storage group for retention based on event category.

See “Storage Rules” on page 138 for more information.

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

Using SmartConnectors to Collect Events

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

SmartMessage

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



Caution

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

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

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



Note

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

Downloading SmartConnectors

For the Enterprise Version of OLI, contact your HP sales representative or customer support for the location to download SmartConnectors.

A restricted set of ArcSight SmartConnectors are supported and available For the Downloadable Version of OLI. You can download these SmartConnectors from the same location from which you downloaded OLI. The configuration guides for the supported SmartConnectors are available at the same web site. To learn more about ArcSight SmartConnectors, visit <http://www.arcsight.com>.

Configuring a SmartConnector to Send Events to OLI

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

To configure a SmartConnector to send events to OLI:

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



Refer to the documentation that came with your SmartConnector for instructions.

- 2 Specify the required parameters. Enter the OLI hostname or IP address and the name of the SmartMessage receiver. These settings must match the receiver in OLI that listen for events from this connector.
 - ◆ To use the preconfigured receiver, specify "SmartMessage Receiver" as the **Receiver Name**.
 - ◆ To communicate between an ArcSight SmartConnector and OLI, configure the SmartConnector to use the port configured for the OLI.
 - ◆ For un-encrypted CEF syslog, enter the OLI hostname or IP address, the desired port, and choose UDP or TCP output.

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

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

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

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

Configuring SmartConnectors for Failover Destinations

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

To configure a failover destination, follow these steps:

- 1 Configure the SmartConnector for the primary OLI 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 OLI selected, choose **Modify**. Then select **Add failover destination**.
- 5 Enter information for the secondary OLI.
- 6 Restart the SmartConnector for the changes to take effect.
- 7 For more information about installing and configuring ArcSight SmartConnectors, refer to the ArcSight SmartConnector User's Guide, or specific SmartConnector Configuration Guides, available from the Protect 724 Community at <https://protect724.arcsight.com>.

Chapter 3: User Interface and Dashboards

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

This chapter includes information on the following topics.

Logging in: see [“Connecting to the OLI User Interface”](#) on page 45

Navigation: [“Navigating the User Interface”](#) on page 46

Summary Dashboard: [“Summary”](#) on page 49

Dashboards: see [“Dashboards”](#) on page 51

Performance monitoring: see [“The Default Monitor Dashboard”](#) on page 62

Connecting to the OLI User Interface

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

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

Logging In

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

To connect and log into OLI:

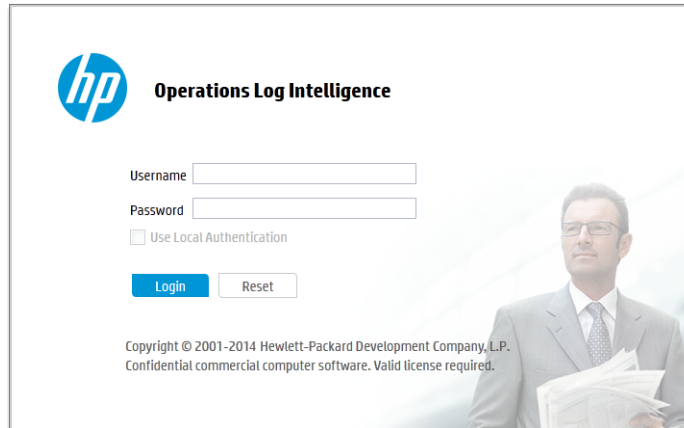
- 1 Use the URL configured during OLI installation to connect to OLI through a supported browser.

Connect using this URL:

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

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

Once you connect, the following Login screen is displayed.



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

Username: admin
Password: password



For security reasons, be sure to change the default credentials as soon as possible after connecting to OLI for the first time. Refer to [“Change Password” on page 266](#) for instructions.

If you forget your system password, contact your system administrator.

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

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

- ◆ Use Local Authentication—The “Use Local Authentication” check box is always displayed, but only becomes active when a login attempt fails. By default, this option is available only for the default admin. For more information on the Use Local Authentication, see [“Local Password Fallback” on page 259](#).

Navigating the User Interface

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

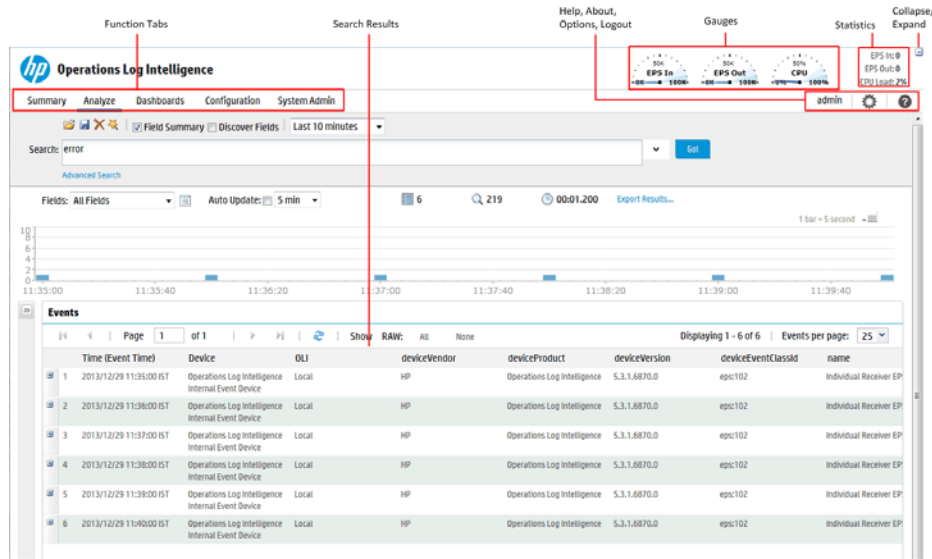


Figure 3-1 Overall layout of the OLI interface

Gauges at the top of the screen provide an indication of the throughput and CPU usage information, available in more detail on the Monitor Dashboard (“Dashboards” on page 51). The range of the gauges can be changed on the Options page. The current logged-in user’s name is shown below the statistics. The gauge and logo bar can be collapsed to allow more room on the screen for search results. Click the icon to collapse the bar, and the icon to expand it.

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

The Summary, Analyze, Dashboards tabs provide access to various OLI functions and data stored on it. The Configuration and System Admin tabs are used for configuring the system administration and configuration settings on OLI.

Help

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

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

Options

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

Additionally, the default start page (home page) for all users and specific start pages for individual users can be set on the Options page. These pages indicate which user interface page is displayed after a user logs in.

Options

System

EPS input rate gauge max

EPS output rate gauge max

Default start page for all users

Personal

Default start page for admin

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

If you want to set...	Configure the...
The same start page for all users	<p>Default start page for all users option to the desired value.</p> <p>This is a global setting for your OLI. To override this setting, configure a different start page for specific users by using the Default start page for <username> option.</p> <p>When you set Default start page for all users option to Dashboards, the Monitor Dashboard is the default dashboard displayed for all users, unless the users have configured other dashboards as their defaults, as described in “To set a dashboard as default:” on page 58.</p>
A different start page for specific users	<p>Default start page for <username> option to the desired value.</p> <p>This setting overrides the global Default start page for all users setting.</p> <p>When this option is set to “Use default for all users”, the global default page (Default start page for all users) value is used for all users.</p>
A specific dashboard for a specific user	<p>Default start page for <username> option to Dashboards.</p>
OR	
A specific dashboard for all users	<p>The Monitor Dashboard is the default dashboard displayed for all users. However, if you want to display a different dashboard for one or more users, set the desired dashboard as the default when logged in as those users. For details, see “To set a dashboard as default:” on page 58.</p>

Logout

Click **user name > Logout** 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 OLI session.

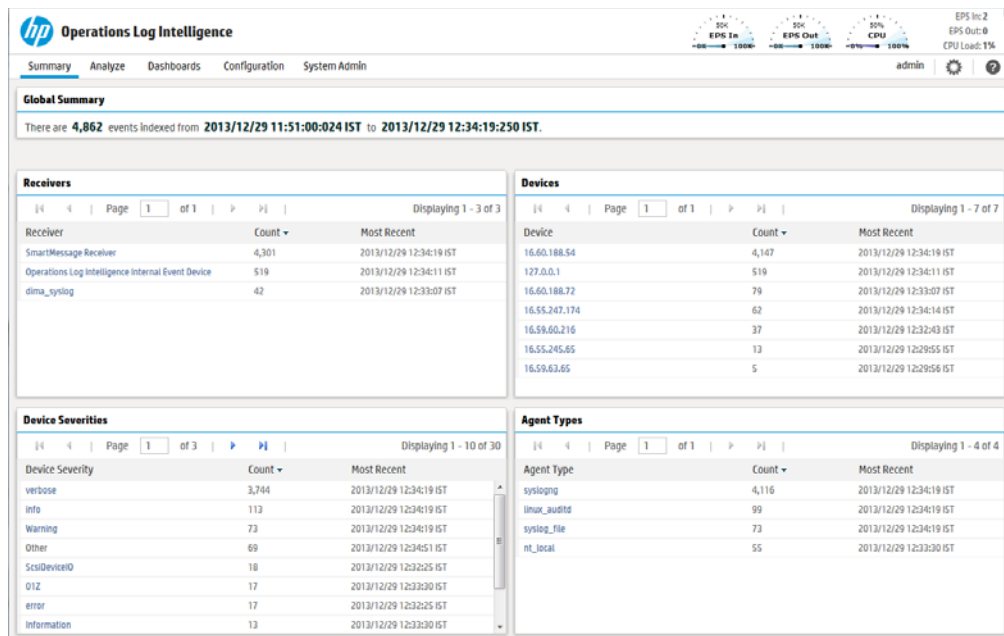
OLI automatically logs you out after a user-configurable length of time (15 minutes by default). To change this length of time, see [“Users/Groups” on page 251](#).



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 OLI session.

Summary

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



Specifically, the Summary page contains the following panels:

- **Global Summary**

The number of events indexed on your OLI during the time period displayed on the screen. The time period is dependent on the retention policy of your OLI, where the start and end times are the time of the oldest events stored on your OLI (that have not aged out due to retention) and the current time, respectively.

- **Receivers**

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

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

- **Devices**

A device is a named event source, comprised of an IP address (or hostname) and a receiver name. Two receivers can receive events from the same IP address, so an IP address alone is insufficient to identify a device. An Event source is the device that directly sends the event to OLI. When an event is sent through a SmartConnector, the event source is the system on which the SmartConnector is running and not the device that sent the event to the SmartConnector

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

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

- **Agent Severities**

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





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

- **Agent Types**


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

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

The Summary page is pre-designed to display the information described above. You cannot change or add other panels to it. If you need to display other information, you can create a custom Dashboard as described in ["Dashboards" on page 51](#). The information displayed on the Summary page is for your local OLI only, and does not include information about peer OLIs even if peers are configured.

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

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



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

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

Dashboards

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

Each Dashboard contains one or more panels of these types:

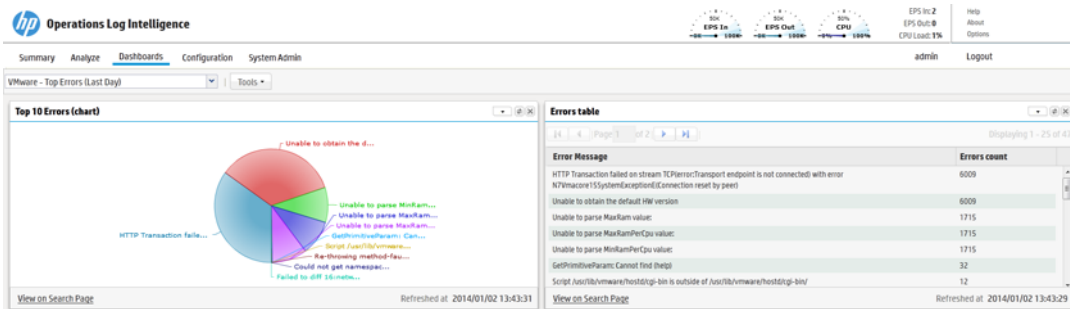
- Search Results
 - The Search Results panels display events that match the query associated with the panel.
- Monitor
 - The Monitor panels display the real-time and historical status of various OLI components such as receivers, forwarders, storage, CPU, and disk.

■ Summary

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

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

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



You can click the “View on Search Page” link in the Search Results panels to go to the Analyze (Search) page and view the event details; the panel query is automatically run and the search results are displayed. Additionally, you can drill-down from any chart to quickly filter down to events with specific field values. To do so, identify the value in the chart on a Search Results Chart panel and click it to drill-down to events that match the value. When you click on a chart value (a bar, or pie section), the query is rerun on the Analyze (Search) page with an additional WHERE operator clause that includes the field name and value you clicked on the chart.



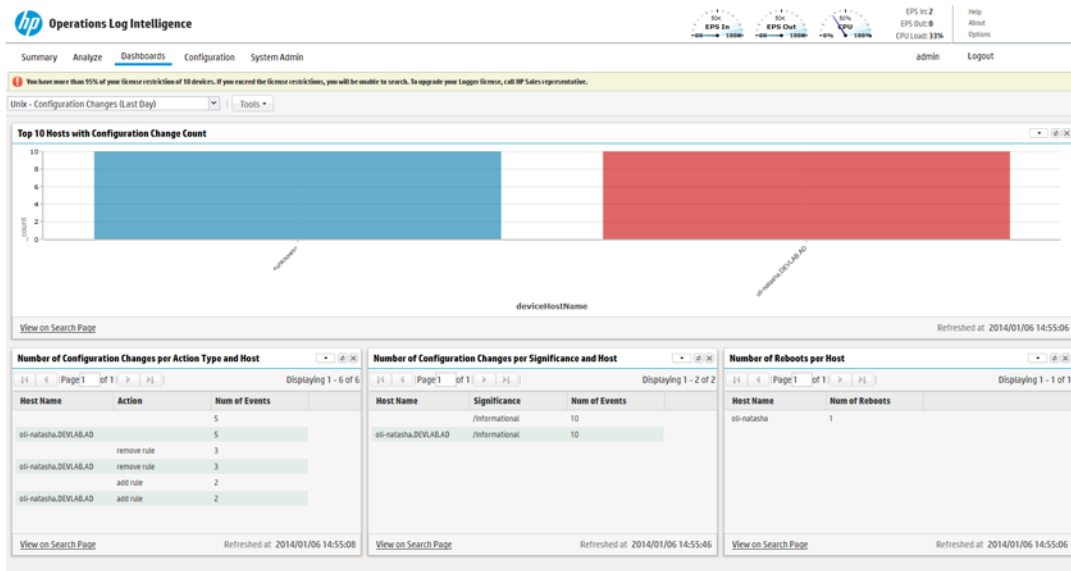
The saved search query associated with the Search Results panel in the dashboard is not modified. If you need to return to the dashboard from the drill-down screen, use the Back function of your browser.

The Monitor dashboard is displayed by default, unless you configure another dashboard as your default. The Monitor dashboard displays the Summary page, which shows the status of CPU Usage, Event Flow, Receivers, Forwarders, and Storage Groups in a summarized view. The other pages available from the Monitor dashboard are Platform, Network, Operations Log Intelligence, Receivers, Forwarders, and Storage. These views are described in detail in “[The Default Monitor Dashboard](#)” on page 62. You cannot change or adjust any panels

available in this dashboard; however you can add specific Monitor panels to a custom dashboard as described in this section.

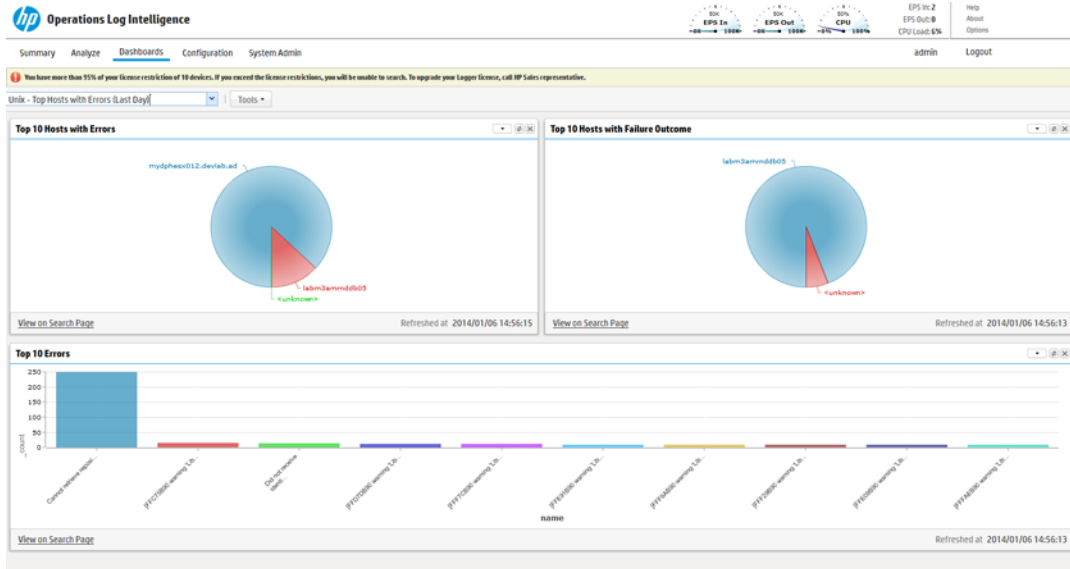
The following dashboards are also available by default:

- **Unix - Configuration Changes (Last Day)**
 - ◆ **Top 10 Hosts with Configuration Change Count** - Displays the top 10 hosts with configuration changes for the last 24 hours as a bar graph.
 - ◆ **Number of Configuration Changes per Action Type and Host** - Displays all configuration changes per action type and host for the last 24 hours in a table sorted by the number of configuration changes.
 - ◆ **Number of Configuration Changes per Significance and Host** - Displays all configuration changes per significance and host for the last 24 hours in a table sorted by the number of configuration changes.
 - ◆ **Number of Reboots per Host** - Displays all reboots per host for the last 24 hours in a table sorted by the number of reboots.



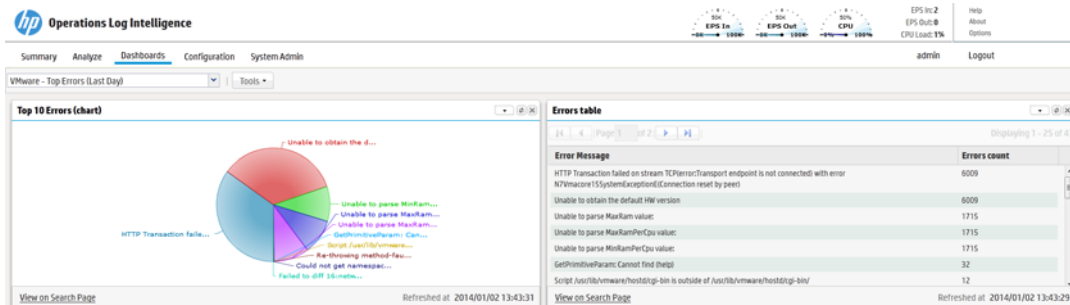
- **Unix - Top Hosts with Errors (Last Day)**
 - ◆ **Top 10 Hosts with Errors** - Displays the top 10 host errors for the last 24 hours as a pie chart.
 - ◆ **Top 10 Hosts with Failure Outcomes** - Displays the top 10 host with failure outcomes for the last 24 hours as a pie chart.
 - ◆ **Top 10 Errors** - Displays the top 10 errors for the last 24 hours as a bar graph.

3 User Interface and Dashboards



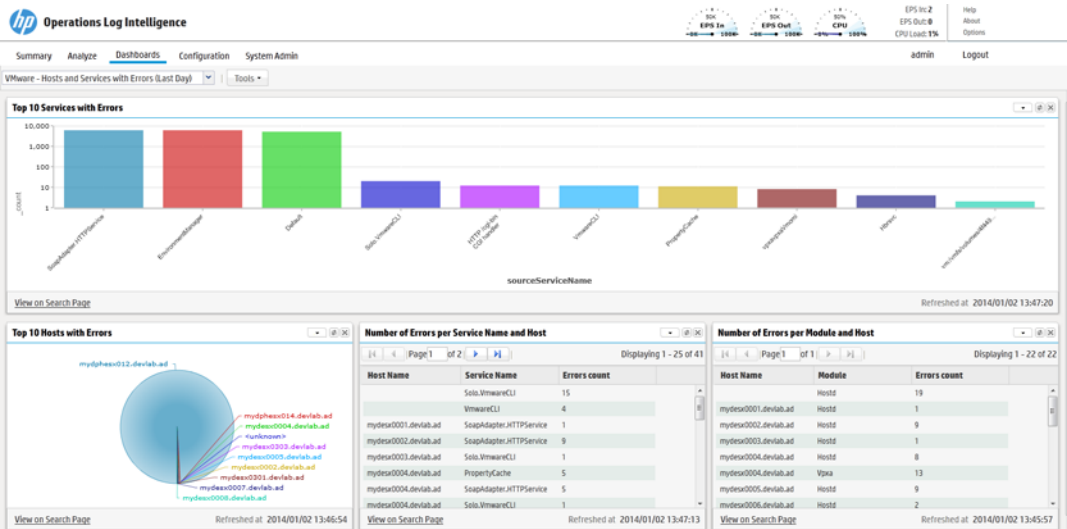
- **VMware - Top Errors (Last Day)**

- ◆ **Top 10 Errors (chart)** - Displays the top 10 errors for the last 24 hours as a pie chart.
- ◆ **Errors table** - Displays all errors for the last 24 hours in a table sorted by the error count.

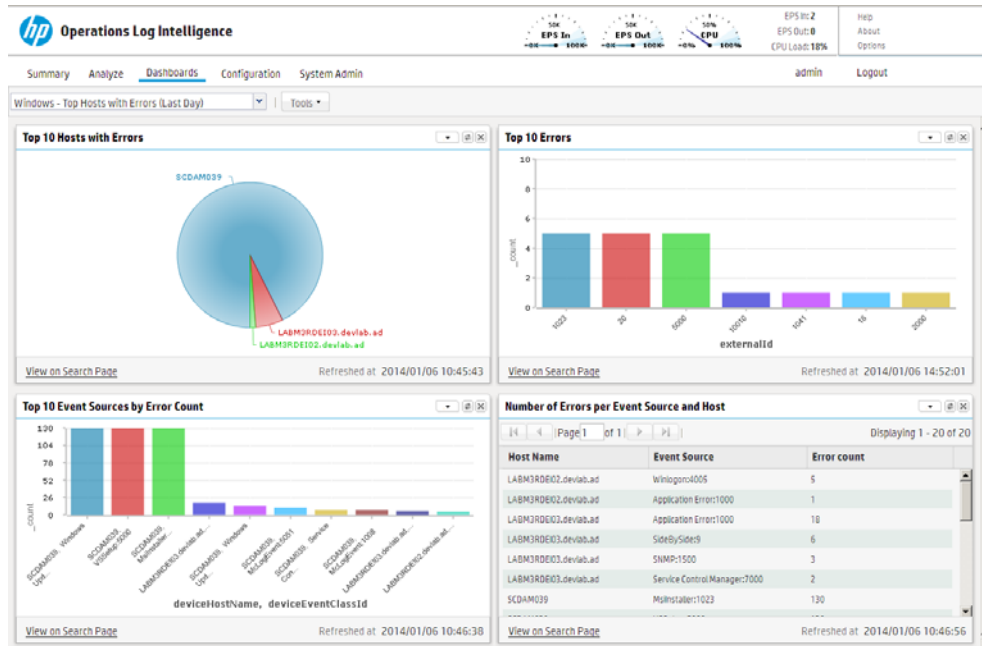


- **VMware - Hosts and Services with Errors (Last Day)** - Displays information in the following panels for the last 24 hours.

- ◆ **Top 10 Services with Errors** - Displays the top 10 services with errors for the last 24 hours as a bar graph sorted by error count.
- ◆ **Top 10 Hosts with Errors** - Displays the top 10 hosts with errors for the last 24 hours as a pie chart.
- ◆ **Number of Errors per Service Name and Host** - Displays all service name and host errors for the last 24 hours in a table sorted by host, service name, and then error count.
- ◆ **Number of Errors per Module and Host** - Displays all module and host errors for the last 24 hours in a table sorted by host, module, and then error count.

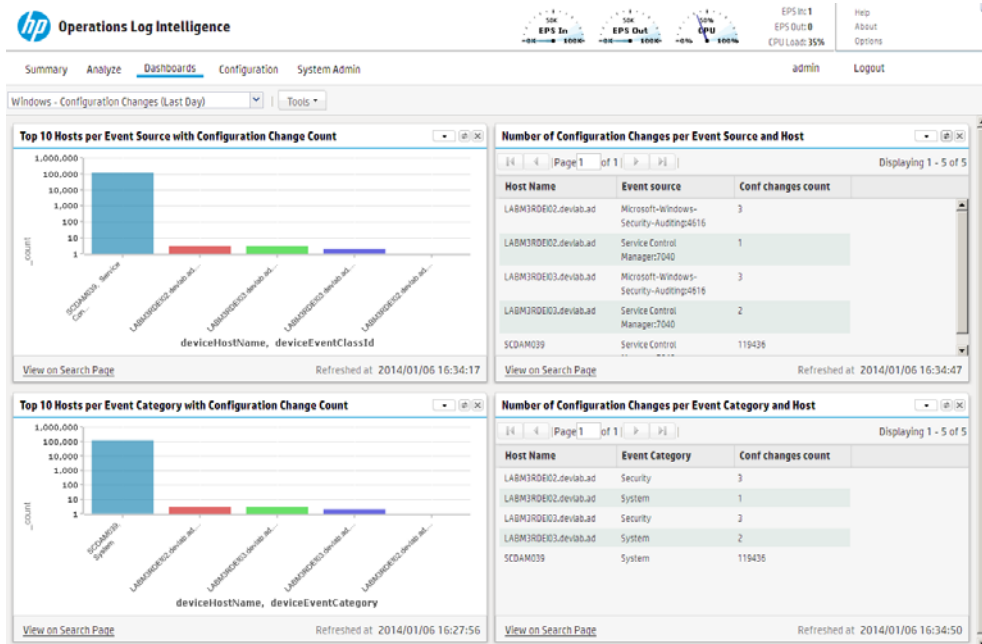


- **Windows - Top Hosts with Errors (Last Day)**
 - ◆ **Top 10 Hosts with Errors** - Displays the top 10 host errors for the last 24 hours as a pie chart.
 - ◆ **Top 10 Errors** - Displays the top 10 errors for the last 24 hours as a bar graph.
 - ◆ **Top 10 Event Sources by Error Count** - Displays the top 10 event source errors by error count for the last 24 hours as a bar graph.
 - ◆ **Number of Errors per Event Source and Host** - Displays the total number of event source and host errors for the last 24 hours in a table sorted by the number of errors.



- **Windows - Configuration Changes (Last Day)**

- ◆ **Top 10 Hosts per Event Source with Configuration Change Count** - Displays the top 10 hosts per event source with configuration changes for the last 24 hours as a bar graph.
- ◆ **Number of Configuration Changes per Event Source and Host** - Displays all configuration changes per event source and host for the last 24 hours in a table sorted by the number of changes.
- ◆ **Top 10 Hosts per Event Category with Configuration Change Count** - Displays the top 10 hosts per event category with configuration changes for the last 24 hours as a bar graph.
- ◆ **Number of Configuration Changes per Event Category and Host** - Displays all configuration changes per event category and host for the last 24 hours in a table sorted by the number of changes.



Creating and Managing Dashboards

Users can create both shared and private dashboards.

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

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

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

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

The following steps outline the process of creating a dashboard:

- 1 Ensure that you have the privileges to create a dashboard.
- 2 Create a dashboard. See [“To add a dashboard:” on page 57](#).
- 3 Add panels to the dashboard you created. See [“To add a panel to a dashboard:” on page 59](#).

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

To add a dashboard:

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

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

- 5 Click **Create**.

The dashboard is created. You must add panels to the dashboard next, as described in [“To add a panel to a dashboard:” on page 59](#).

To edit a dashboard:

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

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

To delete a dashboard:

- 1 Click **Dashboards** from the top-level menu bar.
- 2 Select the dashboard that you want to delete.
- 3 Click the **Tools** drop-down menu and select **Delete Dashboard**.

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

To set a dashboard as default:

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

For all OLI users, the Summary page (accessible from the Summary navigation option in the top-level menu bar) is the default home page. That is, when you log in, the Summary page is displayed first. However, you can configure OLI to display a specific dashboard as the default home page for you. To do so, first configure the option described in [“Options” on page 47](#), then

To select the dashboard you want to display by default:

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

Adding and Managing Panels in a Dashboard

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

These panels provide the same information available through the default Monitor dashboard and the default Summary dashboard, however in a modular form that enables you to choose specific views. For example, if you want to view the EPS for the last 4 hours on all receivers, add the panel Type “Monitor Graph”, and select “(OLI) All EPS Out-All EPS In - 4 hour” as the Graph, or if you want to view the EPS on Forwarders in a table form, select the “Monitor (Forwarders)” panel Type. Similarly, if you want to view only the summary information for receivers on your OLI, add the panel of Type “Summary (Receivers)”. Besides the four Summary panels (Agent Severities, Agent Types, Receivers, and Devices), you can also create a user-defined Summary panel in which you can select *any indexed, non-time field* by which you want to categorize event summary. For example, if you want to add a Summary panel to display event summary categorized by “destinationAddress”, you can add a panel of Type “Summary (User Defined)” for this field if it is indexed on your OLI.

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

configure a receiver, click the Configure link on top of the Monitor (Receiver) panel.

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

The following describes the dashboard panels:

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

To add a panel to a dashboard:

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

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

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

To edit a panel:

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

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

Action	Description
All Panels	
Delete	Removes a panel from a dashboard.
Search Result Panels	
Edit Panel	Change Title, associated saved search, Chart Type, or Chart Limit
Edit Saved Search	Access the Edit Saved search page to edit the associated saved search query
View on Search Page	Runs the panel's query on the Search Results page (Analyze > Search) and displays matching events on that page

Action	Description
Refresh	Refreshes the current contents of the panel. Note: All other panel types are automatically refreshed; therefore, an explicit refresh is not required for them.
Summary Panels - User Defined	
Edit Panel	Change Title or field name by which events are categorized.

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

To delete a panel from a dashboard:

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

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

To change the layout of a dashboard:

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

- 1 Click **Dashboards** from the top-level menu bar.
- 2 Select the dashboard that contains the panel you want to rearrange.

- 3 Click the **Tools** drop-down menu and select **Change Layout**.
- 4 Point your cursor in the blue band that shows the panel title and drag the panel to a different position.
- 5 Click **Save** after you rearrange the panels.

The Default Monitor Dashboard

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

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

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

- 4-hours
- Daily
- Weekly

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

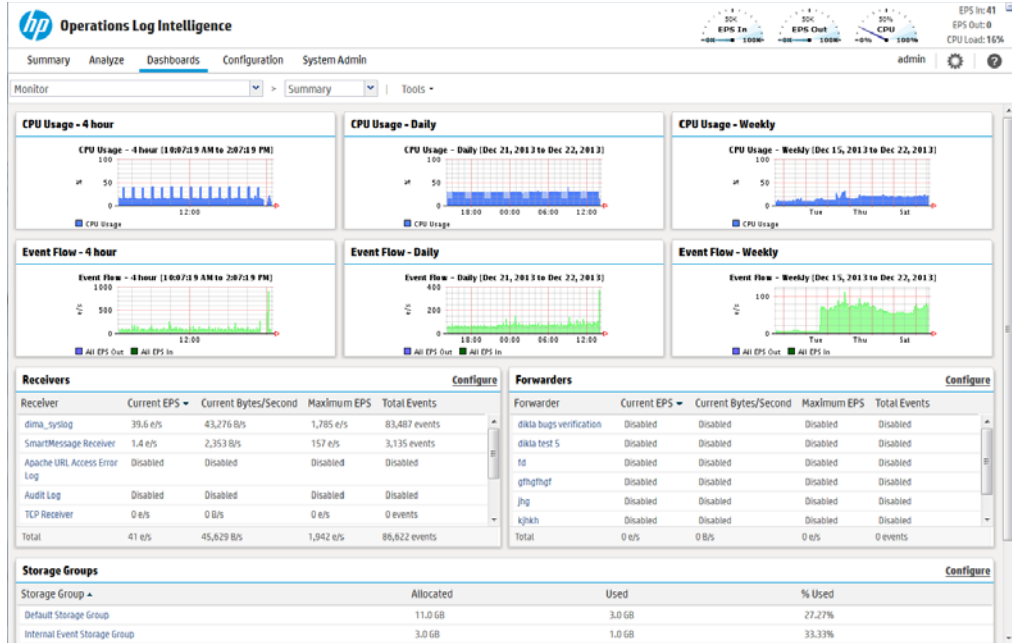


Figure 3-2 The Monitor dashboard displays summary information by default

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

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

Platform

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



Figure 3-3 Monitor dashboard - Platform page

Network

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

Operations Log Intelligence

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



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

Receivers

The Receivers monitor page shows total Events per Second (EPS) received and displays values for each configured receiver. The list of receivers includes all receivers known to the system, including those that are disabled. To create a new receiver, or to enable or disable one, see [“Receivers” on page 140](#).

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 165](#).

Storage

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

For more information about storage groups, see [“Storage Groups” on page 136](#).

3 User Interface and Dashboards



Figure 3-5 Monitor dashboard - Storage page

Chapter 4: Searching and Analyzing Events

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

- [“The Need to Search Events” on page 67](#)
- [“The Process of Searching Events” on page 68](#)
- [“Elements of a Search Query” on page 68](#)
- [“Syntax Reference for Query Expression” on page 81](#)
- [“Using the Search Builder Tool” on page 86](#)
- [“Search Analyzer” on page 90](#)
- [“Regex Helper Tool” on page 91](#)
- [“Search Helper” on page 93](#)
- [“Searching for Events on OLI” on page 96](#)
- [“Understanding the Search Results Display” on page 99](#)
- [“Exporting Search Results” on page 108](#)
- [“Indexing” on page 111](#)
- [“Saving Queries \(Saved Filters and Searches\)” on page 116](#)
- [“System Filters/Predefined Filters” on page 119](#)
- [“Alerts” on page 122](#)
- [“Live Event Viewer” on page 123](#)


The Need to Search Events

When you need to analyze events matching specific criteria, you will need to search for them on the OLI.

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

The Process of Searching Events

Searching through stored events is very simple and intuitive on OLI. It uses a flow-based search language that allows you to specify multiple search commands in a pipeline format. In addition, you can customize the display of search results, view search results as charts, and so on.



You enter the keywords or information you are searching for (referred as queries) in the Search text box, select the time range, and click Go, as shown in the previous figure. OLI 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, such as, `hostA.companyxyz.com`. Or, it can be 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" | chart count by name
```

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

OLI offers several convenient ways to enter a search query—Typing the query in the Search text box, using OLI’s Search Builder tool to create a query, or using a previously saved query (referred to as filter or saved search). When you type a query, the Search Helper facility provides suggestions and possible matches to quickly build a query expression. (See [“Search Helper” on page 93](#) for more information.)

Although a search query on OLI is as simple as entering a keyword to match, you will utilize the full potential of OLI’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 68](#).

Elements of a Search Query

A simple OLI search query consists of these elements:

- Query Expression
- Time range

- Field Set

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

Query Expression

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

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

Indexed Search | Search Operators

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

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

Indexed Search is described in [“Indexed Search” on page 69](#).

Search Operators are described in [“Search Operators” on page 73](#).

Indexed Search

The *Indexed search* uses OLI’s indexing capability to quickly and efficiently search for relevant data, and enables you to specify **keywords** and **field-based expressions** (indexed and non-indexed fields) in a boolean format.

Keyword Search (Full-text Search)

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 81](#).

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, HP 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.

Field-based Search

The OLI schema contains a predefined set of fields. You can add fields that are relevant to the events you collect on your OLI to its schema. A field-based search can only contain fields in OLI's schema. (See additional guidelines at [“Guidelines for Field-based Search Expressions” on page 72.](#)) The OLI indexing capability allows for these *fields* to be indexed. The OLI's search operation utilizes these indexed fields to yield significant search performance gains.

Although you can add indexed and non-indexed fields to a search query, **you will realize the search 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 111.](#) For discussion on field-based query performance, see [“Performance Optimizations for Indexed Fields in Search Queries” on page 91.](#))

Field-based search is case sensitive. Make sure you understand and follow other requirements and guidelines listed in [“Syntax Reference for Query Expression” on page 81.](#)

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

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



Note

If a query includes the boolean operator OR and the metadata identifiers (discussed in [“Constraints” on page 79](#)), the expression to be evaluated with OR must be enclosed in parentheses, as shown in this example:

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

If the expression is not enclosed in parentheses, an error message is displayed on the user interface screen.

A complete list of fields you can specify is available in [“Indexing” on page 111](#) 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	
!=	<code>message!="failed login"</code> <code>message!=failed*login</code> (* means wildcard) <code>"test"</code> <code>message!=failed*login</code> (* is literal in this case)
=	<code>message="failed login"</code> <code>message="failed*login"</code> (* means wildcard)
	<p>Caution: The size of each field in the OLI schema is predetermined. If the string you are searching for is longer than the field-length, you should use a STARTSWITH rather than an = search, and include no more than the number of characters in the field size. To determine the size of a default field, see "Viewing Default Fields" on page 201. To determine the size of a custom field, see "Viewing Custom Fields" on page 201.</p>
>	<p>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.</p>
<	
>=	
<=	
BETWEEN	
IN	
STARTSWITH	<code>message STARTSWITH "failed"</code>
ENDSWITH	<code>message ENDSWITH "login"</code>
CONTAINS	<code>message CONTAINS "foobar"</code>
Numeric / Timestamp Operators	
=	<code>bytesIn = 32</code>
!=	<code>destinationPort != 100</code>
>	<code>bytesIn > 100</code>
>=	<code>endTime >="01/13/2009 07:07:21"</code> <code>endTime >="2009/13/01 00:00:00 PDT"</code> <code>endTime >="Sep 10 2009 00:00:00 PDT"</code>

Field Operator	Example
<	startTime < "\$now - 1d"
<=	startTime <= "\$now - 1d"
BETWEEN	priority BETWEEN 1 AND 5
SQL Operator	
IS	sessionId IS NULL sessionId IS NOT NULL
Boolean Operators	
AND	name="Data List" AND message="Hello" AND 1.2.3.4
OR	(name="TestEvent" OR message="Hello") AND type=2 AND 1.2.4.3
NOT	NOT name="test 123"
List Operator	
IN	priority IN [2,5,4,3] destinationAddress IN ["10.0.20.40", "209.128.98.147"] _deviceGroup IN ["DM1"] _storageGroup NOT IN ["Internal Event Storage Group", "SG1"] _peerLogger IN ["192.0.2.10", "192.0.2.11"]

Guidelines for Field-based Search Expressions

Follow these guidelines when specifying field-based search expressions:

- You can specify any predefined OLI schema field. For example, cat = /Monitor/CPU/Usage. For a complete list, see ["Indexing" on page 111](#).
- You can specify any custom schema field you have added. For example, SSN=333-333-3333. For more information about custom schema fields, see ["Adding or Importing Schema Fields" on page 221](#).
- You cannot specify user-defined fields created through a predefined or user-defined parser in the Indexed Search portion of a query. (The Indexed Search portion of a query is the expression before the first pipeline character.) For example, if the Apache Access Log parser creates a field called SourceHost, you cannot specify the following query expression:

```
SourceHost = "123.456.789"
```

A query expression (Indexed Search | Search Operators) is evaluated from left to right in a pipeline fashion. By design, a parser—predefined or user-defined—is applied to an event when the Search Operators are processed in a search query. Therefore, field creation when a parser is

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

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

```
SourceHost="123.456.789"
```

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

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

Or, if you want to search only the Apache Access Logs for SourceHost="123.456.789", you can specify this expression:

```
| where parser="Apache Access Log" and clientIP="123.456.789"
```

Additionally, you can also run a full-text (keyword) search on "123.456.789", as follows: "123.456.789" | where SourceHost="123.456.789"

Search Operators

The *Search Operators* enable you to further refine the data that matched the indexed search filter.

The search operator, `rex`, is useful for syslog events (raw or unstructured data) or if you want to extract information from a specific point in an event, such as the 15th character in an event.

Other operators such as `head`, `tail`, `top`, `rare`, `chart`, `sort`, `fields`, and `eval` are applied to the CEF fields you specify or the information you extract using the `rex` operator.

See ["" on page 269](#) for a list of search operators and examples of how to use them.

Time Range

An event is timestamped with the OLI receipt time when it is received on the OLI. **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 OLI 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 74](#). The <units> portion, if given, consists of an integer and a single, case-sensitive letter, as shown in [Table 4-2 on page 75](#).


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



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.

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

The predefined field set, Raw Event, displays the whole raw syslog event in a column called `rawEvent`, as shown in the following figure. The event is formatted to fit in the column.

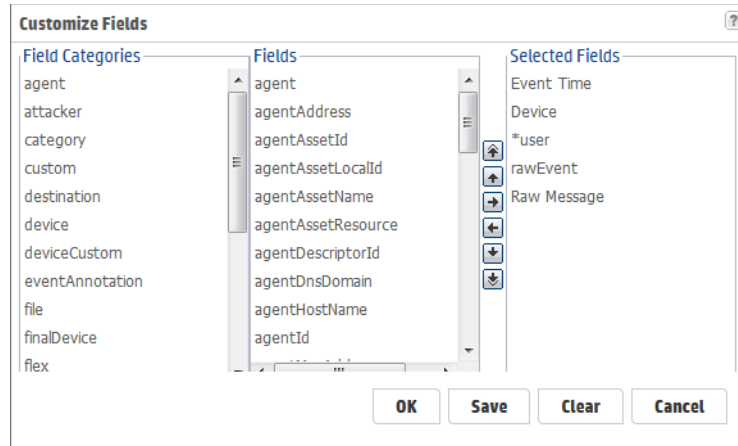
4 Searching and Analyzing Events

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


You must enable the Search Option, “Populate rawEvent field for syslog events”, to see the raw events in the rawEvent column. See [“Tuning Advanced Search Options” on page 197](#) for more information.

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

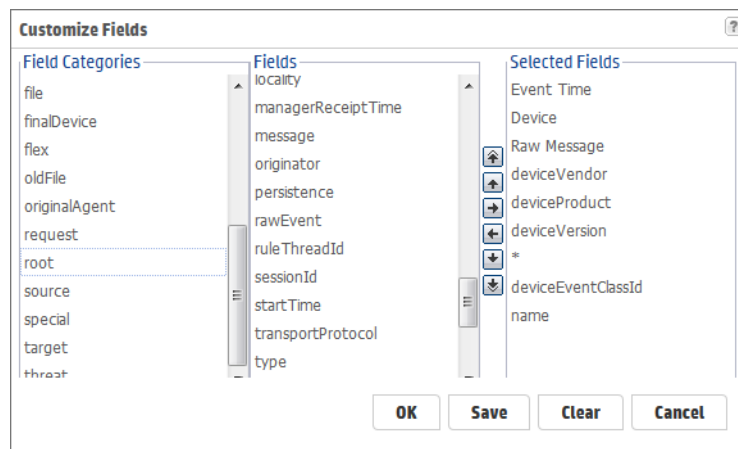
You can also create your own field sets by selecting “Customize...” from the “Fields:” drop-down menu. The OLI user interface offers a simple and intuitive interface to select and move event fields you want to include in a field set, as shown in the following figure.



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

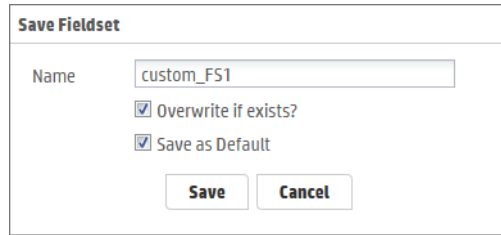


A wildcard field (“*”) is available in the Fields list when you create a custom field set. This field includes all fields available in an event that are not individually listed in the custom field set definition. For example, for the following custom field set definition, the search results will list 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. When saving a custom field set, you can specify a field set as the default for this OLI. If you do so, it is the default field set for all users on that system.

4 Searching and Analyzing Events



The 'Save Fieldset' dialog box contains the following elements:

- Name:
- Overwrite if exists?
- Save as Default
- Buttons: Save, Cancel

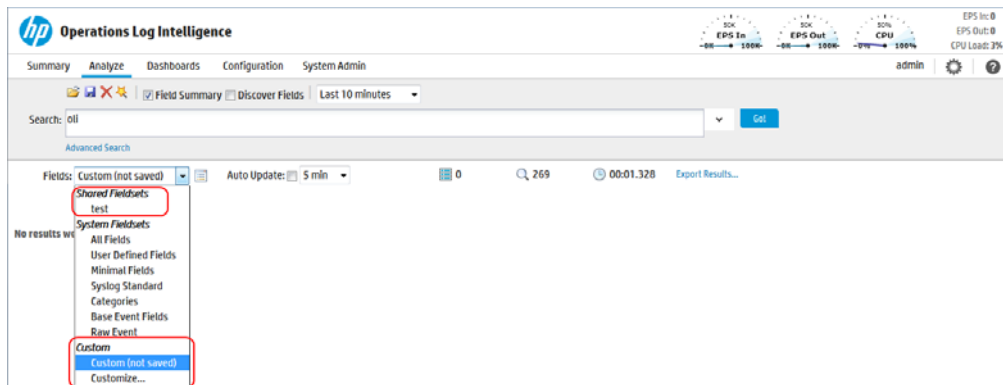
If you save a custom field set, it appears under the *Shared Fieldsets* category and is visible and available to the other users of your OLI, as shown in the following figure. Once a field set is saved, you can edit and delete it.

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 an OLI user’s interface. For example, UserA and UserB are connected to the same OLI 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 affect UserA’s display; UserB will continue to see the search results in the All Fields format.

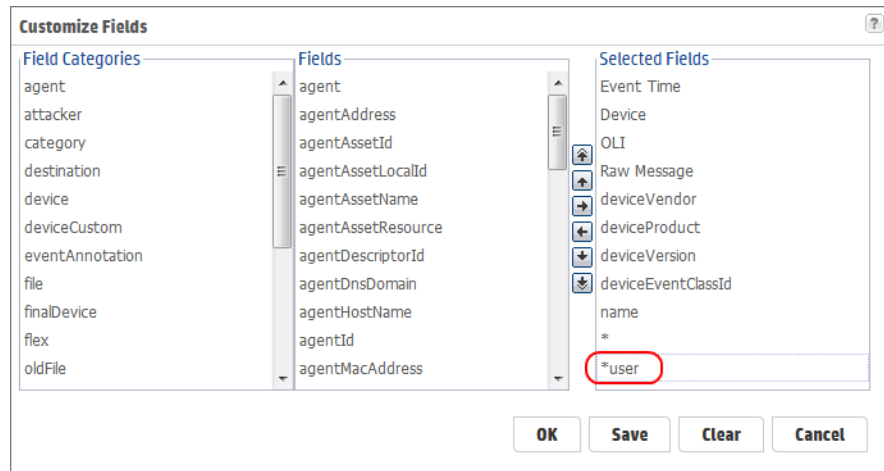


Field sets are not included in the saved filter definition.



For information about deleting custom field sets, see [“Viewing and Deleting Field Sets”](#) on page 200.

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



Constraints

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

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

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

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

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

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

- If a query includes the boolean operator OR and the metadata identifiers (discussed in [“Constraints” on page 79](#)), the expression to be evaluated with OR must be enclosed in parentheses, as shown in this example:

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

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

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

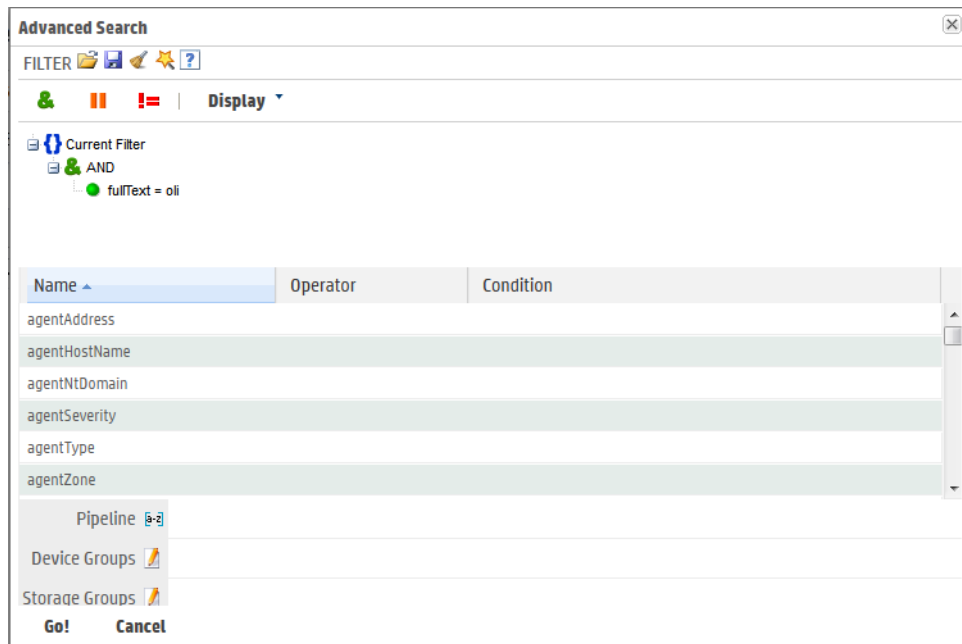
Use OLI's Search Helper to enter a constraint in the Search text box. Once you type `"_s"` (for storage group), `"_d"` (for device group), or `"_p"` (for OLI) in the Search text box, OLI automatically provides a drop-down list of relevant terms and operators from which you can select.



Caution

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

- ◆ Using the Search Builder tool as you can select the constraints in it, as shown in the following figure. (To access the Search Builder tool, click **Advanced** under the Search text box where you type query expression.) For more information about the Search Builder, see ["Using the Search Builder Tool" on page 86](#).



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.

Table 4-3 Query Syntax 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 111 section.) (List of operators in the "Field-based Search" on page 70 section.)	REGEX="<REGEX1>" REGEX="<REGEX2>" ..

Table 4-3 Query Syntax Requirements (Continued)

Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
Operators	<p>Upper-, lower-, or mixed case boolean operators—AND, OR, NOT. If an operator is not specified, AND is used.</p> <p>To search for literal operator AND, OR, NOT, in an event, enclose them in double quotes.</p> <p>Example: “AND”, “or”, “Not”</p> <p>Note: If a query includes the boolean operator OR and the metadata identifiers (discussed in “Constraints” on page 79), the expression to be evaluated with OR must be enclosed in parentheses, as shown in this example:</p> <pre>(success OR fail) _storageGroup IN ["Default Storage Group"]</pre>	<p>Use any operator listed in the “Field-based Search” on page 70 section.</p> <ul style="list-style-type: none"> Unless a value is enclosed between double quotes, a space between values is interpreted as an AND. For example, name=John Doe is interpreted as John AND Doe. If an operator is not specified between multiple field expressions, AND is used. To search for literal operator, enclose the operator in double quotes. Examples: <pre>message STARTSWITH="NOT" message="LOGIN DID NOT SUCCEED"</pre> If a query includes the boolean operator OR and the metadata identifiers (discussed in “Constraints” on page 79), the expression to be evaluated with OR must be enclosed in parentheses, as shown in this example: <pre>(success OR fail) _storageGroup IN ["Default Storage Group"]</pre> 	<p> and the operators described in “Time Range” on page 73.</p> <p>Use this operator to AND multiple regular expressions in one query expression.</p>

Table 4-3 Query Syntax Requirements (Continued)

Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
Nesting (including parenthetical clauses, such as (a OR b) AND c)	<p>Allowed</p> <ul style="list-style-type: none"> Use boolean operators to connect and nest keywords. Metadata identifiers (_storageGroup, _deviceGroup, and _peerLogger), but can only appear at the top level in a query expression). If the query contains a regular expression, the metadata identifiers need to precede the regular expression. 	<p>Allowed</p> <ul style="list-style-type: none"> Use any operator listed in the “Field-based Search” on page 70 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 expressions can be specified in one query using this syntax: REGEX=“<REGEX1>” REGEX=“<REGEX2>” ...</p>
Case sensitivity	<p>Insensitive (Cannot be changed.)</p>	<p>Sensitive* (Can be changed using Tuning options. See “Tuning Advanced Search Options” on page 197.)</p>	<p>Insensitive* (Can be changed using Tuning options. See “Tuning Advanced Search Options” on page 197.)</p>
Wildcard	<p>*</p> <p>Cannot be the leading character; only a suffix or in between a keyword.</p> <p>Examples:</p> <ul style="list-style-type: none"> *log is invalid log* is valid lo*g* is valid 	<p>*</p> <p>Can appear anywhere in the value.</p> <p>Examples: name=*log (searches for ablog, blog, and so on.) name=“*log” name=*log (both search for *log)</p>	<p>*</p> <p>Can appear anywhere.</p>

Table 4-3 Query Syntax Requirements (Continued)

Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
Exact Match/Search string includes an operator or a special character	Enclose keyword in double quotes; Otherwise, keyword treated as keyword*. Example: log (matches log, logging, logger, and so on) "log" (matches only log) Note: See the list of special characters that cannot be searched even when enclosed in double quotes, later in this table.	Enclose value in double quotes Example: message="failed login"	No special requirement.
Escape character	\ Use to escape \. You cannot escape any other character.	\ Use to escape \, ", and *. Examples: <ul style="list-style-type: none"> name=log\ger (matches log\ger) name=logger* (matches logger*) 	\ Use to escape any special character. Example: To search for a term with the character "[": REGEX="logger\[
Escaping wildcard character	Cannot search for * Example: log* is invalid	Can search for * by escaping the character name=log* is valid	Can search for * by escaping the character
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.

Table 4-3 Query Syntax Requirements (Continued)

Behavior	Full Text (Keyword)	Field Search (Indexed)	Regular Expression
Primary Delimiters: Space , ; () [] } " * > < !	You can search for keywords containing primary delimiters by enclosing the keywords in double quotes. Example: "John Doe" "Name=John Doe" "www.hp.com"	You can search for these characters. Enclose value in double quotes if value contains any of these characters. Example: name="John="	<ul style="list-style-type: none"> Cannot contain ^ in the beginning and \$ at the end as a matching character unless the regular expression you specify must look for an event that contains only the pattern you are specifying; for example, REGEX="^test\$" will search for events containing the word "test" (without quotes) only. Special regular expression characters such as \ and ? need to be escaped.
Secondary Delimiters: = . : / \ @ - ? # \$ & - %	You can also search for keywords containing secondary delimiters once you have configured the full-text search options as described in "Full-text Search Options" on page 199 . Example: You can search for hp.com in a URL http://www.hp.com/apps by specifying hp.com as the search string.	You can search for these characters. Enclose value in double quotes if value contains any of these characters. Example: name="John="	<ul style="list-style-type: none"> Cannot contain ^ in the beginning and \$ at the end as a matching character unless the regular expression you specify must look for an event that contains only the pattern you are specifying; for example, REGEX="^test\$" will search for events containing the word "test" (without quotes) only. Special regular expression characters such as \ and ? need to be escaped.

Table 4-3 Query Syntax Requirements (Continued)

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>"mm/dd/yyyy hh:mm:ss"</p> <p>OR</p> <p>"yyyy/mm/dd hh:mm:ss timezone"</p> <p>OR</p> <p>"MMM dd yyyy hh:mm:ss timezone"</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 OLI 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 79](#)). This section describes how to use the tool.

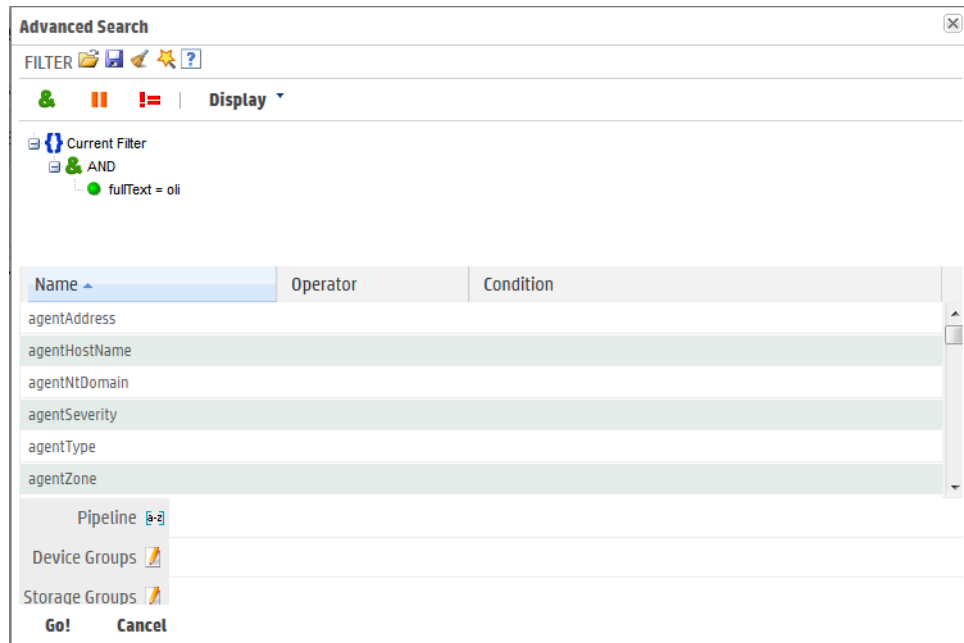
Accessing Search Builder

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**.

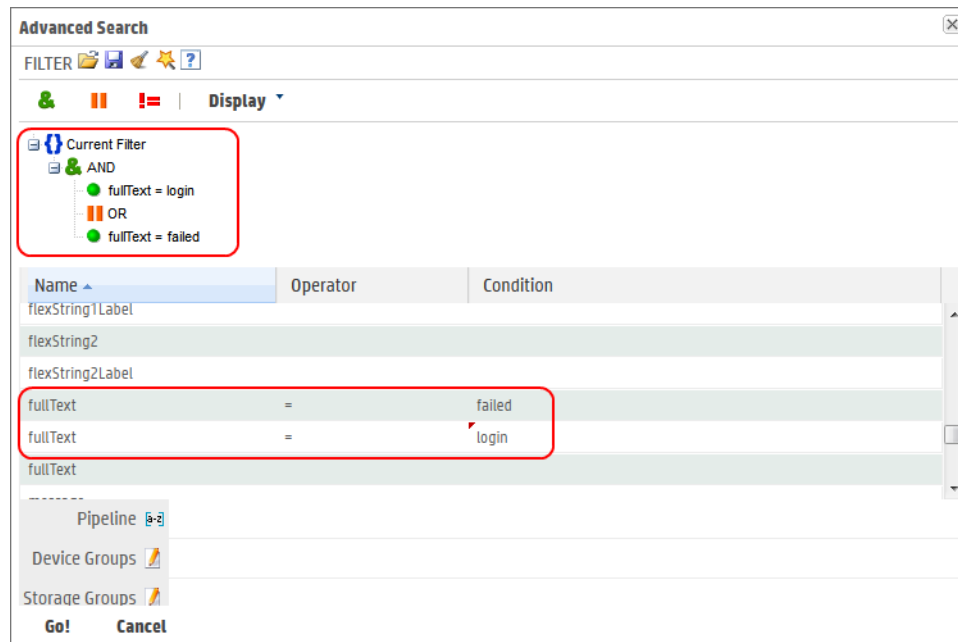
For more information, see [“Saving Queries \(Saved Filters and Searches\)”](#) on page 116 and [“System Filters/Predefined Filters”](#) on page 119.

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

4 Searching and Analyzing Events

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

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


The OLI constraint category is displayed only if OLIs are configured on your OLI.

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**.

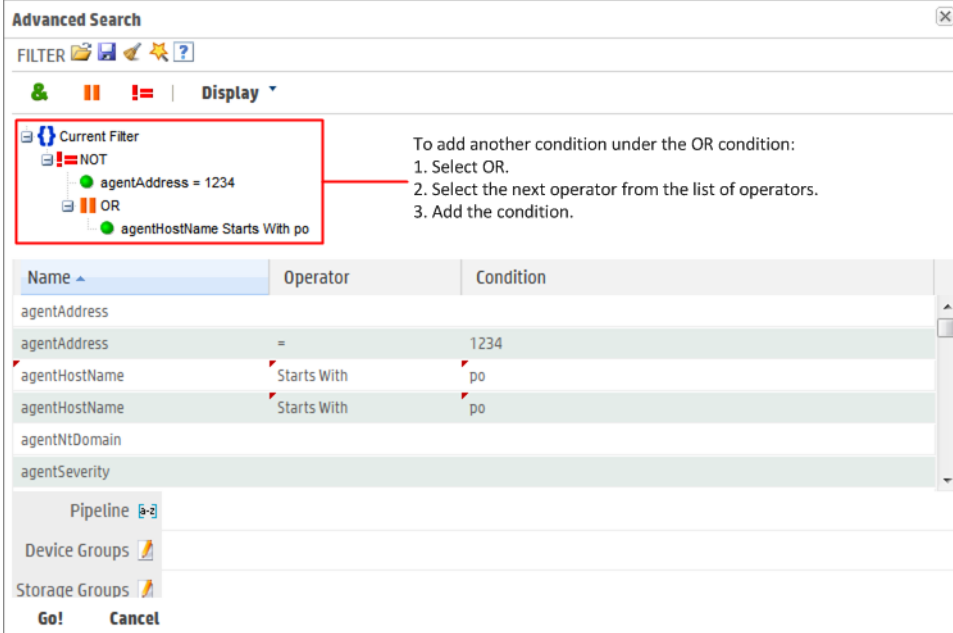
The query is automatically displayed in the Search text box and is ready to be run.

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 116](#).




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 86](#).


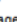
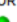



Advanced Search

FILTER 

   | Display ▾

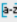
Current Filter


-  NOT
 -  agentAddress = 1234
 -  OR
 -  agentHostName Starts With po


To add another condition under the OR condition:

1. Select OR.
2. Select the next operator from the list of operators.
3. Add the condition.

Name	Operator	Condition
agentAddress		
agentAddress	=	1234
agentHostName	Starts With	po
agentHostName	Starts With	po
agentNtDomain		
agentSeverity		

Pipeline 

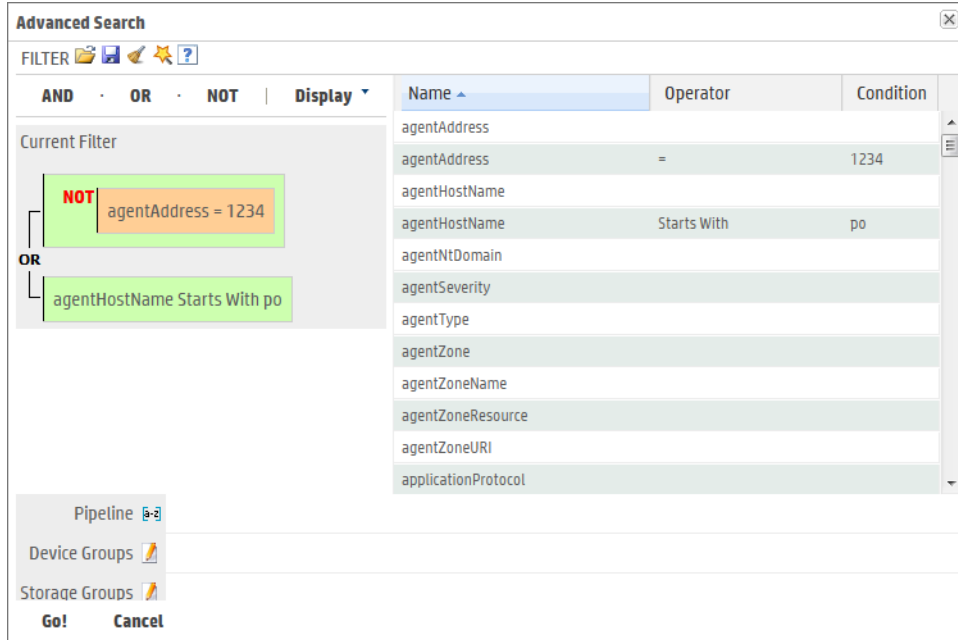
Device Groups 

Storage Groups 

Go! **Cancel**

Alternate Views for Query Building in Search Builder

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




To change views:

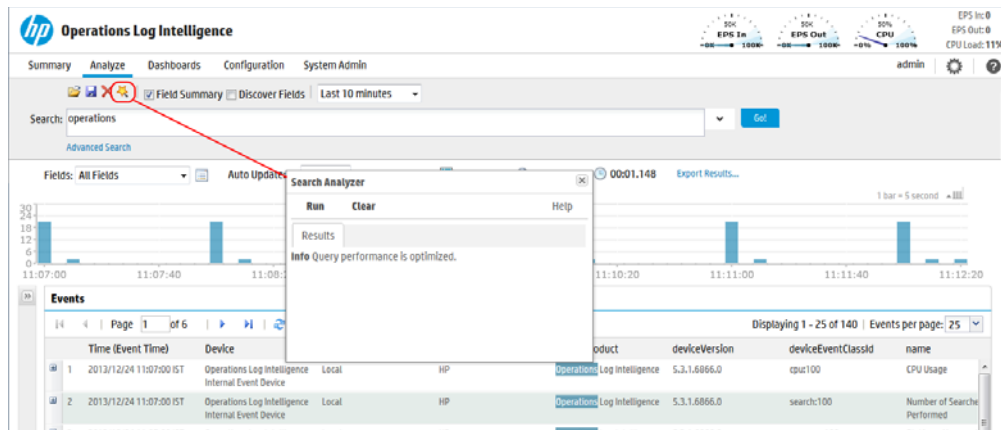
Click **Display** in the Search Builder tool and select the view of your choice.

Search Analyzer

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

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

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



Performance Optimizations for Indexed Fields in Search Queries

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

- When you include indexed and non-indexed fields in a query. Therefore, HP recommends that you identify the fields that you will most commonly use in queries and index all those fields.
- When you 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 OLI 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 OLI is not archived with events.

Regex Helper Tool

The Regex Helper tool enables you to create regular expressions that can be used with the `rex` pipeline operator to extract fields of interest from an event. (For information about `rex`, see [“Search Operators” on page 73](#) or [Appendix D, , on page 325](#).) This tool not only simplifies the task of creating regular expressions for the `rex` operator but also makes it efficient and error free.

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

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



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

- 1 Enter a search query that finds events of interest to you. (For information about running a search, see “Searching for Events on OLI” on page 96.)
- 2 Identify a syslog event that you want to analyze further. For example, in the shown figure, event #1 is the event we will analyze further.
- 3 Click the icon (in the left-most column) for the identified event to expand it and display its raw event.
- 4 Click the icon (next to the word **RAW**) to launch the Regex Helper tool.
- 5 Select the fields that you want to extract.
- 6 Click **OK**.

The screenshot shows the HP Operations Log Intelligence interface. At the top, a search bar contains the query 'tcp'. Below the search bar is a bar chart showing event counts over time. The main area displays a list of events. Event #1 is selected and expanded, showing its raw event text. The Regex Helper tool is open, displaying a list of fields to extract. The fields 'IPAddress_2' and 'Server' are selected. The 'OK' button is highlighted.

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

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

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

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

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

Search Helper

The screenshot shows the Search Helper interface. The search query is "oli deviceE". The interface displays a table of fields with the following columns: Field, Examples, and Help. The fields listed are deviceEventCategory, deviceEventClassId, and deviceExternalId. The Examples column shows "error alert" and "message CONTAINS 'Between' (name='John Doe' OR name='Jane Doe') AND message='success'". The Help section explains that the indexed search uses the 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. Below the Help section, there is a "Suggested Next Search Operators" section listing "cef, rex, extract, regex". At the bottom of the interface, the text "Auto-open is ON" is displayed.

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

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

Search Helper includes following types of information:

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

Autocomplete Search

Search:	ag
agent	4,654
agent:050	1,025
agentAddress	● Field
agentHostName	● Field
agentNtdomain	● Field
agentSeverity	Field
agentType	Field
agentZone	● Field
agentZoneName	● Field
agentZoneResource	● Field
agentZoneURI	● Field

Provides full-text keywords and OLI schema field suggestions based on the currently entered text in the Search box. The suggestions enable you to select a keyword, a field, a field value, a search operator, or a metadata term (`_storageGroup`, `_deviceGroup`, `_peerLogger`) from a list instead of typing it in, thus enabling you to quickly build a query expression.

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

If you type “[|” (the pipeline character), the list of operators available on OLI are displayed.

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

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

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

These guidelines describe the autocomplete feature behavior in detail:

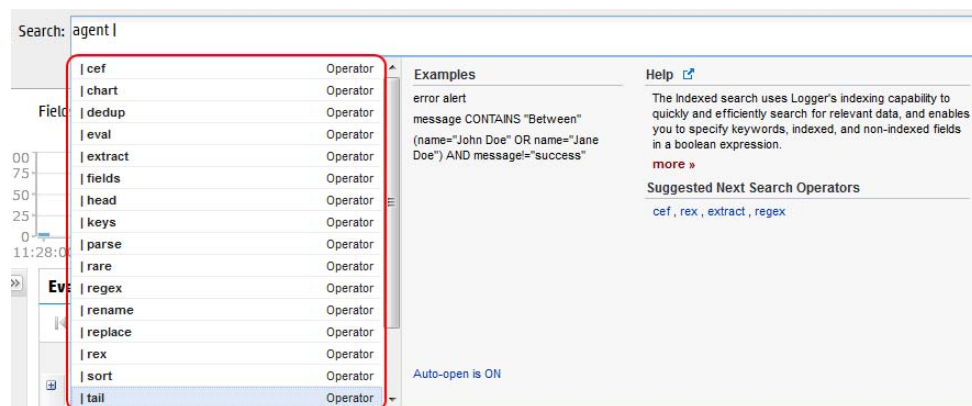
- The autocomplete suggestions and counts are based on the data stored on local OLI only. Peer OLI data is not included.
- The count next to each suggestion represents the total number of values stored on OLI for a field until the last time it was indexed. It is independent of the events that match the query specified in the Search box. For example, even if number of matching events for the search query on keyword “Error” are 5, the count for the keyword can be any other whole number greater than 5 because the count represents the total number of events containing the keyword “Error” while the search results are limited by elements of a search query such as time range, search constraints, and search operators (see [“Elements of a Search Query” on page 68](#)).

Additionally, the count shown is dependent on many other factors; therefore, it might not be exact in certain circumstances.

- Search Group filters (that restrict privileges on storage and device groups) are not enforced on the autocomplete list. Therefore, the list includes keywords, fields, field values, and counts of events in storage and device groups to which a user might not have privileges.
- When an archive is loaded back on OLI, the autocomplete list does not list the full-text keywords or field values that were available when the events were not archived. This happens because index data is not archived along with the event data; therefore, when the event data is loaded back from an archive, this data is treated as unindexed.

Search History

Displays the recently run queries on OLI, thus enabling you to select and reuse previously run queries without typing them again.



Search Operator History

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

Examples

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

Usage


Provides the syntax for the search operator.

Suggested Next Operators

List of operators that generally follow the currently typed query. For example, if you type `logger |`, the operators that often follow are `rex`, `extract`, or `regex`. You can select one of the listed operators to automatically append to the currently

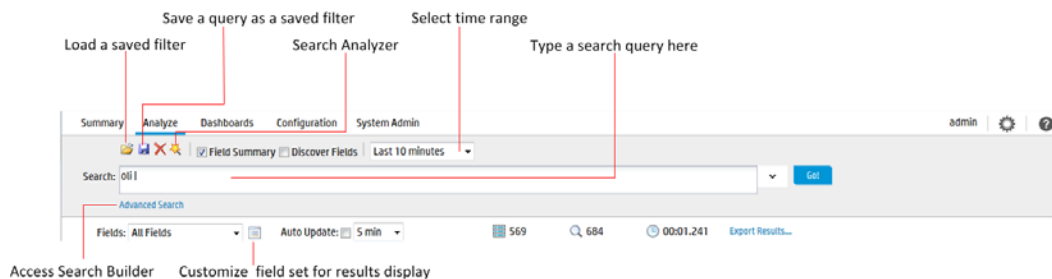
typed query in the Search text box. This list saves you from guessing the next possible operators and manually typing them in.

Help

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

Searching for Events on OLI

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



To search for events on OLI:

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



Refer to [“Query Expression” on page 69](#) for a list of exceptions and invalid characters before you create a query expression.

Note

- 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 68](#).

When you type a query, OLI’s Search Helper enables you to quickly build a query expression by automatically providing suggestions, possible matches, and applicable operators. See [“Search Helper” on page 93](#) for more information.

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


- For a complete list of fields in OLI schema, see [“Indexing” on page 111](#).
- Metadata terms (`_storageGroup`, `_deviceGroup`, `_peerLogger`)

Type `_s` (for storage group), `_d` (for device group), or `_p` (for OLI) in the Search text box to obtain a drop-down list of constraint terms and operators.

- Regular expression term (|REGEX=)




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

- Click **Advanced** to use the Search Builder tool. (See [“Using the Search Builder Tool” on page 86](#) for more information.) Also, use this option to specify device groups, storage groups, and OLI to which search should be limited.
- 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 116](#) and [“System Filters/Predefined Filters” on page 119](#).

- Use the following default values or change them suit your needs:
 - Local Logger:** By default the query is run on the local OLI only. If you want to run the query on the OLIs as well, uncheck the “Local Only” field located to the right of the Go! button.
 - Time Range:** By default, the query is run on the data received in the last two hours on the OLI. 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 73](#).
 - 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 75](#).
- 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 99](#).

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 116](#).

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 197](#).

Searching Peer OLIs (Distributed Search)

When you run a search query, by default, only your local OLI is searched for matching events. However, when specifying a query, you can select an option to

run the search on the peer OLIs. You can also select the OLIs to which the search should be constrained, as described in [“Searching for Events on OLI” on page 96](#).

Follow these guidelines for searching across peers:

- If OLIs do not have identical storage or device group names, a search query operation skips searching for events for those groups on those peers.
- If you added custom schema fields to your OLI schema, those fields must exist on all peers. Otherwise, a search query containing those fields will not run (when run across peers) and return an error. See [“Adding or Importing Schema Fields” on page 221](#).
- A user needs to belong to these user groups with the listed permissions set to perform searches and view their search results:
 - ◆ OLI Search Group with “Search for events on remote peers” user right set (checked).
 - ◆ OLI Rights Group with the “View registered peers” user rights set (checked).
- When an OLI becomes unavailable during a search operation, the one of the following errors might be displayed:

```
[OLI IP address] Error: Get Query Statistics  
[OLI IP address] Error: Remote exception (does not authorize the  
request. Please check if remote has relationship with your OLI)  
These error messages can occur when the OLI cannot be reached. Restore  
the relationship and run the search again.
```

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

Tuning Search Performance

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

- Enable field-based indexing for all fields that occur in your events. When events are indexed, OLI can quickly and efficiently search for relevant data. By default, a recommended set of fields are indexed on your OLI; you might need to add additional fields, as described in [“Indexing” on page 111](#).
- The amount of time it takes to search depends on the size of the data set that needs to be searched through, the complexity of the query, and whether the search is distributed across peers. To limit the data set, ensure that time range within which the events must be searched does not result in a query that needs to scan multi-millions of events. Additionally, limiting search to specific storage groups typically results in better search performance than when the storage groups are not specified.
- Reduce other load on the system when your query needs to run, such as scheduled jobs, large number of incoming events.


Understanding the Search Results Display

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





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



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

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

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

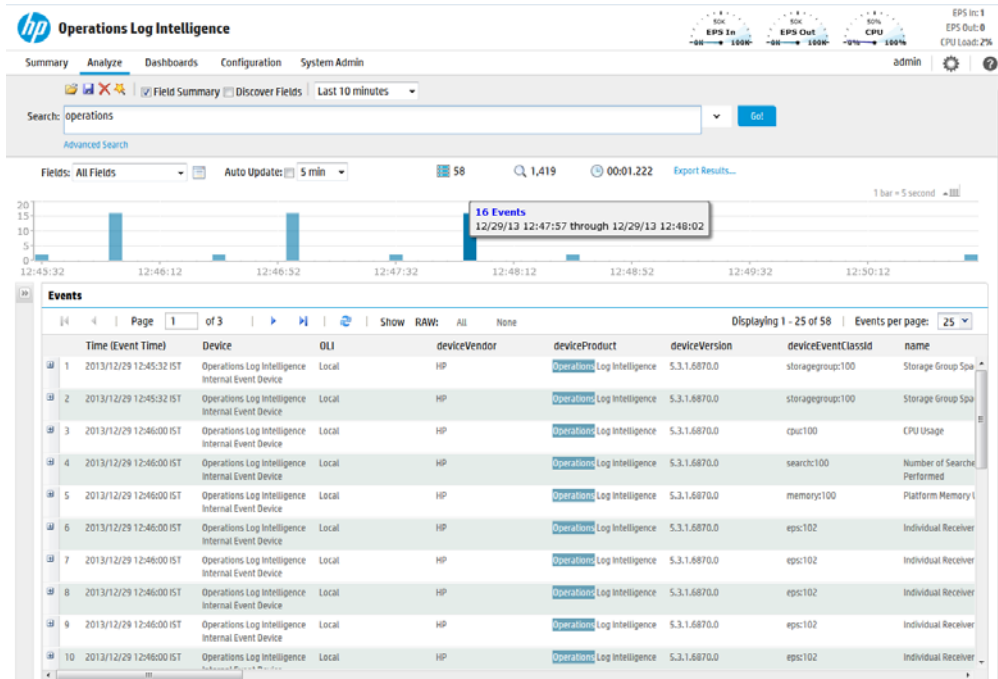
To see the next screen of events, click ; or  to go to the last page. Once you are past the first screen of events, you can click  to go back to the previous screen; or  to go to the first page.

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

The Search Results page displays a histogram that provides a graphical representation of the events that match a search query. The distribution is based on the time range specified in the query. That is, the X-axis represents event time and Y-axis represents the number of matching events, as shown in the following figure. You can randomly drill-down to events in a specific time period by clicking the histogram bar representing the time period.

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

4 Searching and Analyzing Events



Events are shown in table form, one row per event. Terms that match your query are highlighted in blue to make it easy to see why an event matched the query. To view the raw event of a listed event, click the icon to the left of the matching event. You can also view the Syslog raw events in a formatted column called `rawEvent` if you have enabled the “Populate rawEvent field for syslog events” option on the Search Options page, as discussed in [“Tuning Advanced Search Options” on page 197](#). Also, see [“Field Set” on page 75](#) to learn more about the `rawEvent` field.

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

- **Ctrl** + click the term in search results to add to search query
Adds the selected term to the search query, and reruns the search.
- **Alt** or **Shift** + click the term in search results to add to search query
Adds a NOT to the term, and reruns the query thus eliminating the events that match the term you selected. You can add multiple NOT conditions by holding

the Alt key and selecting terms in search results. When multiple conditions are added, they are joined by AND operators.

A Field Summary panel is displayed on the left side of the matched events. This section lists the fields that occur in matching events and the number of unique values for each in those events. For more information about Field Summary, see [“Understanding Field Summary” on page 105](#).

User-defined Fields in Search Results

When a search query matches events that were received from a defined source type and were parsed using a pre-defined or user-defined parser, the search results include a parser field, and may include fields for the source type, and source, depending on the setting in the Search Options tab. For more information, see [“Tuning Advanced Search Options” on page 197](#).

The following table describes the purpose of these fields.

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

User-defined fields are also created when a search query includes operators such as `rex`, `extract`, and `rename`. See [Appendix A, , on page 269](#) for information on these operators.

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

Viewing Search Results using Field Sets

By default, the Search Results are displayed using the All Fields field set, which displays all fields contained in an event. Once you select another field set, it



becomes your default view until you change it the next time. For a detailed discussion about field sets, see [“Field Set” on page 75](#).

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

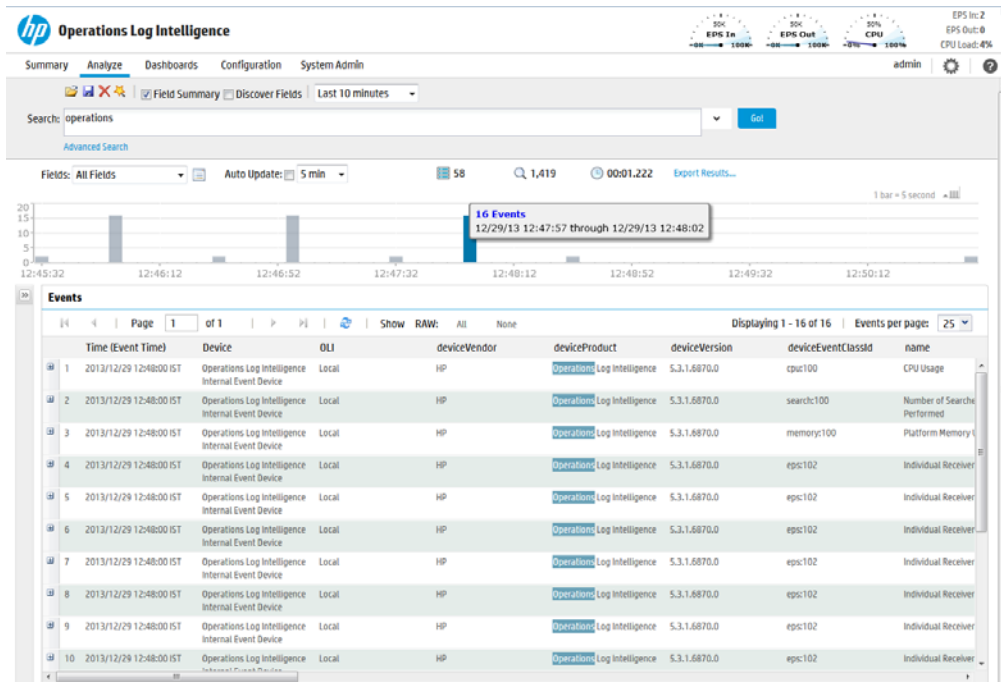
- Even though the rawEvent column displays the raw event, this column is not added to the OLI database and is not indexed. Therefore, you can only run a keyword (full-text) or regular expression to search on the event.
- You can use the Regex Helper tool to identify strings from the raw syslog events in the rawEvent column that you want to add to a query. (You cannot use the Regex Helper for CEF events displayed in the rawEvent column.) See [“Regex Helper Tool” on page 91](#) for details about the Regex Helper tool.

Using the Histogram

Use the following guidelines to effectively and efficiently use histograms:

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

- ◆ To deselect a selected bar, click it.



- A histogram is progressively built and displayed as events match a search query. If the search query needs to scan a large amount of data or a large time period, the histogram displayed initially might refresh multiple times while the query is running. To view the complete (and final) histogram of a search query, wait until the query has finished running (that is, the screen does not display the circular “waiting” icon anymore).
- The time range on the X-axis might not match the time range specified in the search query because the start and end times on the X-axis are determined by the event times of the first and last matching events of the search query.
- The first one million matching events are plotted on the histogram. If a search query matches more than one million events, an informational message is displayed on the screen.

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

Multi-line Data Display

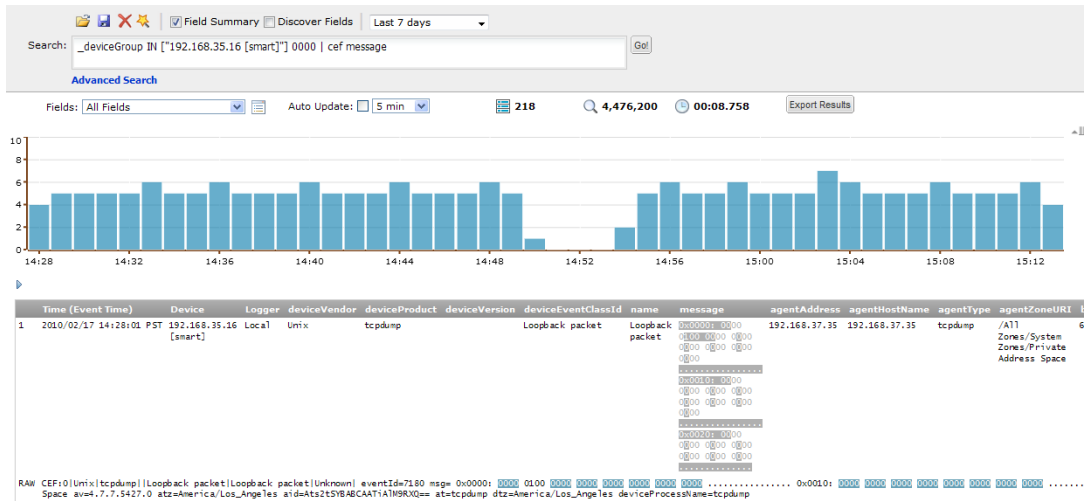
An event message might span multiple lines separated by characters such as newline (\n) or carriage return (\r). For example,

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

```

4 Searching and Analyzing Events

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



Auto Updating Search Results

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

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

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

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

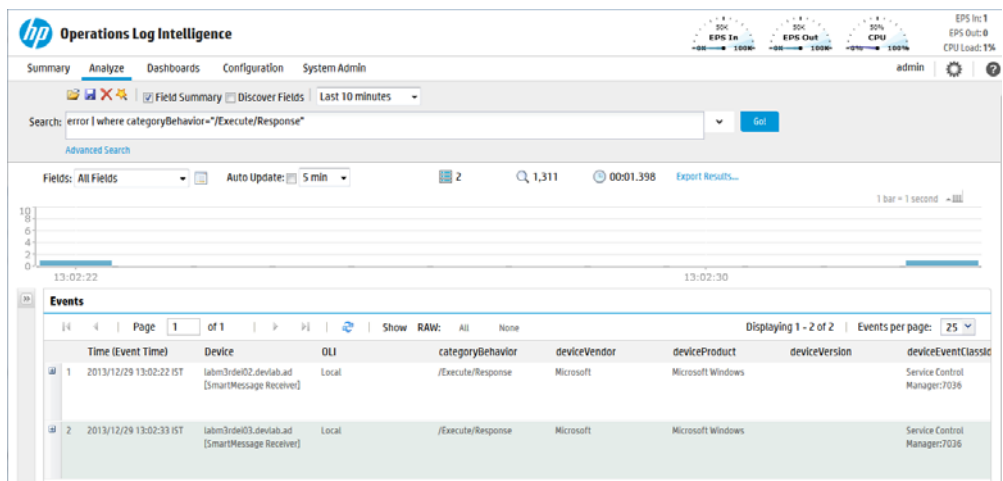
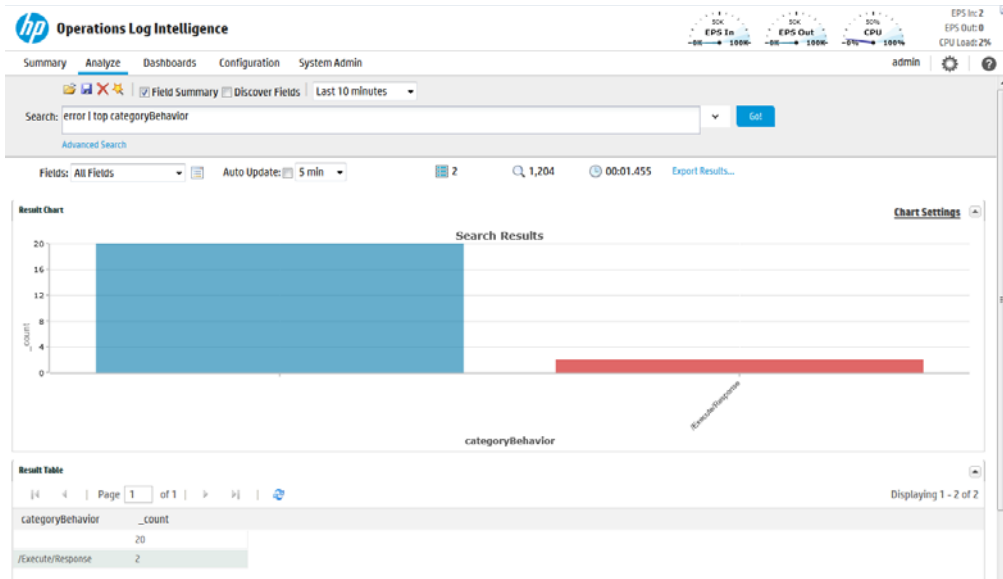
To auto update search results:

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

Chart Drill Down

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

/Execute/Response, click the second column to display events shown in the second figure.



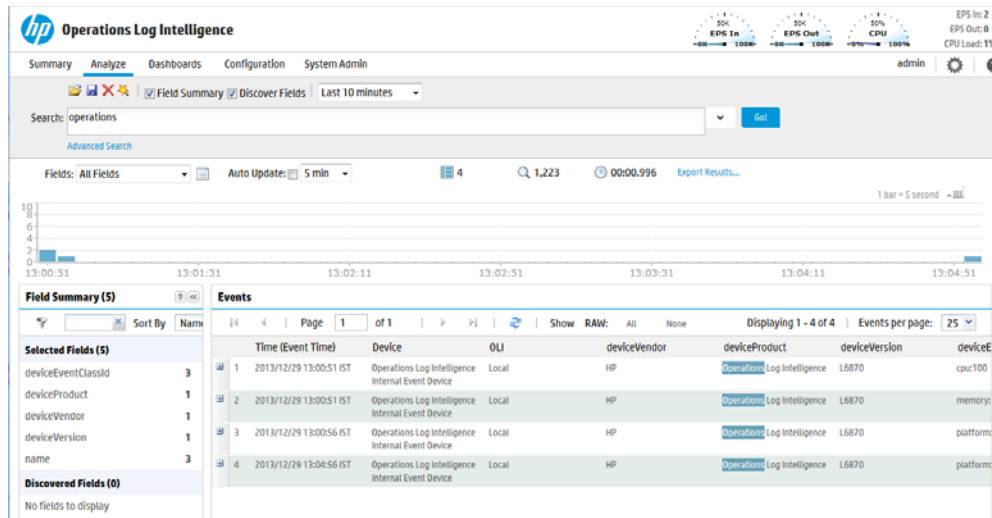
When you click on a chart value (a column, bar, or pie section), the existing search query is modified to include the WHERE operator with the field name and value, and automatically rerun. If you need to return to the original query from the drill-down screen, use the Back function of your browser.

Understanding Field Summary

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

4 Searching and Analyzing Events

The Field Summary panel contains two sections: Selected Fields and Discovered Fields. The Selected Fields section lists the CEF fields, while the Discovered Fields section lists the non-CEF fields discovered in raw events (described later in this section). By default, the Selected Fields list contains these fields: deviceEventClassId, deviceProduct, deviceVendor, deviceVersion, and name; you can edit this list to suit your needs, as described in [“To change the default Selected Fields list:” on page 107](#). For both lists, by default, the top 10 values for each field are listed.



The screenshot shows the HP Operations Log Intelligence interface. At the top, there are navigation tabs: Summary, Analyze, Dashboards, Configuration, and System Admin. The current view is 'Field Summary' with 'Discover Fields' checked. A search box contains 'operations' and shows '1,223' results. Below the search, there are controls for 'Fields: All Fields', 'Auto Update: 5 min', and 'Export Results...'. The main area is divided into two sections: 'Field Summary (5)' and 'Events'. The 'Field Summary' section lists 'Selected Fields (5)' and 'Discovered Fields (0)'. The 'Events' section displays a table with 4 rows of event data.

Time (Event Time)	Device	OLI	deviceVendor	deviceProduct	deviceVersion	deviceEventClassId
1 2013/12/29 13:00:51 IST	Operations Log Intelligence Internal Event Device	Local	HP	Operations Log Intelligence	1.6870	cpu:100
2 2013/12/29 13:00:51 IST	Operations Log Intelligence Internal Event Device	Local	HP	Operations Log Intelligence	1.6870	memory:
3 2013/12/29 13:00:56 IST	Operations Log Intelligence Internal Event Device	Local	HP	Operations Log Intelligence	1.6870	platform:
4 2013/12/29 13:04:56 IST	Operations Log Intelligence Internal Event Device	Local	HP	Operations Log Intelligence	1.6870	platform:

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

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



The screenshot shows the 'Field Summary Options' configuration page. It contains two dropdown menus: 'Use Field Summary' is set to 'Yes' and 'Discover fields' is set to 'No'.

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

To auto discover fields, the raw event must contain data in the “key=value” format, and none of these characters can be the first character of the “value”: comma, space, tab and semicolon. For each “key=value” pair found in a raw event, a new

field of the name “key” is created. The Field Summary includes a summary of the values for all the new fields under the Discovered Fields section. The discovered fields are assigned the type “String” by default. The auto-discovery capability works only if at least 2,500 of the first 10,000 matching events contain “key=value” pairs. If this threshold is not met, auto discovery is automatically turned off. However, this threshold does not apply if there are less than 10,000 matching events; in that case, fields are discovered regardless.

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

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


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

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

The screenshot shows a window titled "Field Summary (5)". At the top, there is a search filter text box with an "x" icon and a "Sort By" dropdown menu currently set to "Name". Below this, the "Selected Fields (5)" section contains a table with the following data:

Field Name	Count
deviceEventClassId	14
deviceProduct	1
deviceVendor	1
deviceVersion	2
name	14

Below the selected fields, the "Discovered Fields (0)" section is empty, displaying "No fields to display".


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

To change the default Selected Fields list:

- 1 Define or update an existing custom field set to include fields you want the Selected Fields list to contain. See [“Field Set” on page 75](#) for information on creating custom field sets.
- 2 Select the custom field set you defined to view search results.

After running a search query, if you select a different field set, the Field Summary panel displays the following message:

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



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

Exporting Search Results

You can export search results in these formats:



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

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

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

Events Type and Category

[Start:26/Mar/2010 9:47:48 -0700 End:26/Mar/2010 9:57:48 -0700]

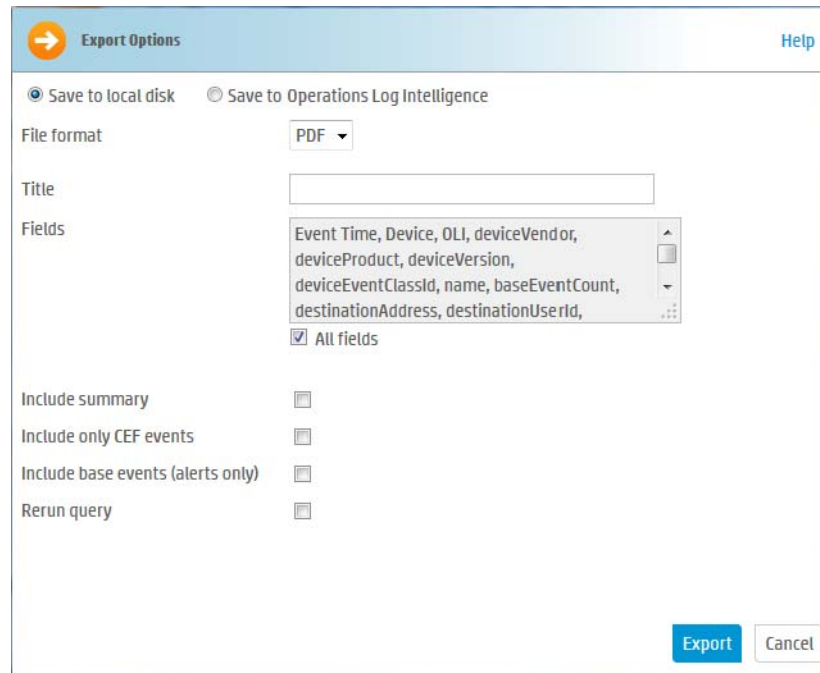


To export search results:

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

Option	Description
Save to local disk	The file is saved to a local system from which you are accessing OLI or is it sent to the browser for viewing or saving.
Save to Operations Log Intelligence	The file is written to the OLI's local storage.

Option	Description
File Format	<p>CSV, for comma-separated values file.</p> <p>PDF, for a report-style file that contains search results as charts and in tables. Charts are only included in the PDF file if the search query contains an operator that creates charts, such as <code>chart</code>, <code>top</code>, and so on.</p>
Title	<p>(Optional, available only when the File Format is “PDF”)</p> <p>A meaningful name that appears on top of the PDF file. If no title is specified, “Untitled” is included.</p>
Fields	<p>A list of event fields that will be included in the exported file.</p> <p>By default, all fields are included.</p> <p>You can enter fields or edit the displayed fields by deselecting All Fields.</p> <p>To export fields created as a result of <code>rex</code>, <code>extract</code>, <code>rename</code>, or <code>eval</code> operators, or field created when a parser is applied to an event, ensure that <code>*user</code> is selected in the Fields list.</p>
Chart Type (for PDF only)	<p>(Available only when a chart is available in search results)</p> <p>Type of chart to include in the PDF file. You can select from: Column, Bar, Pie, Area, Line, Stacked Column, Stacked Bar.</p> <p>Note: If the Chart Type is different from the chart displayed on the Search Results screen, the value selected for this option overrides the one shown in the screen. Therefore, the exported PDF contains the chart you specify for this option and not the one shown on the screen.</p>



Option	Description
Chart Result Limit (for PDF only)	(Available only when a chart is available in search results) Number of unique values to plot. Default: 10 If the configured Chart Result Limit is less than the number of unique values for a query, the top values equal to the Chart Result Limit are plotted. That is, if the Chart Result Limit is 5 and 7 unique values are found, the top 5 values will be plotted.
Include Summary	Include an event count in the exported search results.
Include Only CEF Events	Only include CEF events in the exported search results.
Include Base Events	Include base events in the exported search results.
Rerun Query	Reruns the search.

- 4 Click **Export**.

Scheduling an Export Operation

The time it takes to export search results is proportional to the number of events being exported. Therefore, for a large number of events, HP recommends that you schedule the export operation to be performed at a later time by saving the query and time parameters as a Saved Search, and then scheduling a Saved Search Job. For more information about Saved Search Jobs, see [“Scheduled Saved Search” on page 193](#).

Indexing

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

- Full-text indexing—Each event is tokenized and indexed. See [“Full-text Indexing \(Keyword Indexing\)” on page 112](#).
- Field-based indexing—Event fields are indexed based on a predetermined schema. See [“Field-based Indexing” on page 112](#).

How indexing works

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

All events are timestamped with the receipt time when received on the OLI. The default fields are automatically indexed. For the remaining fields, OLI uses the receipt time of an event and the time when a field was added to the index to determine whether that event will be indexed. If the receipt time of the event is

equal to or later than the time when the field was added to the index, the event is indexed; otherwise, it is not.



Indexing information is not archived.

Full-text Indexing (Keyword Indexing)

For full-text indexing, each event (CEF or non-CEF) received on OLI is scanned and divided into keywords and stored on the OLI. The full-text search options control the manner in which an event is tokenized as described in [“Full-text Search Options” on page 199](#).

Full-text indexing is automatically enabled at OLI initialization time. You cannot disable it. For details about enabling full-text indexing, see [“Enabling Indexing” on page 115](#).

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 OLI’s field search method utilize these indexed fields to yield significant search performance gains.

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

A list of the default index fields, along with their field descriptions is available from the OLI Configuration menu. For instructions on how to view the default OLI Schema fields, see [“Viewing Default Fields” on page 201](#).



- HP strongly recommends that you index fields that you will be using in search queries.
 - The `requestUrl` field is available for search queries; however, this field cannot be indexed.
 - The fields created when a predefined or user-defined rex parser parses the non-CEF events cannot be indexed using the field-based indexing capability. See [“Parsers” on page 159](#) for more information about rex parsers.
-

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

The following fields are available for indexing. The fields that OLI starts indexing automatically after OLI initialization are indicated in **bold** font. In addition to the

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

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

Index Fields		
destinationServiceName	deviceReceiptTime	sourceTranslatedAddress
destinationTranslatedAddress	deviceSeverity	sourceUserId
destinationUserPrivileges	deviceVendor	sourceUserName
destinationUserId	deviceVersion	sourceUserPrivileges
destinationUserName	deviceZone	sourceZone
destinationZone	deviceZoneName	sourceZoneName
destinationZoneName	deviceZoneResource	sourcezoneResource
destinationZoneResource	deviceZoneURI	sourceZoneURI
destinationZoneURI	endTime	startTime
deviceAction	eventId	transportProtocol
deviceAddress	externalId	type
deviceCustomDate1	fileName	vulnerabilityExternalID
deviceCustomDate1Label		VulnerabilityURI

Guidelines for Field-based Indexing

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

- Events are indexed by the fields in the “Indexed fields” list (on the Search Indexes page) and the default event metadata fields—event time, OLI receipt time, and device address.
- You can index up to 123 fields on OLI. This number includes the custom schema fields you may have added to your OLI.
- 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, OLI might not immediately start indexing on that field. Therefore, allow some time between adding a field and using it in the search query. If OLI is in the process of indexing on a field and you use that field to run a search query, the search performance for that operation will be slower than expected.
- If an event field contains data of unexpected type (for example, a string when an integer is expected), the data is ignored. Therefore, search for that data value will not yield any results. For example, if the port field contains a value 8080A (alphanumeric) instead of 8080 (numeric), the alphanumeric value is ignored.
- For 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 an OLI but not on its peers for a specific time range, a distributed search

will run slower on the OLIs. However, it will run at optimal speed on the local OLI. Therefore, the search performance in such a setup will be slow.

- Although the `requestUrl` field is available for search queries, it cannot be indexed. Including this field in such queries will result in the query running slower than a search performed on indexed data.

Enabling Indexing

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

Adding Fields to Field-based Index

To add fields to the field-based index:

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

Search Indexes | Search Options | Fieldsets | Default Fields | Custom Fields | Running Tasks | Parsers

Edit Search Index

Important

Once a field is indexed or full text indexing enabled, it cannot be changed. Significantly exceeding ArcSight's default recommended indexed fields could result in performance degradation in certain situations. If you need to exceed the default number of fields, only index those additional fields which are necessary for your environment.

To add indexed fields, select one or more fields below

Indexable fields

- agentAddress
- agentHostName
- agentNIDomain
- agentZoneURL
- customerName
- destinationDnsDomain
- destinationMacAddress
- destinationTranslatedAddress
- destinationUserPrivileges
- deviceCustomDate1
- deviceCustomDate2
- deviceCustomDate2Label
- deviceCustomNumber1Label
- deviceCustomNumber2Label
- deviceCustomNumber3Label
- deviceCustomString1Label
- deviceCustomString2Label
- deviceCustomString3Label
- deviceCustomString4Label
- deviceCustomString5Label
- deviceCustomString6Label
- deviceExternalId
- externalId
- flexDate1
- flexDate1Label

Use ctrl-click to select or deselect items

Indexed fields

- deviceVendor
- deviceProduct
- deviceVersion
- deviceEventClassId
- name
- agentSeverity
- agentType
- applicationProtocol
- baseEventCount
- bytesIn

All recommended fields have already been indexed

Full text indexing is enabled

Apply Changes

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

Saving Queries (Saved Filters and Searches)

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

- As a filter

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

- As a saved search

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

For information about Saved Search Alerts, see [“Creating and Managing Saved Search Alerts” on page 179](#).

Saving a Query

To save a query:

- 1 Define a query as described in [“Searching for Events on OLI” on page 96](#) or [“Using the Search Builder Tool” on page 86](#).
- 2 Click the Save icon (📁) and enter a name for the query in the Name field, as shown in the following figure.

- 3 In the Save as field, select whether you want to save this query as a filter, as a saved search, or as a Dashboard panel.

If you select to **save as a Saved Search**, you can either keep the saved query as Saved Search or change it to a Scheduled Alert by specifying a schedule based on which the query runs periodically and generates alerts.

If you choose to **schedule the Saved Search**, you can either specify the schedule in the following screens or skip it for now.

If the search query includes an aggregation operator such as chart or top, a third option to save the query for a **Dashboard panel** is also displayed. If you select this option, you need to enter the following parameters.

Save query
Help

Panel Title

Save as Filter Saved search Dashboard panel

Saved search New saved search
 Unix - Configuration - Configuration
 Dikla test

Saved search name

Dashboard Dikla test
 New dashboard

Panel type or Add both types

Chart type

Chart limit

Parameter	Description
Title	Enter a meaningful name for the panel that will be added to the Dashboard.
Saved search	Select an existing saved search from the drop-down box that will be overwritten with this query. OR Select "New saved search" to create a new saved search query. Enter the new name in the text box.
Dashboard	Select an existing Dashboard from the drop-down box to which the Search Results panel will be added. OR Select "New dashboard" to add the Search Results panel to a new Dashboard. Enter the name of the new Dashboard in the "Dashboard Name" field.
Panel type	Select the type of panel: <ul style="list-style-type: none"> • Chart—Displays search results in a chart form • Table—Displays search results in a table form • Chart and Table—Adds two panels, one for displaying search results in the chart form and the other for displaying search results in the table form

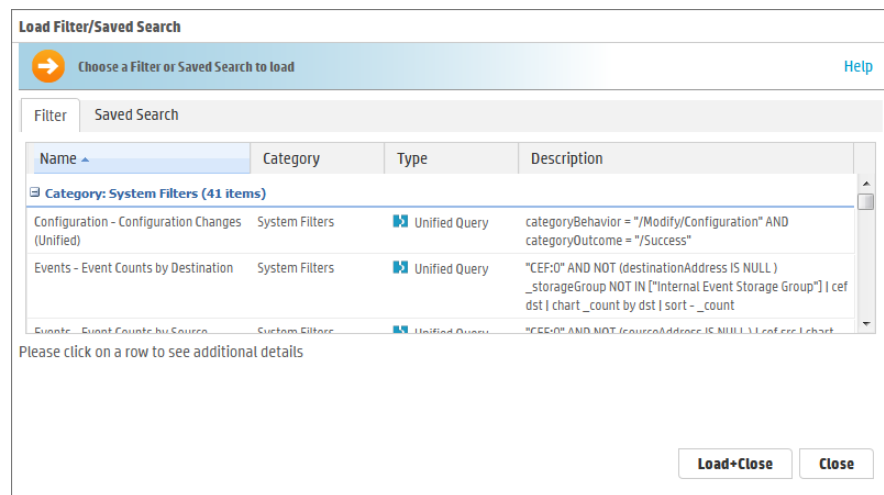
Parameter	Description
Chart type	Type of chart to display matching events. You can select from: Column, Bar, Pie, Area, Line, Stacked Column, Stacked Bar. Default: Column
Chart limit	<i>Only applicable to Search Result Chart panels.</i> Number of unique values to plot. Default: 10

- 4 Click **Save**.

Using a Saved Filter or a Saved Search

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

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



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

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

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

System Filters/Predefined Filters

Your OLI ships with a number of predefined filters, also known as system filters. These filters define queries for commonly searched events. For example, unsuccessful login attempts or the number of events by source. Filter queries are available as Unified queries and as Regular Expression queries. Unified queries can be used for searching while Regular Expression queries are for defining alerts and forwarders.



Note

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

The following is a list of all the system filters. To use a predefined system filter, follow instructions in [“Using a Saved Filter or a Saved Search” on page 118](#).

Table 4-4 System Filters

Category	Unified Query Filters	Regular Expression Query Filters
Login Status use case	All Logins	All Logins (Non-CEF) All Logins (CEF format)
	Successful Logins	Successful Logins (Non-CEF) Successful Logins (CEF format)
	Failed Logins	
Configuration	Configuration Changes	System configuration changes (CEF format)
Events use case	High and Very High Severity Events	High and Very High Severity CEF events
	Event Counts by Source	
	Event Counts by Destination	All CEF events
Network use case	DHCP Lease Events	
	Port Links Up and Down	

Table 4-4 System Filters (Continued)

Category	Unified Query Filters	Regular Expression Query Filters
Connector System Status use case	Protocol Links Up and Down CPU Utilization by Connector Host Disk Utilization by Connector Host	
UNIX Server use case	Memory Utilization by Connector Host CRON related events	
Windows Events use case	IO Errors and Warnings Password Changes SAMBA Events SSH Authentications User and Group Additions User and Group Deletions Account Added to Global Group Account Added to Global Group (CEF) Audit Policy Change Audit Policy Change (CEF) Change Password Attempt Change Password Attempt (CEF) Global Group Created Global Group Created (CEF) Windows Events (CEF)	

Table 4-4 System Filters (Continued)

Category	Unified Query Filters	Regular Expression Query Filters
System Alerts	<p>The following filters search for specific internal alert events, which are written in CEF format to a special Internal Storage Group. These filters are available for both search methods. In addition to the following filters, you can define your own alerts based on the system health events listed in “System Health Events” on page 267.</p> <p>NOTE: Although these filters are displayed on the OLI, these do not apply to it.</p>	
	CPU Utilization Above 90 Percent	CPU Utilization Above 90 Percent
	CPU Utilization Above 95 Percent	CPU Utilization Above 95 Percent
	Disk Failure	Disk Failure
	Root Partition Below 10 Percent	Root Partition Below 10 Percent
	Root Partition Below 5 Percent	Root Partition Below 5 Percent
	Device Configuration Changes	Device Configuration Changes
	Filter Configuration Changes	Filter Configuration Changes
	High CPU Temperature	High CPU Temperature Bad Fan
	Power Supply Failure	Power Supply Failure
	RAID Controller Issue	RAID Controller Issue
	RAID Status Battery Failure	RAID Status Battery Failure
	RAID Status Disk Failure	RAID Status Disk Failure
	Storage Configuration Changes	Storage Configuration Changes
	Storage Group Usage Above 90%	Storage Group Usage Above 90%
	Storage Group Usage Above 95%	Storage Group Usage Above 95%
	Zero Events Incoming	Zero Events Incoming
	Zero Events Outgoing	Zero Events Outgoing

Using a System Filter

To use a predefined system filter, follow instructions in [“Using a Saved Filter or a Saved Search” on page 118](#).

Alerts

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

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

Audit events for alerts are only written to the Internal Storage Group.



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

Viewing Alerts

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

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

When you create Alerts (see [“Alerts” on page 172](#)), 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 73](#) for more detail.

Receiving Alerts for Events

To receive alerts:

- 1 Configure the OLI’s SMTP with the desired e-mail address destination (see [“SMTP” on page 236](#)) or create an SNMP Destination (see [“Sending Notifications to SNMP Destinations” on page 185](#)) or Syslog Destination (see [“Sending Notifications to Syslog Destinations” on page 186](#)).



Number of destinations per alert:

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

- 2 Create a query to find the events of interest; save the query as a filter. (See [“Saving Queries \(Saved Filters and Searches\)” on page 116](#).)
- 3 Create an Alert that uses the new filter and specify match count and threshold (see [“Alerts” on page 172](#).) Enable the new Alert.

The screenshot displays the HP Operations Log Intelligence interface. At the top, there are navigation tabs: Summary, Analyze, Dashboards, Configuration, and System Admin. The 'Analyze' tab is active. Below the navigation, there are several controls: 'Show: All Alerts', 'Within: Last 24 hours', 'Base Event Fields: All Fields', 'Export Results...', and 'Auto Update: 5 min'. The main content area shows a list of alerts. The first alert is at 2014/01/06 21:46:01 IST, labeled 'testalert 1', with a status of 'Paused'. Below the alert list, there are two sections labeled 'Base Event (1 found)'. Each section contains a table with columns: Time (Event Time), device, OLI, deviceVendor, deviceProduct, deviceVersion, deviceEventCause, name, baseEventCause, descriptionAddress, deviceAddress, deviceCustomNumber1, deviceCustomNumberLabel, deviceCustomNumber2, and deviceCustomNumber3. The first base event is for 'Percent Used' on a 'Storage Group' at 2014/01/05 17:56:07 IST. The second base event is for 'Space Used' on a 'Storage Group' at 2014/01/06 21:46:01 IST. At the bottom of each table, there is a 'RAW' field containing a complex query string.

Base Event Fields

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

Go, Export, and Auto Update Options

The **Go** and **Export Results** buttons and the **Auto Update** option accomplish the same tasks in both the Search and Alert pages. For more information, see [“Searching for Events on OLI” on page 96](#), [“Understanding the Search Results Display” on page 99](#), [“Viewing Alerts” on page 122](#), and [“Advanced Search Options” on page 97](#).

Live Event Viewer

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

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

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

4 Searching and Analyzing Events

The screenshot shows the Search Composer interface with the following components and annotations:

- Filter Configuration:** A table with columns 'Exclude/Include Terms' and 'Search Terms'. It contains entries for 'failure' (Search For), 'admin' (Exclude From Search), and another 'Search For' entry.
- Logical NOT:** A checkbox labeled 'Choose if a logical NOT should be applied to the search term and enter the text to search for'.
- Where do you want to look?:** Fields for 'Device Groups' and 'Storage Groups'.
- Start Button:** A blue button at the bottom left.

Annotations on the right side of the image:

- Select a saved filter.
- Save query as saved search query.
- Clear all filters.
- Add filter.
- Remove filter.
- Enter the search criteria here
- Specify constraints here
- Click Start to start view Live Event View

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

The screenshot shows the Search Results interface with the following components and annotations:


- Control Buttons:** Play, Pause, Stop, Clear, and Export buttons at the top left.
- Filter operation:** A dropdown menu at the top center.
- Search Status:** 'Searching' indicator, '260' events found, '277' events scanned, 'Display Buffer 1000', and '12:47' timer.
- Event List:** A list of raw event logs with details like timestamp, source, and severity.

Annotations on the right side of the image:

- Filter specified in Search Composer
- Search timer
- Number of events to display in Viewer
- Events scanned
- Events found
- Current state
- Play, Pause, Stop, Clear, Export
- Number of matching events

The following list highlights the features of Search Results display:

- Events are displayed in the raw event format and not in the columnar, table form as displayed in the Search Results page (Analyze > Search) when you run a search query.
- A user can launch a maximum of one Live Event Viewer. There can be a maximum of five Live Event Viewers running on OLI at any time.
- The regular expression search method is used to identify matching events. Therefore, you can specify regular expressions as the search term in the Search Composer.

- Buffer Size defines the maximum number of events displayed in the Viewer. By default, the Buffer Size is 1000, however, it can be set to any number between the range of 20 and 5000.
- By default, the search is run for 15 minutes and then stopped to preserve system resources. If you need to run the search for longer than 15 minutes, click the  icon next to the countdown timer to reset the timer to 15 minutes.
- When you click Pause, the Search Results display is frozen. However, the search operation continues in the background and the new matching events are buffered until a maximum of 1000 events have been buffered or the search timer, which continues to count down even when the Search Results display is frozen, reaches 00:00.


If the timer has not reached 00:00, you can click Play to resume the paused search operation. When you click Play, the buffered events are displayed. The newly found events are appended to the previously found events on the Search Results display screen.

- When you click Stop, the search for matching events and the count down of the search timer stop. When you click Play, the search is started afresh—the currently displayed events are cleared from the Search Results screen, the search timer is reset to 15 minutes, and the search starts again.
- You must stop the search operation to export the matching events.

To launch a Live Event Viewer:




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


- 1 Click **Analyze > Live Event Viewer** from the top-level menu bar.
- 2 In the Search Composer tab, enter the search terms or click the  icon to select a saved filter.


You can enter search terms that the event must contain (Search For:) or terms that the events must not contain (Exclude From Search:). Click the “Search For:” field to display a drop-down list from which you can select “Exclude From Search:”.

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

To add additional filters, click the  icon.

To remove a filter line, click the  icon.

To remove all filters, click the  icon.

- 3 Enter constraints to limit your search to specific device groups, devices, or storage groups in the “Where do you want to look?” section. Click the  icon to display a list from which you can choose the constraints.
- 4 Click **Start**.

The search results are automatically displayed in the Search Results display screen.

To update the Live Event Viewer query:

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

To export Search Results display:

- 1 Make sure you have stopped the Live Event Viewer. To do so, click the (■) icon in the Search Results display window.
- 2 Click the (↓) icon to open the Export Options window.
- 3 Follow [Step 3](#) onward in “[To export search results:](#)” on page 109 to export the displayed search results.

Chapter 5: Configuration

This chapter describes the Configuration tab, in which you create and manage receivers, forwarders, devices, device groups, and filters. Receivers, devices, and other resources created by one user are visible to all other users, although subject to user group privileges. Resources are shared by all sessions.

This chapter includes information on the following areas of OLI configuration:

- [“Devices” on page 127](#)
- [“Event Archives” on page 131](#)
- [“Storage” on page 136](#)
- [“Event Input” on page 140](#)
- [“Event Output” on page 165](#)
- [“Alerts” on page 172](#)
- [“Scheduled Tasks” on page 187](#)
- [“Filters” on page 189](#)
- [“Saved Searches” on page 192](#)
- [“Search” on page 197](#)
- [“Peer OLIs” on page 203](#)
- [“Configuration Backup and Restore” on page 207](#)
- [“System Maintenance” on page 210](#)
- [“License Information” on page 227](#)
- [“Retrieve Logs” on page 229](#)
- [“Content Management” on page 229](#)

Devices

The Devices section manages both Devices and named collections of devices called device groups.

Devices

A device is a named event source, comprising of an IP address (or hostname) and a receiver name. Two receivers can receive events from the same IP address, so IP address alone is insufficient to identify a device. Event source is the device that directly sends the event to OLI. When an event is sent through a SmartConnector, the event source is the system on which the SmartConnector is running and not the device that sent the event to the SmartConnector.

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

Figure 5-1 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
labm3rdei02.devlab.ad [SmartMessage Receiver]	16.55.245.65	SmartMessage Receiver	System	
labm3rdei03.devlab.ad [SmartMessage Receiver]	16.59.60.216	SmartMessage Receiver	System	
oii-natasha.devlab.ad [SmartMessage Receiver]	16.60.188.54	SmartMessage Receiver	System	
Operations Log Intelligence Internal Event Device	127.0.0.1	Not Applicable	System	System
verticap.devlab.ad [dima_syslog]	16.60.188.72	dima_syslog	System	
verticap.devlab.ad [SmartMessage Receiver]	16.60.188.72	SmartMessage Receiver	System	
vmamqa201.devlab.ad [SmartMessage Receiver]	16.55.247.174	SmartMessage Receiver	System	
vmamqa23.devlab.ad [SmartMessage Receiver]	16.59.63.65	SmartMessage Receiver	System	

Figure 5-1 Devices page

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



Even though there is no limit on the number of devices (IP address of event source + receiver) you can create on OLI, there is limit on the number of source devices (event sources) from which your OLI can receive events. The limit on maximum number of source devices is imposed by the license you obtained when you purchased your OLI.

To pre-define a device:


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

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Devices** in the sub-menu, and then click the **Devices** tab. A display similar to that shown in Figure 5-1 appears.
- 3 Click **Add**.
- 4 Enter a name, an IP address, and select a receiver for the new device.
- 5 Click **Save** to add the new device, or **Cancel** to abandon it.


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

To edit a device:

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

- 2 Click **Devices** in the sub-menu, and then click the **Devices** tab. A display similar to that shown in [Figure 5-1](#) appears.
- 3 Locate the device that you want to edit and click the Edit icon () on that row.
- 4 Change the Name or IP address for the device.
- 5 Click **Save** to update the device group, or **Cancel** to abandon your changes.

To delete a device:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Devices** in the sub-menu, and then click the **Devices** tab. A display similar to that shown in [Figure 5-1](#) appears.
- 3 Locate the device that you want to delete and click the delete icon () on that row.

Deleting a device does not block the source IP address from sending events. If new events are received, autodiscovery recreates the device.

- 4 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the device.

Device Groups

Device groups allow you to categorize named source IP addresses called devices. The Device Groups page lists all device groups with edit and delete icons and includes the ability to create new device groups.



Device groups can be associated with storage rules that define in which storage group events from a specific devices are stored. Doing so enables you to retain event data from different sources for different lengths of times (because you can define different retention policies on different storage groups). For more information about storage rules, see [“Storage Rules” on page 138](#).

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

To create a device group:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Devices** in the sub-menu, and then click the **Device Groups** tab.
- 3 Click **Add**. A display similar to that shown below appears.

Devices **Device Groups**

Add Device Group

You may assign one or more devices to a device group.

If you wish to add a device which is not yet created, you must first go to the [Devices](#) page and create it.

To select or deselect devices, ctrl-click each device name.

Name

Devices

- labm3rdei02.devlab.ad [SmartMessage Receiver]
- labm3rdei03.devlab.ad [SmartMessage Receiver]
- oli-natasha.devlab.ad [SmartMessage Receiver]
- verticap.devlab.ad [dima_syslog]
- verticap.devlab.ad [SmartMessage Receiver]
- vmamqa201.devlab.ad [SmartMessage Receiver]
- vmamqa23.devlab.ad [SmartMessage Receiver]

Use ctrl-click to select or deselect items


4 Enter a name for the new device group. Click to select devices from the list. Press and hold the **Ctrl** key when clicking to add additional devices to the selection. To select a range of devices, click to select the first device, then press and hold the **Shift** key while clicking the last device.

5 Click **Save** to create the new device group, or **Cancel** to abandon it.

To edit a device group:

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

1 Click **Devices** in the sub-menu, and then click the **Device Groups** tab.

2 Locate the device group that you want to edit and click the Edit icon () on that row.


3 Change the Name, add, or remove devices from the selection. Ctrl-Click devices that are not selected to select them, or Ctrl-Click selected devices to remove them from the selection.

4 Click **Save** to update the device group, or **Cancel** to abandon your changes.

To delete a device group:

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

2 Click **Devices** in the sub-menu, and then click the **Device Groups** tab.

3 Locate the device group that you want to delete and click the delete icon () on that row. Deleting a device group does not affect the set of devices.

- 4 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the device group.

Event Archives

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



Caution

Ensure that both Configuration Backups (for configuration settings) and Event Archives (for data) run on a regular basis and are stored in a remote location. In the event of catastrophic failure, you will need to restore the most recent Configuration Backup and Event Archive. For information on Configuration Backups, see [“Configuration Backup and Restore” on page 207](#).

The location is a directory (either local or a mount point that you have already established on the machine on which the OLI is installed).

Events in each storage group are archived separately. That is, one archive file is created for each storage group, for each day. In addition, you can bulk archive events—that is, specify a range of dates to archive events in a single archive operation.

Archiving events from each storage group to a separate archive location enables you to keep data in specific storage groups longer than others. You need to specify these locations when you configure the Archive Storage Settings before archiving any events

You need to enter a complete path where the archive file will be written in the Archive Path field. This path could be a local directory or a mount point that is already established on the machine on which the OLI is installed.

OLI uses the receipt time of an event to determine its archival day. For example, an event with timestamp of 11:55:00 p.m. on December 7 is received at 12:01:00

a.m. on December 8 on the OLI. This event is archived in the archive file created for December 8th and not December 7th. When an archive operation occurs, one archive file per storage group is created at the location specified in Archive Storage Settings. Each archive file contains events from 12:00:00 a.m. to 11:59:59 p.m. for a single storage group of any given day. When you specify a range of dates, one archive file per storage group, for each specified day is created.

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



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

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

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

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

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

The source type information (if associated with an event) is preserved when the event is archived. For information on creating and using source types, see [“Source Types” on page 155](#).

Guidelines for Archiving Events

- Be sure to run Configuration Backups as well as Event Archives regularly, and to store them in a remote location. In the event of catastrophic failure, you will need to restore the most recent Configuration Backup and Event Archive. For information on Configuration Backups, see [“Configuration Backup and Restore” on page 207](#).

- If you need to archive a large number of events (in the order of tens of GB), HP recommends that you archive during the off-peak hours to prevent impacting the performance of your OLI.
- Multiple archiving operations such as loading, unloading, archiving, and deletion of archives can occur simultaneously. Therefore, you can initiate the loading of an existing archive, while an archive operation is in progress.
- Only one manual archive job can run at a time. However, a scheduled archiving operation can run in parallel with a manual job.
- You cannot re-archive the events that have been archived already. If you try to do so, the OLI reports an error.
- Do not move the archived files from their archive location. The archives that have been moved from the originally archived location cannot be loaded on to the OLI. If you need to delete the archives, use the OLI user interface to do so.
- If an archive job fails, you need to initiate it manually. To do so, delete the failed archive and archive it manually. To get notified of a failed archive, configure an alert for this audit event: Event Archive Failed. For more information about this event, see [Appendix B, , on page 293](#). For more information about configuring alerts, see [“Alerts” on page 172](#).
- You can cancel an in-progress archive operation that was manually initiated at any time using the Cancel link that displays on top of the Event Archives page.

Archiving Events

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

An archive storage location must be established on the OLI before you can archive its events. This is a one-time configuration. To establish an archive storage location, see [“Archive Storage Settings” on page 135](#).

To add an Event Archive:

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

Name	Day	Month	Year	Storage Group	Status	Mount	Mount Path
archive [2012-11-28] [Default Storage Group]	28	11	2012	Default Storage Group	Archived	test1	
archive [2012-11-28] [StorageGroup2]	28	11	2012	StorageGroup2	Archived	test1	
archive [2012-11-28] [StorageGroup3]	28	11	2012	StorageGroup3	Archived	test1	
archive [2012-11-28] [StorageGroup4]	28	11	2012	StorageGroup4	Archived	test1	
archive [2012-11-28] [StorageGroup5]	28	11	2012	StorageGroup5	Archived	test1	
archive [2012-11-27] [Default Storage Group]	27	11	2012	Default Storage Group	Failed	test1	

- 3 Click **Add** in the Event Archives tab, in the right panel.
- 4 Enter a meaningful name in the Name field for the new Event Archive and specify the Start and End dates in the format m/dd/yy, where m is month

number, dd is the day of the month (with a leading zero if necessary), and yy is the two-digit year number.

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

Start Date: 8/12/12

End Date: 8/13/12

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

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

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



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

To delete an Event Archive:

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

Scheduled Event Archive

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

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 236](#) before you schedule an event archive.

To schedule a daily event archive:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Archives** in the left panel.
- 3 Click the **Daily Task Settings** tab in the right panel.

- 4 Select a time from the “Time For Daily Archive to Start” list. Scheduled archives must start on the hour, and midnight and 1:00 AM are not on the list to allow your OLI to receive all of the previous day’s events.
- 5 Select the storage groups whose events should be included in the scheduled archive.
- 6 Click **Save** to schedule daily event archive, or click on another tab or sub-menu to cancel.

Archive Storage Settings

Event archives are saved to the specified directory, which can be a path to a local directory or to a mount point on the machine on which the OLI is installed. To establish a mount point, see your system’s operating system documentation.

To perform Archive Storage Setting setup:

- 1 If you intend to use an NFS or CIFS mount point, ensure that the external storage point is mounted on the machine on which OLI is installed. See your system’s operating system documentation for more information.
- 2 Click **Configuration** from the top-level menu bar.
- 3 Click **Event Archives** in the left panel.
- 4 Click the **Archive Storage Settings** tab in the right panel.
- 5 Specify a mount location and an archive path for each storage group. You can specify a different path for each storage group, thus enabling the OLI to archive events to a different location for each storage group.

You need to enter a complete path where the archive file will be written in the Archive Path field. This path could be a local directory or a mount point that is already established on the machine on which the OLI is installed.



Note

You must configure settings for all storage groups on the Archive Storage Settings tab even if you do not intend to archive all of them.

Event Archives
Daily Task Settings
Archive Storage Settings

Important
Before you may archive any event data, you must setup the event archive settings.

Storage Group	Default Storage Group
Archive Path	<input style="width: 80%;" type="text"/>
Storage Group	Internal Event Storage Group
Archive Path	<input style="width: 80%;" type="text"/>
<input style="background-color: #0070c0; color: white; padding: 5px 15px; border: none; cursor: pointer;" type="button" value="Save"/>	

- 6 Click **Save**.

Loading and Unloading Archives

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

To load or unload an Event Archive:

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

Storage

Different storage groups support the implementation of multiple retention policies. Each group can have a different policy, and storage rules determine which storage group is used for events from specific device groups. See [“Retention Policy” on page 23](#). The Storage section has three tabs: Storage Groups, Storage Rules, and Storage Volume.



Storage Groups					
Storage Groups Storage Rules Storage Volume					
Name	Maximum Age (Days)	Maximum Size (GB)	Creator	Last Editor	
Default Storage Group	180	3	admin	admin	
Internal Event Storage Group	365	3	System	System	

Figure 5-2 Storage Groups page

Storage Groups

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

OLI can have a maximum of 6 storage groups—two that pre-exist on your OLI (Internal Storage Group and Default Storage Group) and four that you can create. You can add the additional storage groups (up to the maximum of six) at any time.

HP recommends that you create the four additional storage groups in addition to the two that preexist, so that you have five storage groups available for event storage and one for OLI's internal events.

To add additional storage groups, follow the instructions in [“Adding Storage Groups” on page 219](#).

Once a storage group is created, it cannot be deleted however its size can be increased or decreased any time. If you are decreasing the size of the storage group and the new size is lesser than the currently used space on the storage group, you will need to delete data to achieve the new size. OLI UI guides you in this situation to delete sufficient data.

To edit (including resizing) a storage group:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Storage** in the left panel. The Storage Groups tab displays the available storage groups.
- 3 Identify the storage group you want to modify and click the Edit icon (✎) for it. The Storage Groups tab displays the Edit Storage Group pane.

Storage Groups Storage Rules Storage Volume

Edit Default Storage Group

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

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

Maximum Age (Days) 180

Maximum Size (GB) 12

Current Size (GB) 1

Save Cancel

- 4 Change the desired parameters such as the name of the storage group, or increase or decrease Maximum Age or Maximum size.



The names of the Internal Storage Group and Default Storage Group cannot be modified.

Note

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, OLI displays a message indicating that reducing storage group size in this situation will require you to delete existing data.

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

- a Set the Maximum Age value to the number indicated in the above message. Doing so, triggers deletion of events.

- b Refresh the Edit Storage Group screen. When the “Current Size” value is less than or equal to the storage group size you want to set, go to the next step. Otherwise, keep refreshing the screen periodically.



The “Current Size” value changes as data are deleted, which can take some time. Therefore, you need to wait before proceeding to the next step.

- c Set the Maximum Size value to suit your needs.
- d If you wish, restore the Maximum Age setting (that you changed in Step b) to the original value.

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



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

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

Storage Rules

Storage rules create a mapping between device groups and storage groups. Doing so enables you to store events from specific sources to a specific storage group. Additionally, you can configure these storage groups with different retention policies, and thus retain event data based on the source of incoming events. For example, all events from firewall devices can be subject to a short retention period. To accomplish this, manually assign the firewall devices to a device group and then create a storage rule that maps the device group to a storage group with the desired short retention period.



Events that are not subject to any storage rule are sent to the Default Storage Group.

Before you add a storage rule, make sure that the storage group to which you want to store the events and the device group that contains the devices whose events you want to store exist. For information on how to create device groups, see [“Device Groups” on page 129](#).

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

To add a storage rule:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Storage** in the left panel, and then click the **Storage Rules** tab.
- 3 Click **Add**. The following page is displayed.

Storage Groups | **Storage Rules** | Storage Volume

Add Storage Rule

Storage Group:

Device Group:


Priority:

- 4 Enter the following parameters:


Parameter	Description
Storage Group	Select a storage group from the drop-down list. The storage groups must already be set up before any storage rules are added. You can only add storage groups at the time of OLI initialization.
Device Groups	Select one or more device groups to associate with the specified storage group. You may associate several device groups with a single storage group.
Priority	An integer that indicates the new rule's priority. The number must be unique for each storage rule. The smaller the number, the higher the rule's priority.

- 5 Click **Save** to add the new storage rule, or **Cancel** to quit.

To edit or reorder a storage rule:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Storage** in the left panel.
- 3 Click the **Storage Rules** tab.
- 4 Find the storage rule that you want to edit and click the Edit icon () on that row.
- 5 Change the information in the form--for example, change the priority value to reposition the storage rule in the table--and click **Save**.

To delete a storage rule:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Storage** in the left panel.
- 3 Click the **Storage Rules** tab.
- 4 Find the storage rule that you want to delete and click the Remove icon ().
- 5 Click **OK** to confirm the delete.

Storage Volume

The Storage Volume tab displays the mount location and current storage volume settings. To increase the Storage Volume size, go to the System Maintenance page. You must have admin-level privileges to perform this operation. For more information, see [“Storage Volume Size Increase” on page 217](#).



Note

Storage volume can be extended after initialization, but its size cannot be reduced.

Event Input

The event input section allows you to set up how the data that comes into OLI is described. You can configure receivers to listen for and capture event data, set up source types that identify the type of log file that captured events come from, and create parsers that extract field-value pairs when searching the data.

Receivers

From the Receivers tab, you can set up and configure the receivers that will capture event data for your OLI, and populate each event with information about its origin. Some receivers capture streaming events transmitted over the network by devices, applications, services, and so on. Other types of receivers monitor individual files for events or monitor files selected from a directory tree, based on a pattern you specify. Since receivers can only receive events of a single source type, you should set up separate receivers for each type of log file. To start receiving events, direct your event sources to the default receivers. For more information about the default receivers, see [“Receivers” on page 25](#).

Receiver types include UDP, TCP, SmartMessage, and three types of file based receivers, File Transfer, File Receiver, and Folder Follower Receiver.

You can configure the following types of receivers for OLI:

- **UDP Receiver:** UDP receivers listen for User Datagram Protocol messages on the port you specify. OLI comes pre-configured with a UDP Receiver on port 514 or 8514, enabled by default. This port may vary based on the port numbers available at installation time.
- **CEF UDP Receiver:** UDP receivers that receive events in Common Event Format.
- **TCP Receiver:** TCP receivers listen for Transmission Control Protocol messages on the port you specify. OLI comes pre-configured with a TCP receiver on port 515 or 8515, enabled by default. This port may vary based on the port numbers available at installation time.
- **CEF TCP Receiver:** TCP receivers that receive events in Common Event Format.
- **File Receiver:** Depending on the type of OLI, file receivers read log files from a local file system, Network File System (NFS), Common Internet File System (CIFS), or Storage Area Network (SAN). File receivers read single or

multi-line log files. They provide a snapshot of a log file at a single point in time.

- **Folder Follower Receiver:** Folder follower receivers actively read the log files in a specified directory as they are updated. If the source directory contains different types of log files, you can create a receiver for each type of file that you want to monitor. OLI comes pre-configured with folder follower receivers for OLI's Apache Access Error Log, the system Messages Log, and Audit Log (when auditing is enabled). You must enable these receivers in order to use them.
- **File Transfer:** File Transfer receivers read remote log files using SCP, SFTP or FTP protocol. These receivers can read single- or multi-line log files. You can schedule the receiver to read a file or batch of files periodically.



The SCP, SFTP, and FTP file transfer receivers depend on the FTP (File Transfer Protocol) SCP (Secure Copy Protocol) and SFTP (SSH file transfer protocol) clients installed on your system.

- **SmartMessage Receiver:** SmartMessage receivers listen for encrypted messages from ArcSight SmartConnectors. OLI comes pre-configured with a SmartMessage receiver with the name "SmartMessage Receiver". To use this receiver to receive events from a SmartConnector, set the **Receiver Name** to be "SmartMessage Receiver" when configuring the SmartConnector's destination.

File Based Receivers

File based receivers include File Receiver, File Transfer Receiver and Folder Follower Receiver. You can set them up as multiline receivers, and configure them to use source types with associated parsers to extract data from captured events.



When a receiver cannot read the file it logs from, such as when the file or folder is deleted or renamed, OLI records a message in `current/arcsight/logger/logs/Logger_receiver.log`.

Multi-line Receivers

TCP and UDP receivers interpret line break characters, such as `\r` or `\n`, as the end of the event. If the input event contains embedded `\r` or `\n` characters, the event will be treated as more than one event. If your events span more than one line, you may want to use a multi-line receiver. Multi-line receivers include the File Transfer, File Receiver and Folder Follower Receivers.

A multi-line receiver can read events that span more than one line, such as a server log. You could set up the receiver to handle stack traces reported in the log by reading the entire stack trace as a single event instead of reading each line separately.

When creating a multi-line receiver, you must specify a regular expression that the receiver should use to detect the start of a new event in the log file. Each new event starts where the characters in the log file match the regular expression.

For example, in the following log file, each event starts with a timestamp embedded within square brackets ([yy-MM-dd HH:mm:ss.SSS]); therefore, you can use this regular expression to identify each event:

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

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

- For multi-line file receivers and file transfer receivers, the regular expression that identifies the beginning of a new event must be specified in the receiver's *Event Start* field.
- For multi-line folder follower receivers, the regular expression that identifies the beginning of a new event must be specified in the *Multiline Regex* field of the source type associated with that receiver.

For information on creating and using receivers, see [“Working with Receivers” on page 143](#). For information on creating and using source types, see [“Source Types” on page 155](#).

Folder Follower Receivers

When you want to monitor active files as they are updated, use a folder follower receiver. After you set up a folder follower receiver and enable it, it will monitor the specified files in that directory and continuously upload new events to OLI. Folder follower receivers recognize file rotation.

Overview of the steps to monitor a directory:

- 1 Determine the types of logs you need to monitor.
- 2 Determine whether the out-of-box source types or source type/parser pairs will satisfy your needs. For more information, see [“Source Types” on page 155](#), and [“Parsers” on page 159](#).
 - ◆ If so, proceed to [Step 3](#).
 - ◆ If not, create the parsers and sources types that you need.
 - i Select an appropriate parser or set of parser for the log files in the directory you want to follow. If the out-of-box parsers do not provide what you need, create appropriate parsers.
 - ii Assign a source type for each parser. If the out-of-box source types do not provide what you need, create appropriate source types.
- 3 Create the folder follower receivers required to monitor the logs in the directory, selecting the source type determined in [Step 1](#). For more information, see [“Working with Receivers” on page 143](#).
- 4 Enable the receivers.

- 5 Optionally, to forward log file events, set up and configure one or more forwarders. For more information, see [“Forwarders” on page 165](#).

Using Source Types with File Follower Receivers

OLI uses the parser associated with the source type you select for a receiver to extract fields and their respective values from the received events. These fields are parsed at search time. For more information on using source types and parsers, see [“Source Types” on page 155](#), and [“Parsers” on page 159](#).

When creating a file follower receiver, you must select a source type appropriate to monitor a specific type of log file. After you select the source type for the file follower receiver, ensure that the parser associated with it works with your source files.

Events from different versions of the same source type can be in different formats. Similarly, events from different source types of the same vendor might be formatted differently. Therefore, if the source type you choose from the OLI UI does not exactly match the specifications of your source type, the associated parser will not parse events correctly, and the search results will not display any parsed fields.

To confirm whether the source type has a valid parser for your source type, after you have set up the receiver, check whether the incoming events are parsed. To determine this, run a search and review the “parser” field in the search results. The parser used in the search will be displayed in the parser column of the search results. If the event was parsed, this field contains the name of the parser. If the event was not parsed successfully, this field contains “Not parsed”. If no parser is defined for the source type or if there is no source type, the field is blank.

Working with Receivers

Several receivers come set up on your OLI. You can add other receivers as needed. The maximum number of receivers that can be created on OLI is limited by system resources—memory, CPU, disk input/output and possibly network bandwidth.

The receiver ports available on your OLI may vary from the image below.

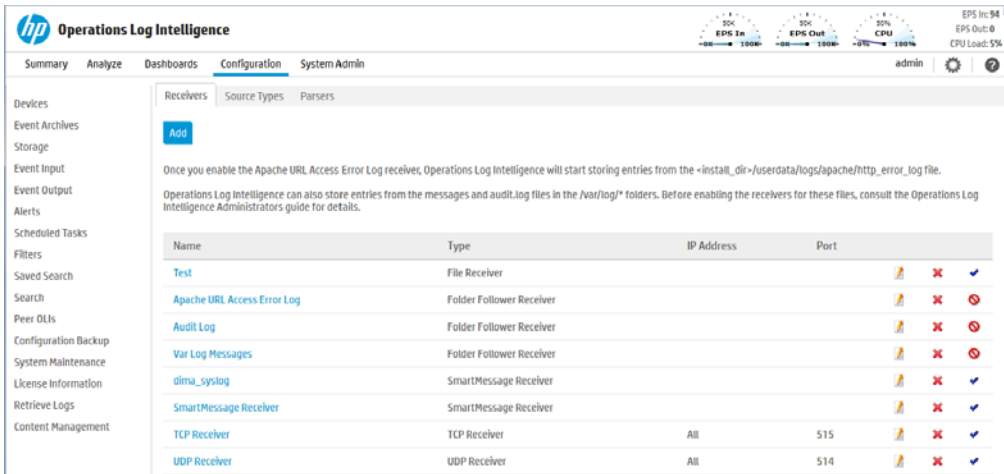


Figure 5-3 Receivers tab

To create a receiver:

Before creating a receiver of type File Receiver, for the OLI, the file system from which the log files will be read needs to be mounted on the system on which you have installed OLI.

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Receivers** tab (right panel).
The Receivers tab, shown in [Figure 5-3 on page 144](#), displays the current receivers and their status. You can sort the fields by clicking the column headers.
- 3 Click **Add**.
- 4 Enter a name for the new receiver. SmartMessage receiver names are used when configuring the associated ArcSight SmartConnectors.
- 5 Choose the receiver type. Select UDP Receiver, TCP Receiver, CEF UDP Receiver, CEF TCP Receiver, File Receiver, Folder Follower Receiver, File Transfer, or SmartMessage Receiver.



The receiver type cannot be changed after the receiver is created.

- 6 Click **Next** to edit receiver parameters.
The fields displayed in the Edit Receiver dialog box vary according to the type of OLI and the type of receiver.
- 7 Fill in the appropriate fields. Refer to the following tables for field descriptions.
 - ◆ [Table 5-1, “Parameters used in UDP Receivers, TCP Receivers, CEF UDP Receivers, and CEF TCP Receivers,” on page 146](#)

- ◆ [Table 5-2, “Parameters used in File Receivers,” on page 147](#)
 - ◆ [Table 5-3, “Parameters used in Folder Follow Receivers,” on page 149](#)
 - ◆ [Table 5-4, “Parameters used in File Transfer Receivers,” on page 150](#)
 - ◆ [Table 5-5, “Parameters used in SmartMessage Receivers,” on page 153](#)
- 8 Click **Save**.
 - 9 New receivers are initially disabled. You must enable them in order to use them.

To enable or disable a receiver:





Note

Before enabling the following preconfigured folder follower receivers for OLI, ensure that the files are readable by the non-root user that you installed with or specified during installation.

- `/var/log/messages`
- `/var/log/audit/audit.log`

- 1 Click **Configuration** ofrom the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Receivers** tab (right panel).

The Receivers tab, shown in [Figure 5-3 on page 144](#), displays the current receivers and their status. You can sort the fields by clicking the column headers.

- 3 Locate the receiver that you want to enable or disable.
 - ◆ If the receiver is currently disabled, click the disabled icon () to enable it.
 - ◆ If the receiver is currently enabled, click the enabled () icon to disable it.




Tip

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

To edit a receiver:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Receivers** tab (right panel).

The Receivers tab, shown in [Figure 5-3 on page 144](#), displays the current receivers and their status. You can sort the fields by clicking the column headers.

- 3 Locate the receiver that you want to update and click the Edit icon () on that row.

The fields displayed in the Edit Receiver dialog box vary according to the type of OLI and the type of Receiver.

- 4 Edit the appropriate fields. Refer to the following tables for field descriptions.
 - ◆ [Table 5-1, “Parameters used in UDP Receivers, TCP Receivers, CEF UDP Receivers, and CEF TCP Receivers,” on page 146](#)
 - ◆ [Table 5-2, “Parameters used in File Receivers,” on page 147](#)
 - ◆ [Table 5-3, “Parameters used in Folder Follow Receivers,” on page 149](#)
 - ◆ [Table 5-4, “Parameters used in File Transfer Receivers,” on page 150](#)
 - ◆ [Table 5-5, “Parameters used in SmartMessage Receivers,” on page 153](#)
- 5 Click **Save**.

To delete a receiver:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Receivers** tab (right panel).
 The Receivers tab, shown in [Figure 5-3 on page 144](#), displays the current receivers and their status. You can sort the fields by clicking the column headers.
- 3 Locate the receiver that you want to delete and click the Remove icon (✖) on that row.
- 4 Click **OK** to confirm the delete.

Receiver Parameters

The following tables describe the parameters used when creating and editing receivers:

- [Table 5-1, “Parameters used in UDP Receivers, TCP Receivers, CEF UDP Receivers, and CEF TCP Receivers,” on page 146](#)
- [Table 5-2, “Parameters used in File Receivers,” on page 147](#)
- [Table 5-3, “Parameters used in Folder Follow Receivers,” on page 149](#)
- [Table 5-4, “Parameters used in File Transfer Receivers,” on page 150](#)
- [Table 5-5, “Parameters used in SmartMessage Receivers,” on page 153](#)

Fill in the following fields when creating or editing UDP Receivers, TCP Receivers, CEF UDP Receivers, and CEF TCP Receivers.

Table 5-1 Parameters used in UDP Receivers, TCP Receivers, CEF UDP Receivers, and CEF TCP Receivers

Parameter	Description
Name	The name of the receiver, used for status monitoring.
IP/Host	Select one of the OLI's network connections for the receiver to listen to, or select All to listen on both network connections. Note: If localhost (127.0.0.1) appears in the list, it means that the OLI hostname has not been configured.

Table 5-1 Parameters used in UDP Receivers, TCP Receivers, CEF UDP Receivers, and CEF TCP Receivers (Continued)

Parameter	Description
Port	<p>If you installed OLI as a root user, you can use any available port. The default UDP Receiver is pre-configured on port 514. If that port is not available, then the next higher available port is chosen.</p> <p>If you installed OLI as a non-root user, you can only use a port numbers greater than 1024. The default UDP Receiver is pre-configured on port 8514. If that port is not available, then the next higher available port is chosen.</p>
Encoding	Select a character encoding, such as US-ASCII, Big5, or EUC-KR, from the pulldown list. CEF UDP, CEF TCP, and SmartMessage receivers must use US-ASCII or UTF-8 encoding.
Source Type	<p>Select from the pulldown list of log file types, including:</p> <ul style="list-style-type: none"> • Apache HTTP Server Access • Apache HTTP Server Error • Juniper Steel-Belted Radius • Microsoft DHCP Log • IBM DB2 Audit <p>Additionally, you can define your own source type, based on the needs of your company. See “Source Types” on page 155.</p> <p>A receiver can only receive events of a single source type. Set up separate receivers for each type of log file.</p> <p>Note: CEF TCP and CEF UDP receivers are set to the CEF source type, and cannot be changed. Currently, there is no parser associated with the CEF source type.</p>

Fill in the following fields when creating or editing File Receivers.

Table 5-2 Parameters used in File Receivers

Parameter	Description
Name	The name of the receiver, used for status monitoring.
RFS Names	Select from the pulldown list of NFS or CIFS mount names.
Folder	<p>Choose “Local” and then specify the directory on your OLI where the remote file system is mounted in the “Folder” field.</p> <p>To mount a remote file system on the system on which you have installed OLI, see its operating system’s documentation.</p>

Table 5-2 Parameters used in File Receivers (Continued)

Parameter	Description
Source Type	<p>Select from the pulldown list of log file types, including:</p> <ul style="list-style-type: none"> • Apache HTTP Server Access • Apache HTTP Server Error • Juniper Steel-Belted Radius • Microsoft DHCP Log • IBM DB2 Audit <p>Additionally, you can define your own source type, based on the needs of your company. See “Source Types” on page 155.</p> <p>A receiver can only receive events of a single source type. Set up separate receivers for each type of log file.</p>
Mode	<p>Select one of the following:</p> <p>Delete - delete the log file once it has been processed</p> <p>Rename - rename the log file once it has been processed. The file is named by appending the Rename Extension.</p> <p>Persist - OLI remembers which files have been processed and only processes them once.</p>
Rename Extension	The suffix to append to log files that have been processed.
Character Encoding	Select a character encoding, such as US-ASCII, Big5, or EUC-KR, from the pulldown list. CEF UDP, CEF TCP, and SmartMessage receivers must use US-ASCII or UTF-8 encoding.
Delay after seen	<p>Number of seconds to wait after a source file is first seen until it is processed. This allows the entire file to be copied to OLI or (in the case of File Receiver) copied to the remote file system, before processing begins.</p> <p>The default is 10 seconds.</p> <p>Note: For File Transfer Receivers, this parameter should be set to a larger value if large files are expected. The default, 10 seconds, does not allow enough time for a large file, such as 1 GB.</p>
Date/time locale	Select a locale from the pulldown list, such as English (United States), Chinese (Hong Kong), Chinese (Taiwan), and so on.
Date/time zone	<p>Required if the timestamp in the log file does not specify a time zone.</p> <p>For File Transfer and File Receivers, this parameter is ignored if either Date/time format or Date/time location regex are blank.</p> <p>You can see the time zone configured on the OLI (System Admin > Settings > Platform > Time/NTP).</p> <p>OLIs use the system time.</p>

Table 5-2 Parameters used in File Receivers (Continued)

Parameter	Description
Event/time location	<p>A regular expression describing which characters represent the timestamp in the log file. For example:</p> <pre>.*\[(.*)\].*</pre> <p>This regular expression specifies that the timestamp is found inside the first set of square brackets on each line. The first capturing group (the part of the regex in parentheses) is that part that is then parsed using the Date/time format.</p> <p>The default is "" (no timestamp in log file).</p>
Event/time format	<p>Required if the log file contains timestamps in the same format for each event. If not specified (or if the Date/time location regex is blank), each event in the file will be stamped with the date that the file itself was first seen by OLI (not its file system timestamp).</p> <p>See Table 5-7 for a list of format specifiers.</p> <p>The default is "" (no timestamp in log file).</p>
Multiline (regex)	<p>A regular expression that specifies the start of a new event in a log file. Specify this expression to enable the receiver to read multi-line log files. Each new event starts at the point where the regular expression is matched to the characters in the log file. For example,</p> <pre>^\[\d+-\d+-\d+ \d+:\d+,\d+ \].*</pre> <p>This regular expression matches timestamps such as:</p> <pre>[2010-12-06 13:09:46,818]</pre> <p>When this field is left blank (""), each line in the log file is treated as a single event.</p> <p>The default is "" (each line in the log file is a single event).</p>

Fill in the following fields when creating or editing Folder Follower Receivers,

Table 5-3 Parameters used in Folder Follow Receivers

Parameter	Description
Name	The name of the receiver, used for status monitoring.
Local Folder	Specify the local folder to process.

Table 5-3 Parameters used in Folder Follow Receivers (Continued)

Parameter	Description
Wildcard (regex)	<p>A regular expression (regex) describing the log files to read.</p> <p>Note: This is a regular expression, not a typical file wildcard like “*.*”.</p> <p>The default is <code>.*</code>, meaning all files.</p> <p>Examples:</p> <p>To include all files ending with <code>.process</code>, you could use: <code>.*\.process</code></p> <p>To monitor only <code>*.properties</code> files, you could use: <code>.*\.properties</code></p> <p>To include only <code>.log</code> files with eight digit filenames, you could use: <code>\d{8}.log</code></p>
Blacklist (regex)	<p>A regular expression (regex) describing the name of the log files to ignore. Files are not monitored if they match this expression.</p> <p>Note: This is a regular expression, not a typical file wildcard like “*.*”.</p> <p>Example:</p> <p>To exclude files that end in <code>.txt</code>, you could use: <code>.*\.txt</code></p> <p>To monitor all files except <code>*.txt</code>, you could use: Wildcard: <code>.*</code> Blacklist: <code>.*\.txt</code></p>
Character encoding	<p>Select a character encoding, such as US-ASCII, Big5, or EUC-KR, from the pulldown list. CEF UDP, CEF TCP, and SmartMessage receivers must use US-ASCII or UTF-8 encoding.</p>
Date/time zone	<p>Required if the timestamp in the log file does not specify a time zone.</p> <p>For File Transfer and File Receivers, this parameter is ignored if either Date/time format or Date/time location regex are blank.</p> <p>You can see the time zone configured on the OLI (System Admin > Settings > Platform > Time/NTP).</p> <p>OLIs use the system time.</p>

Fill in the following fields when creating or editing File Transfer Receivers,

Table 5-4 Parameters used in File Transfer Receivers

Parameter	Description
Name	The name of the receiver, used for status monitoring.
Protocol	Select SCP, SFTP or FTP protocol.

Table 5-4 Parameters used in File Transfer Receivers (Continued)

Parameter	Description
Port	<p>If you installed OLI as a root user, you can use any available port. The default UDP Receiver is pre-configured on port 514. If that port is not available, then the next higher available port is chosen.</p> <p>If you installed OLI as a non-root user, you can only use a port numbers greater than 1024. The default UDP Receiver is pre-configured on port 8514. If that port is not available, then the next higher available port is chosen.</p>
IP/Host	<p>Select one of the OLI'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 OLI hostname has not been configured.</p>
User	A user on the host with privileges to view and read the source log files. If the protocol is FTP, you can specify the special user, "anonymous."
Password	The password of the specified User. The password must not be empty, even in the case of anonymous FTP (although in this case, the password will be ignored.)
File Path	<p>The path and the name of the log file(s) to be read. You can use wild cards like ? and * (for example, "*.log" or "Log-?.txt") in the path name and the file name. Separate directories with forward slashes (/).</p> <p>Separate multiple file specifications with commas.</p> <p>Example: /tmp/SyslogData/syslog.log.gz, /security/logs/*/ , /security/ log?/admin/special/</p>
Schedule	<p>If no schedule is specified, the File Transfer will occur just once.</p> <p>Choose Everyday or Days of Week from the first pulldown menu.</p> <p>If Everyday, select Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the text box.</p> <p>If Days of Week, enter the days (day 1 is Sunday) in the text box. Then choose Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the second text box.</p> <p>For example, to read log files every day at 2 a.m., select Everyday in the first pulldown menu, then choose Hour of Day from the second pulldown menu and enter 2 in the text box. To read the log files every day at 2 a.m. and 3 p.m., enter 2,15 in the text box.</p> <p>For another example, to read log files Tuesdays and Thursdays at 10 p.m., select Days of Week from the first pulldown menu and enter 3,5 for days. Then choose Hour of Day from the second pulldown menu and enter 22 in the text box.</p> <p>Make sure you are familiar with the information in "Impact of Daylight Savings Time Change on Logger Operations" on page 236 before you schedule a file transfer.</p>
Zip Format	Choose gzip, zip, or none.

Table 5-4 Parameters used in File Transfer Receivers (Continued)

Parameter	Description
Source Type	<p>Select from the pulldown list of log file types, including:</p> <ul style="list-style-type: none"> • Apache HTTP Server Access • Apache HTTP Server Error • Juniper Steel-Belted Radius • Microsoft DHCP Log • IBM DB2 Audit <p>Additionally, you can define your own source type, based on the needs of your company. See “Source Types” on page 155.</p> <p>A receiver can only receive events of a single source type. Set up separate receivers for each type of log file.</p>
Character Encoding	<p>Select a character encoding, such as US-ASCII, Big5, or EUC-KR, from the pulldown list. CEF UDP, CEF TCP, and SmartMessage receivers must use US-ASCII or UTF-8 encoding.</p>
Delay after seen	<p>Number of seconds to wait after a source file is first seen until it is processed. This allows the entire file to be copied to OLI or (in the case of File Receiver) copied to the remote file system, before processing begins.</p> <p>The default is 10 seconds.</p> <p>Note: For File Transfer Receivers, this parameter should be set to a larger value if large files are expected. The default, 10 seconds, does not allow enough time for a large file, such as 1 GB.</p>
Event Time locale	<p>Select a locale from the pulldown list, such as English (United States), Chinese (Hong Kong), Chinese (Taiwan), and so on.</p>
Date/time zone	<p>Required if the timestamp in the log file does not specify a time zone.</p> <p>For File Transfer and File Receivers, this parameter is ignored if either Date/time format or Date/time location regex are blank.</p> <p>You can see the time zone configured on the OLI (System Admin > Settings > Platform > Time/NTP).</p> <p>OLIs use the system time.</p>
Event Time location regular expression	<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>

Table 5-4 Parameters used in File Transfer Receivers (Continued)

Parameter	Description
Event time format	<p>Required if the log file contains timestamps in the same format for each event. If not specified (or if the Date/time location regex is blank), each event in the file will be stamped with the date that the file itself was first seen by OLI (not its file system timestamp).</p> <p>See Table 5-7, "Date/time format specification," on page 154 for a list of format specifiers.</p> <p>The default is "" (no timestamp in log file).</p>
Multiline (regex)	<p>A regular expression that specifies the start of a new event in a log file. Specify this expression to enable the receiver to read multi-line log files. Each new event starts at the point where the regular expression is matched to the characters in the log file. For example,</p> <pre>^\[\d+-\d+-\d+ \d+:\d+,\d+\].*</pre> <p>This regular expression matches timestamps such as:</p> <pre>[2010-12-06 13:09:46,818]</pre> <p>When this field is left blank (""), each line in the log file is treated as a single event.</p> <p>The default is "" (each line in the log file is a single event).</p>

Fill in the following fields when creating or editing SmartMessage Receivers,

Table 5-5 Parameters used in SmartMessage Receivers

Parameter	Description
Name	The name of the receiver, used when configuring an associated ArcSight SmartConnectors.
Encoding	Select a character encoding, such as US-ASCII, Big5, or EUC-KR, from the pulldown list. CEF UDP, CEF TCP, and SmartMessage receivers must use US-ASCII or UTF-8 encoding.

Date and Time Specification

To specify the date and time format so that it can be parsed from a file receiver, (File Receiver, Folder Follower Receiver, or File Transfer), refer to [Table 5-7 on page 154](#).

Internally, OLI 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 OLI. Pattern letters are usually repeated, as their number determines the exact presentation:

The examples in [Table 5-6 on page 154](#) 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 5-6 Date/time examples

Source	Date and Time Pattern
2001.07.04 AD at 12:08:56 PDT	yyyy.MM.dd G 'at' HH:mm:ss z
Wed, Jul 4, '01	EEE, MMM d, 'yy
12:08 PM	h:mm a
12 o'clock PM, Pacific Daylight Time	hh 'o'clock' a, zzzz
0:08 PM, PDT	K:mm a, z
02001.July.04 AD 12:08 PM	yyyyy.MMMMM.dd GGG hh:mm aaa
Wed, 4 Jul 2001 12:08:56 -0700	EEE, d MMM yyyy HH:mm:ss Z
010704120856-0700	yyMMddHHmmssZ
2001-07-04T12:08:56.235-0700	yyyy-MM-dd'T'HH:mm:ss.SSSZ

Table 5-7 Date/time format specification

Symbol	Meaning	Presentation	Examples
G	Era designator	(Text)	AD
y	Year	(Number)	2006 or 06
M	Month in year (1-12)	(Month)	July or Jul or 07
w	Week in year (1-52)	(Number)	39
W	Week in month (1-5)	(Number)	2
D	Day in year (1-366)	(Number)	129
d	Day in month (1-31)	(Number)	10
E	Day in week	(Text)	Tuesday or Tue
F	Day in week of month		
a	Am/pm marker	(Text)	AM or PM
H	Hour in day (0-23)	(Number)	0
k	Hour in day (1-24)	(Number)	24
K	Hour in am/pm (0-11)	(Number)	0
h	Hour in am/pm (1-12)	(Number)	12
m	Minute in hour (0-59)	(Number)	30

Table 5-7 Date/time format specification (Continued)

Symbol	Meaning	Presentation	Examples
s	Second in minute (0-59)	(Number)	55
S	Millisecond (0-999)	(Number)	978
z	Time zone	(Text)	Pacific Standard Time, or PST, or GMT-08:00
Z	Time zone	(RFC 822)	-0800 (indicating PST)

Source Types

Source types identify the kind of event that comes from a specific data source. For example, an event could come from an Apache access log, a simple syslog, or the log of an application you created. You can use parsers to parse event data from a specified source type.

Once events are associated with a source type, if the source type is associated with a parser, the events are parsed by that parser when you run a search that matches those events. The search result displays the matching parsed event fields in columns, similar to the CEF events. (Use the “User Defined Fields” field set to view these events.) For more information, see [“Parsers” on page 159](#).

The source of the event, the source type, and the parser will be displayed in the column list of the search results if any row is fetched from a search which contains a non-CEF source type.

The following columns are displayed in the search results when a source type is used:

- **Source**—The name of the log file from which the event was received. For example, `/opt/mnt/testsoft/web_server.out.log`. If no source was applied when the event was received, this field is blank. You can control whether this field is displayed from the Search Options tab. See [“Tuning Advanced Search Options” on page 197](#) for how to set this option.
- **Source Type**—The type of file from which the event was received, as defined on the Source Type page (**Configuration > Event Input > Source Types**). If no source type was applied when the event was received, this field is blank. You can control whether this field is displayed from the Search Options tab. See [“Tuning Advanced Search Options” on page 197](#) for how to set this option.
- **Parser**—If the event was parsed, this field contains the name of the parser. If the event was not parsed successfully, this field contains “Not parsed”. If no parser is defined for the source type or if there is no source type, the field is blank.

Working with Source Types

OLI provides a number of source types with pre-configured parsers. Additionally, you can define new source types and assign parsers to them. This lets you choose the set of fields you want to extract for a given kind of event. Only one parser can be associated with a source type, however, multiple source types can be associated with a parser. Out-of-box source types cannot be edited or deleted, but you can copy them to make similar source types to meet your needs. You can edit or delete custom source types, as desired. The source types available on your OLI may vary from the image below.

Name	Description	Parser	Event Time Location	Event Time Format	Multiline Regex	Locale
Apache_access	Apache Access Log	Apache_access	*\{[*?]\}.*	dd/MMM /yyyy:HH:mm:ss Z		English (United States)
Apache_error	Apache Error Log	Apache_error	*\{[*?]\}.*	EEE MMM dd HH:mm:ss yyyy		English (United States)
Apache HTTP Server Access (for connector forwarder)	Apache HTTP Server Access log type used to be forwarded to streaming connector					English (United States)
Apache HTTP Server Error (for connector forwarder)	Apache HTTP Server Error log type used to be forwarded to streaming connector					English (United States)
audit_log	Syslog for Audit Log files	audit_log				English (United States)
Bluecoat_proxy	Bluecoat Proxy SG	Bluecoat_proxy	*\{lw(3)\s\d+\s\d\d:\d\d \d:\d\d\}.*	MMM dd HH:mm:ss		English (United States)

Figure 5-4 Source Types tab

The following source types have associated parsers:

Source type	Description
Apache_access	Apache Access Log
Apache_error	Apache Error Log
audit_log	Syslog for Audit Log files
Bluecoat_proxy	Bluecoat Proxy SG
Cisco_PIX	Cisco PIX
IBM_DB2	IBM DB2 9.x Audit Log
Juniper_NSM	Juniper NSM 2009 Syslog
Microsoft_DHCP	Microsoft DHCP for 2008 v6 log files

Source type	Description
syslog	Simple Syslog
TippingPoint_SMS	Tipping Point SMS 2.5 Syslog
VMWare_ESX	VMWare ESX Syslog

To add a source type:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Source Types** tab (right panel).
The Source Types tab, shown in [Figure 5-4 on page 156](#), displays the current source types. You can sort the fields by clicking the column headers.
- 3 Click **Add**.
- 4 Fill in the fields to define the source type:

Table 5-8 Source Type Fields

Field	Description
Name	The name of the source type.
Description	A description of the source type.
Parser	The parser you want to associate with this source type. If the parser you need does not appear in the drop-down list, you can add one. For information on how to add a parser, see “Parsers” on page 159 .
Event Time Location	A regular expression describing the timestamp in the log file. For example: <code>.*\ [(.*?)\] .*</code> This expression specifies that the timestamp is found inside the first set of square brackets on each line. The first capturing group (the part of the regex in parentheses) is the part that is then parsed using the Date/time format. You can specify that there is no timestamp in the log file with <code>''</code> .
Event Time Format	A regular expression describing the date and time format in the log file. For example, <code>dd/MMM/yyyy:HH:mm:ss Z</code> You can specify that there is no timestamp in the log file with <code>''</code> . For more information about event time, see “Time Range” on page 73 and “Date and Time Specification” on page 153 .

Table 5-8 Source Type Fields (Continued)


Field	Description
Multiline Regex	A regular expression describing how to recognize when adjacent lines are of the same event or when a new event starts. For example if each event starts with the date in the format, yy-MM-dd HH:mm:ss.SSS you could use (\d+-\d+-\d+\d+:\d+:\d+.\d+) to indicate the start of a new event.
Locale	Select a locale from the pulldown list, such as English (United States), Chinese (Hong Kong), Chinese (Taiwan), and so on. This is locale of the data OLI should find in the file.

- 5 Click **Save**.

To edit a source type:


- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Source Types** tab (right panel).

The Source Types tab, shown in [Figure 5-4 on page 156](#), displays the current source types. You can sort the fields by clicking the column headers.

- 3 Locate the source type that you want to update and click the Edit icon () on that row.



Note

The Edit icon () is not available for out-of-box source types. You can copy the source type and make a similar one instead.

- 4 Edit the fields as appropriate.

Source type fields are documented in [Table 5-8 on page 157](#).

- 5 Click **Save**.
- 6 Disable and then re-enable any receivers that use this source type.




Note

Changes in source type are not reflected in the associated receivers until you have re-enabled them.

To copy a source type:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Source Types** tab (right panel).

The Source Types tab, shown in [Figure 5-4 on page 156](#), displays the current source types. You can sort the fields by clicking the column headers.

- 3 Locate the source type that you want to copy and click the Copy icon () on that row.
- 4 Enter a name for the new source type and edit the fields as appropriate.

Source type fields are documented in [Table 5-8 on page 157](#).

- 5 Click **Save**.

To delete a source type:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Source Types** tab (right panel).

The Source Types tab, shown in [Figure 5-4 on page 156](#), displays the current source types. You can sort the fields by clicking the column headers.

- 3 Locate the source type that you want to delete and click the Remove icon (✖) on that row.



The Remove icon (✖) is not available for out-of-box source types. You can only remove source types that you added.

- 4 Click **OK** to confirm the removal.

Parsers

Parsers enable you to extract and manipulate raw events (non-CEF data) from different sources in your network environment. Once you have parsed event fields, you can easily search for data, chart it, and perform other operations on it. One user with in-depth knowledge of the events can create the parser, and then all users who look at those events will get the benefit of that work.

Parsers provide you with a simple way to read events. Instead of looking at raw event data and trying to figure out what it means, you can use a parser to extract portions of non-CEF events into fields. However, the fields created by the parser are available only for search operations, and are not added to the OLI schema.

You can use a parser either of the following ways:

- Use the parser with a source type—You can associate the parser with a source type to extract any set of fields in any kind of event. For more information, see and [“Source Types” on page 155](#).
- Use the parse command in a search—During a search, you can use the parse command to extract fields from events and use other search operators (such as where, chart, top, etc) to further refine the search or manipulate the data in the fields. This is particularly useful for IT operations and other customers who need to extract and manipulate raw event data.

Using Parsers with Source Types

OLI provides a number of pre-configured parsers with associated source types. You can also define new parsers and associate them with source types. Only one parser can be associated with a source type, however, multiple source types can use the same parser. Out-of-box parsers cannot be edited or deleted, but you can copy them to make a similar parser to meet your needs. You can edit or delete custom parsers as desired.

Name	Parser Type	Description	Definition
Apache_access	Rex Parser	Apache Access Log Parser	{<SourceHost>[S+]{<Identity>[S+]?[S-]*[!@<Date>.*?]}[S+]{<Method>[S+]{<URL>.*?}[S+]{<HTTPVersion>.*?}[S+]{<ReturnCode>[S+]?[<Length>[S+]?}
Apache_error	Rex Parser	Apache Error Log Parser	{[!@<timestamp>[S+]{S+ \d{1,2} \d{4}}[S+]{!@<severity>[a-z]*}[S+]{!@<client ip>[S+]{<clientip>[!@.]*}[S+]{!@<File>[S+]{!@<does not exist>[S+]{!@<filename>[S+]}?<message>.*}
audit_log	Rex Parser	Syslog parser for Audit Log files	type={<type>[S+]{S+msg=audit{!@<epoch>[S+]{!@<eventid>[S+]}?}[S+]{!@<start>[S+]{!@<version>[S+]{!@<format>[S+]}?}[S+]{!@<normal>[S+]{!@<halt>[S+]{!@<send>[S+]{!@<audit>[S+]{!@<res>[S+]}?}[S+]{!@<res>[S+]}?}[S+]{!@<audit...>
Bluecoat_proxy	Rex Parser	Bluecoat Proxy SG Parser	{<Service>ProxySG}{<ID>[A-Fa-f0-9]*}{<Message>.*}{!@<errorcode>[S+]}{<severity>NORMAL_EVENT SEVERE_ERROR CONFIGURATION_EVENT}{<error_filename>[S+]}{<line>[S+]}
Cisco_PIX	Rex Parser	Cisco PIX Parser	{<priority>[S+]}{<timestamp>[S+]{!@<id>[S+]{!@<devicehostname>[S+]}{!@<product>PIX(FWSM ASA)}{<messageclass>[!@.]*}?{<pixseverity>[S+]}{<pixmessageid>[S+]}{<pixmessage>.*}
IBM_DB2	Rex Parser	IBM DB2 9.x Audit Log Parser	{<Timestamp>[!@.]*}[!@<Category>[S+]}{!@<Event>[S+]}{!@<Correlator>[S+]}{!@<Status>[S+]}{!@<Userid>[S+]}{!@<Authid>[S+]}
Juniper_NSM	Rex Parser	Juniper NSM 2009 Syslog parser	{<logDayid>[S+]}{<logRecordid>[S+]}{<nsmReceivedTime>[S+]}{<deviceGeneratedTime>[S+]}{<deviceDomain>[S+]}{<deviceDomainVersion>[S+]}{<deviceHostName>[S+]}{<...>
Microsoft_DHCP	Rex Parser	Microsoft DHCP parser for 2008 v6 log files	{<Eventid>[S+]}{<Date>[S+]}{<Time>[S+]}{<EventName>[S+]}{<Address>[S+]}{<HostName>[S+]}{<ErrorCode>[S+]}{<DuidLength>[S+]}{<DuidBytes>[S+]}{<UserName>.*}
		Syslog parser for syslog files	{<timestamp>[S+]{!@<host>[S+]}{!@<module>[S+]}{!@<...>

Figure 5-5 Parsers tab

Using the Parse Command

The `parse` command can be used to invoke a parser on any non-CEF events that are returned by a search. It applies the definition of the parser, such as the regular expression of a rex parser, to each event. Then it adds the fields that are extracted by that regular expression to the fields that are being passed through. For a REX parser, this is functionally the same as having a `rex` command with the same regular expression as the definition of the parser, so you can think of a REX parse command as invoking a saved rex expression.

For more information about the `parse` command, see [“parse” on page 281](#). For information about searching in general, see [“Searching and Analyzing Events” on page 67](#).

Working with Parsers

You can define two types of parsers—a REX parser or an Extract parser. Before adding the parser, you need to define the query you want to use for parsing events.

For a Rex parser, one way to do this is to use the `rex` search operator to test and adjust a regular expression until it returns the desired fields from the events that you want it to handle. Then copy the rex expression and paste it into the parser's Definition field. For an Extract parser, use the `extract` operator. For more information about the search operators, see [“parse” on page 281](#), [“rex” on page 286](#), and [“extract” on page 277](#).

The parser used in a search will be displayed in the Parser column of the search results. If the event was parsed, this field contains the name of the parser. If the

event was not parsed successfully, this field contains “Not parsed”. If no parser is defined for the source type or if there is no source type, the field is blank.

To add a parser:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Parsers** tab (right panel) or Click **Search** (left panel) > **Parsers** tab (right panel).

The Parsers tab, shown in [Figure 5-5 on page 160](#), displays the current parsers. You can sort the fields by clicking the column headers.

- 3 Click **Add**.
- 4 Enter a name for the parser.
- 5 Choose the Parser Type from the drop-down list.
- 6 Click **Save**.

The fields displayed in the Edit Parser dialog box according to the type of parser.

- 7 Fill in the fields for the parser.

Table 5-9 Parser fields


Field	Description
Name	The name of the parser. Enter a new name if you want to change the existing name.
Description	A meaningful description of the purpose of the parser.
Rex parsers only	
Definition	The rex expression that you want to use to parse events.
Extract parsers only	
Pair Delimiter	The characters separate key/value pairs within an event. Enter only the separator characters, for example: \\,
Key/Value Delimiter	The characters that separate the key from the value. Enter only the delimiter character, for example: =
Fields	The list of field names to use when parsing events. Enter the field names, separated by comma (.). For example, to parse events like: foo=abc, bar=xyz, baz=def Enter: foo,bar,baz

- 8 Click **Save**.


To edit a parser:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Parsers** tab (right panel) or Click **Search** (left panel) > **Parsers** tab (right panel).

The Parsers tab, shown in [Figure 5-5 on page 160](#), displays the current parsers. You can sort the fields by clicking the column headers.

- 3 Locate the parser that you want to update and click the Edit icon () on that row.



The Edit icon () is not available for out-of-box parsers. You can copy the parser and make a similar one instead.

- 4 Edit the parser fields as appropriate.


The fields displayed in the Edit Parser dialog box according to the type of parser. Parser fields are documented in [Table 5-9 on page 161](#).

- 5 Click **Save**.

To copy a parser:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Parsers** tab (right panel) or Click **Search** (left panel) > **Parsers** tab (right panel).

The Parsers tab, shown in [Figure 5-5 on page 160](#), displays the current parsers. You can sort the fields by clicking the column headers.

- 3 Locate the parser that you want to copy and click the Copy icon () on that row.

The fields displayed in the Edit Parser dialog box according to the type of parser.

- 4 Enter a name for the new parser and edit the fields as appropriate.

Parser fields are documented in [Table 5-9 on page 161](#).

- 5 Click **Save**.

To delete a parser:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Parsers** tab (right panel) or Click **Search** (left panel) > **Parsers** tab (right panel).

The Parsers tab, shown in [Figure 5-5 on page 160](#), displays the current parsers. You can sort the fields by clicking the column headers.

- 3 Locate the parser that you want to delete and click the Remove icon (✖) on that row.



The Remove icon (✖) is not available for out-of-box parsers. You can only remove parsers that you added.

- 4 Click **OK** to confirm the removal.

Example: Creating an Extract Parser

Suppose you want to create a parser to find the contents of the INT, MAC, DST, and SRC fields of a log like the one below.

```
Jul 12 14:30:31 n15-214-128-h92 kernel: IN=eth2
MAC=00:24:e8:60:cb:82:00:50:56:92:2a:d5:08:00 SRC=15.214.157.89 |
DST=15.214.128.92 LEN=52 TOS=0x00 PREC=0x00 TTL=128 ID=21408 DF
PROTO=TCP SPT=56978 DPT=443 WINDOW=8192 RES=0x00 SYN URGP=0
```

```
Jul 12 14:30:31 n15-214-128-h92 kernel: IN=eth2 |
MAC=00:24:e8:60:cb:82:00:50:56:92:2a:d5:08:00 | SRC=15.214.157.89 |
DST=15.214.128.92 LEN=52 TOS=0x00 PREC=0x00 TTL=128 ID=21408 DF
PROTO=TCP SPT=56978 DPT=443 WINDOW=8192 RES=0x00 SYN URGP=0
```

```
Jul 12 14:30:31 n15-214-128-h92 kernel: IN=eth2 |
MAC=00:24:e8:60:cb:82:00:50:56:92:2a:d5:08:00 | SRC=15.214.157.89 |
DST=15.214.128.92 LEN=52 TOS=0x00 PREC=0x00 TTL=128 ID=21408 DF
PROTO=TCP SPT=56978 DPT=443 WINDOW=8192 RES=0x00 SYN URGP=0
```

In this sample log, the field values are indicated with an equal sign (=), and fields are delimited by pipe (|) and colon (:). You could use the following query to search for the contents of the IN, MAC, DST, and SRC fields.

```
extract pairdelim= "|:" kvdelim= "=" fields= "IN,MAC,DST,SRC"
```

The following steps describe how to make an extract parser using that query.

To create an example extract parser:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Parsers** tab (right panel) or Click **Search** (left panel) > **Parsers** tab (right panel).
- 3 Click **Add**.

The Add Parser dialog box opens.

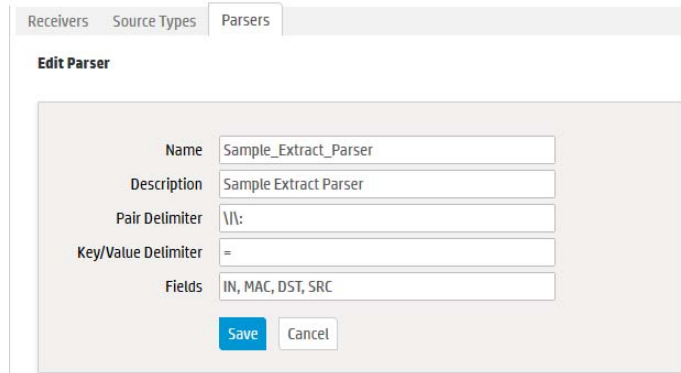
- 4 Enter a Name and select the Parser Type. For the example, enter:

Name: Sample_Extract_Parser

Parser Type: Extract Parser

- 5 Click **Save**.

The Edit parser dialog box opens.



- 6 Enter the Pair Delimiter, Key value, and Fields for the parser. For the example, enter:

Pair Delimiter: \\| :

Key/Value Delimiter: =

Fields: INT, MAC, DST, SRC



You need to escape the pipe (|) and the colon (:) with backslash (\).

- 7 Click **Save**. The Parsers tab displays the new parser.

Name	Parser Type	Description	Definition
Sample_Extract_Parser	Extract Parser	Sample Extract Parser	pair delimiter [\s-] key/value delimiter [-] fields [IN, MAC, DST, SRC]
Apache_access	Rex Parser	Apache Access Log Parser	(?<SourceHost->[^\s-]*) (?<Identity->[^\s-]*) (?<Date->[^\s-]*) (?<Method->[^\s-]*) (?<URL->[^\s-]*) (?<HTTP/>[^\s-]*) (?<HTTPVersion->[^\s-]*) (?<ReturnCode->[^\s-]*) (?<Length->[^\s-]*)
Apache_error	Rex Parser	Apache Error Log Parser	\[(?<timestamp->[^\s-]*) (?<id->[^\s-]*) (?<severity->[^\s-]*) (?<client->[^\s-]*) (?<clientip->[^\s-]*) (?<file->[^\s-]*) (?<exists->[^\s-]*) (?<filename->[^\s-]*) (?<message->.*)\]
audit_log	Rex Parser	Syslog parser for Audit Log files	type=(?<type->[^\s-]*) msg=audit(?<epoch->[^\s-]*) (?<eventid->[^\s-]*) (?<auditd_start->[^\s-]*) (?<version->[^\s-]*) (?<format->[^\s-]*) (?<auditd->[^\s-]*) (?<normal->[^\s-]*) (?<sendmsg->[^\s-]*) (?<audit->[^\s-]*) (?<res->[^\s-]*) (?<auditd->[^\s-]*)
Bluecoat_proxy	Rex Parser	Bluecoat Proxy SG Parser	(?<Service->ProxySG): (?<ID->[A-Fa-f0-9-]) (?<Message->[^\s-]*) (?<errorcode->[^\s-]*) (?<severity->NORMAL EVENT SEVERE_ERROR CONFIGURATION_EVENT) (?<error_filename->[^\s-]*) (?<line->[^\s-]*)
Cisco_PIX	Rex Parser	Cisco PIX Parser	(?<priority->[^\s-]*) (?<timestamp->[^\s-]*) (?<id->[^\s-]*) (?<deviceHostNames->[^\s-]*) (?<Product->PIX FW5MIA5A) (?<MessageClass->[^\s-]*) (?<PixSeverity->[^\s-]*) (?<PixMessageId->[^\s-]*) (?<PixMessage->.*)
IBM_DB2	Rex Parser	IBM DB2 9.x Audit Log Parser	(?<Timestamp->[^\s-]*) (?<Correlator->[^\s-]*) (?<Status->[^\s-]*) (?<UserId->[^\s-]*) (?<AuthId->[^\s-]*)
Juniper_NSM	Rex Parser	Juniper NSM 2009 Syslog parser	(?<logDayid->[^\s-]*) (?<logRecordid->[^\s-]*) (?<nsmReceivedTime->[^\s-]*) (?<deviceGeneratedTime->[^\s-]*) (?<deviceDomainVersion->[^\s-]*) (?<deviceHostName->[^\s-]*) (?<deviceDomain->[^\s-]*) (?<deviceDomainVersion->[^\s-]*) (?<deviceHostName->[^\s-]*) (?<deviceDomain->[^\s-]*)
Microsoft_DHCP	Rex Parser	Microsoft DHCP parser for 2008 v6 in file	(?<EventId->[^\s-]*) (?<Date->[^\s-]*) (?<Time->[^\s-]*) (?<EventName->[^\s-]*) (?<Address->[^\s-]*) (?<HostName->[^\s-]*) (?<ErrorCode->[^\s-]*) (?<DuidLength->[^\s-]*) (?<DuidBytes->[^\s-]*) (?<UserName->[^\s-]*)

Event Output

Use the Event Output section to manage the forwarders that send stored events to other destinations, including ArcSight Manager.

Name	Type	Filter Type	IP/Host	Port	Query
Forwarder1	UDP Forwarder	Unified Query	123.123.23.123	514	((deviceVendor = "Microsoft" AND deviceProduct = "Microsoft Windows") OR (deviceVendor = "SaberNet" AND deviceProduct = "NT syslog") ...)

Figure 5-6 Event Output screen

Forwarders

Forwarders send all events, or events which match a particular filter, on to a particular host. The ability to define a different filter for each forwarder allows OLI to divide traffic among several destinations. For example, because OLI can handle much higher event rates than ArcSight Manager, OLI might be used to forward events to a number of ArcSight Managers. Forwarder filters make it possible to split the flow between the Managers, using one forwarder for each Manager. Additionally, forwarding enables you to send a subset of events to other destinations for further processing while maintaining all events on OLI for long-term storage.

The forwarding filter is a query that searches for matching events, optionally within a time range. You can create two types of forwarder filters—continuous and time-range bound.

A continuous filter constantly evaluates the incoming events and forwards the matching ones to the specified destination.

A time-range bound filter uses a time range in addition to the specified condition to determine whether an event should be forwarded to the destination. If the event falls within the specified time range and matches the specified condition, it is forwarded; otherwise, it is not. The OLI receipt time of an event is used to determine whether an event will be forwarded to a destination when a forwarder filter specifies a time range by which events are evaluated for forwarding. Once a forwarder has forwarded all events within a time range, it does not forward any more events.

A forwarder only forwards events from the OLI that it is configured on; it cannot forward events from peer OLIs.

A forwarder's operation can be paused and resumed at any point in time. When a forwarder resumes operation, forwarding resumes from the last checkpoint that was established before the forwarding operation was paused.

You can also disable and re-enable a forwarder. When you re-enable a forwarder, all previously established checkpoints are removed and forwarding starts over again as per the forwarder configuration—forwarders with continuous filters start from the current time, while forwarders with time-range bound filters start from beginning of the configured time range.

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

- **UDP:** UDP forwarders forward events by using the User Datagram Protocol.
- **TCP:** TCP forwarders forward events by using the Transmission Control Protocol.
- **Connector Forwarder:** Connector forwarders send events to the OLI Streaming Connector.

As a best practice, do not add more than 10 regular expression forwarders on an OLI. Even though each additional forwarder improves the forwarding rate, the relation is not proportional. In high EPS (events per second) situations or situations where other resource-intensive features are running in parallel (alerts, and several search operations) and the forwarding filter is complex, adding too many forwarders may reduce performance because forwarders have to compete for the same OLI resources besides competing for the same built-in connector for forwarding.

You can also specify indexed search queries (known as Unified Queries). Doing so enables you to take advantage of the indexing technology to quickly and efficiently search for events to forward.



Unified query-based forwarders forward events once they have been indexed. Therefore, these forwarders can exhibit “bursty” behavior because indexing occurs in batches on OLI. You might notice the bursty behavior in the EPS out gauge (on top of the OLI interface screen)—the gauge will display high EPS level as a burst of data is forwarded and then drop back to normal level.

To create a forwarder:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Output** in the left panel.
- 3 In the Forwarders tab, click **Add** to display the following form.

- 4 Enter a name for the new forwarder and choose the forwarder type appropriate for your need: UDP Forwarder, TCP Forwarder, or Connector Forwarder type.
- 5 Select the type of forwarding filter you will specify for this forwarder—Unified or Regular Expression. Select “Unified” if you want to specify an indexed search query or “Regular Expression” to specify a regular expression query.
- 6 Click **Next**.
- 7 Enter additional, type-specific information as described in [Table 5-10, “Forwarder Parameters,”](#) on page 168. Click **Save**.
- 8 New forwarders are initially disabled. Click the disabled icon (🚫) to enable the new forwarder.

Table 5-10 Forwarder Parameters

Parameter	Forwarder Types	Description
Name	All	The name that you entered in the previous screen is displayed automatically. If you want to change the name, make the change on this screen.
Query	All	<p>Enter the query to used to filter events that the forwarder will forward, or select a filter from the Filter list below. Forwarder queries can be constrained by device groups and storage groups, but not by Peers.</p> <p>If you selected Unified Query in the previous screen, enter an indexed search query that includes full-text and field-based indexed fields. You can click the Advanced Search link to access the Search Builder tool to build an indexed query. (See “Accessing Search Builder” on page 86 for more information.)</p> <p>The unified query you specify must follow these guidelines, or you will not be able to save the query and thus the forwarder:</p> <ul style="list-style-type: none"> Queries in the following format are valid; all other formats are not allowed. <p style="margin-left: 40px;"><code>(full-text terms field search)* regex</code></p> <p>That is, the query must only contain full-text (keyword) and field-based query elements; it cannot contain any aggregation search operators, or operators that process the searched data further to refine the search. For example, chart, sort, eval, top, and so on.</p> <p>Therefore, this is a valid query: failed message CONTAINS “failed device”</p> <p>However, this is an invalid query: failed message CONTAINS “failed device” sort deviceEventCategory</p> <p>The query can contain the <code>regex</code> operator after a pipeline character (<code>()</code>). Therefore, this is a valid query for a forwarder: failed message CONTAINS “failed device” regex deviceEventCategory = “fan”</p> All search terms (except the “regex” portion) in a query must be indexed. If a query contains full-text (keyword) terms, full-text indexing must be enabled. Similarly, if the query contains a field, field-based indexing must be enabled and the specified field must be indexed. <p>If you selected Regular Expression in the previous screen, specify a regular expression in this text box. See “Searching for Events on OLI” on page 96.</p>



Table 5-10 Forwarder Parameters (Continued)

Parameter	Forwarder Types	Description
Filters	All	<p>Instead of specifying a unified query, you can select a filter from the Filters list. The Filters list contains all saved filters and the predefined system filters on your OLI. Select a filter that meets the validity guidelines described in “Query” on page 168. Otherwise, the user interface will display an error when you save the forwarder definition. <i>You can only select one unified query filter per forwarder.</i></p> <p>Similarly, when creating a regular expression based filter, select a filter from this list. <i>You can select multiple filters for a regular expression based forwarder.</i></p>
Filter by time range	All	<p>If you are creating a continuous filter, which continuously evaluates incoming events and forwards the matching ones, skip this parameter. In this case, the query is run continuously and forwarding continues until you pause it.</p> <p>If you are creating a time range bound filter, check this box to specify a time range of events that the forwarder will forward. If you enter a time range, the forwarder sends events that are within that time range and stops.</p> <p>When you check this box, the Start and End dates and Time fields are displayed.</p> <p>Start must be earlier than End. Specifying a time in the future changes that field to the current time. For example, specifying a Start of the current day at 7 a.m. and an End of current day at 7 p.m. will produce events with timestamps from 7 a.m. to the time the filter is saved (that is, earlier than 7 p.m.).</p>
Source Type	Connector	<p>Select from the pulldown list of log file types, including:</p> <ul style="list-style-type: none"> • Apache HTTP Server Access • Apache HTTP Server Error • IBM DB2 Audit • Juniper Steel-Belted Radius • Microsoft DHCP Log • Other <p>Note: The Source type must be the same in receiver, forwarder, and SmartConnector.</p> <p>A receiver can only receive events of a single source type. Set up separate receivers for each type of log file.</p>

Table 5-10 Forwarder Parameters (Continued)

Parameter	Forwarder Types	Description
Preserve Syslog Timestamp	UDP, TCP	Set to true to preserve the syslog timestamp. The default is true--the timestamp is the original receipt time of the event. If set to false, original timestamp is replaced with OLI's receipt time.
Preserve Original Syslog Sender	UDP, TCP	Set to true to send the event as-is, without inserting OLI's IP address in the hostname (or equivalent) field of the syslog event. The default is true. If set to false, OLI's information is inserted in the hostname (or equivalent) field of the syslog event.
IP/Host	UDP, TCP, Connector	The IP address or host name of the destination that will receive forwarded events. Note: You cannot configure an OLI forwarder to send data to the same OLI on which it is configured.
Port	UDP, TCP, Connector	The port on the destination to which the forwarder will forward events. The default port is 514.
Connection Retry Timeout	TCP, Connector	The time, in seconds, to wait before retrying a connection. The default is 5 seconds.

To edit a forwarder:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Output** in the left panel.
- 3 In the Forwarders tab, locate the forwarder you want to edit.
- 4 If the forwarder is enabled, click the enabled () icon to disable it.
- 5 Click the Edit icon ().

The following screen shows the Edit Forwarder screen for a regular expression based forwarder. The Edit Forwarder screen for a Unified Query forwarder lists the Unified Query based filters and the Query text box only allows you to specify one query.

The screenshot shows the 'Edit Forwarder' configuration interface. At the top, there are tabs for 'Forwarders' and 'Certificates'. The main area is titled 'Edit Forwarder'. It contains the following fields and options:

- Name:** Forwarder1
- Query:** An empty text box with a dropdown arrow.
- Filters:** A list box containing the following items:
 - Configuration - Configuration Changes (Unified)
 - Events - Event Counts by Destination
 - Events - Event Counts by Source
 - Events - High and Very High Severity Events (Unified)
 - Logins - Successful Logins (Unified)
 - Network - DHCP Lease Events
 - Network - Port Links Up and Down
 - Network - Protocol Links Up and Down
 - SystemAlert - CPU Utilization Above 90% (Unified)
 - SystemAlert - CPU Utilization Above 95% (Unified)
- Filter by time range:** An unchecked checkbox.
- Preserve Syslog Timestamp:** A dropdown menu set to 'true'.
- Preserve Original Syslog Sender:** A dropdown menu set to 'true'.
- IP/Host:** An empty text box.
- Port:** A text box containing '514'.
- Buttons:** 'Save' (blue) and 'Cancel' (white) buttons.

Figure 5-7 Specifying Query Terms, Filters, and other forwarder parameters

- 6 Edit the information in the form, as described in [Table 5-10 on page 168](#), and click **Save**.
- 7 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 Output** in the left panel.
- 3 Locate the forwarder that you want to delete in the table.
- 4 If the forwarder is enabled, click the enabled (✔) icon to disable it.
- 5 Click the Delete icon (✖).
- 6 Click **OK** to confirm the delete.

To pause a forwarder:


- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Output** in the left panel.
- 3 Locate the forwarder that you want to pause from the list of forwarders.
- 4 Click the Pause icon (⏸).

To resume a forwarder:

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

- 3 Locate the forwarder whose operation you want to resume.
- 4 Click the Resume icon (.


To disable a forwarder:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Output** in the left panel.
- 3 Locate the forwarder that you want to disable.
- 4 Click the enabled icon (.

To enable or re-enable a forwarder:



Wait a few minutes to disable a forwarder that was just enabled. Likewise, wait before enabling a forwarder that has just been disabled. Background tasks initiated by enabling or disabling a forwarder can produce unexpected results if they are interrupted.

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Event Output** in the left panel.
- 3 Locate the forwarder that you want to enable or re-enable.
- 4 Click the disabled icon (.

Alerts

Alerts respond to events or specified event patterns with optional notification. Event patterns are specified events that occur above a particular frequency (a threshold number of events in a specified time period).

Alerts can be generated for internal events such as storage capacity warnings or for user-determined event patterns such as an alert is generated when five events from a specific device contain the word “unauthorized” within a five minute interval.

OLI provides two types of alerts:

- Real time alerts, discussed in [“Configuring and Managing Real Time Alerts” on page 176](#).
- Saved Search Alerts, discussed in [“Creating and Managing Saved Search Alerts” on page 179](#).

The following table compares the two types of alerts.

Real Time Alerts	Saved Search Alerts
<p>No limit on the number of alerts that are defined.</p> <p>A maximum of five alerts can be enabled at any time.</p>	<p>Any number of alerts can be defined. All defined alerts are enabled and effective, however, a maximum of 50 alerts can run concurrently.</p>
<p>No limit on the number of configured e-mail destinations; however, you can only set one SNMP and one Syslog destination.</p>	<p>No limit on the number of configured e-mail destinations; however, you can only set one SNMP and one Syslog destination.</p>
<p>Only regular expression queries can be specified for these alerts.</p>	<p>Queries for these alerts are defined using the flow-based search language that allows you to specify multiple search commands in a pipeline format, including regular expressions. Aggregation operators such as chart and top cannot be included in the search query.</p>
<p>Alerts are triggered in real time. That is, when specified number of matches occurs within the specified threshold, an alert is immediately triggered.</p>	<p>These alerts are triggered at scheduled intervals. That is, when a specified number of matches occurs within the specified threshold, an alert is triggered at the next scheduled time interval.</p>
<p>To define a real time alert, you specify a query, match count, threshold, and one or more destinations.</p>	<p>To define a Saved Search Alert, you specify a Saved Search (which is a query with a time range), match count, threshold, and one or more destinations.</p>
<p>A time range is not associated with the queries defined for these alerts. Therefore, whenever the specified number of matches occurs within the specified threshold, an alert is triggered.</p>	<p>A time range (within which events should be searched) is specified for the query associated with these alerts. Therefore, specified number of matches within the specified threshold (in minutes) must occur within the specified time range. You can also use dynamic time range (for example, \$Now-1d, \$Now, and so on).</p>
	<p>For example, if a Saved Search query has these start and end times:</p>
	<p>Start Time: 5/11/2010 10:38:04</p>
	<p>End Time: 5/12/2010 10:38:04</p>
	<p>And, the number of matches and threshold are the following:</p>
	<p>Match count: 5</p>
	<p>Threshold: 3600</p>
	<p>Then, 5 events should occur in one hour anytime between May 11th, 2010 10:38:04 a.m. and May 12th, 2010 10:38:04 for this alert to be triggered.</p>

Alert Triggers and Notifications

An alert is triggered if a specified number of matches occurs within the specified threshold (time interval in seconds). When an alert is triggered, OLI creates an alert event containing the triggering events or event IDs, and sends notification through previously configured destinations—e-mail addresses, SNMP server, Syslog server, and ArcSight Manager.

By default, only alert notifications sent to e-mail destinations include all matching events that triggered the alert. You can configure your OLI to include matched events for SNMP and Syslog. However, that kind of configuration is only possible through the command-line interface of the OLI; therefore, please contact customer support for instructions.

When are Alert events triggered?

You also specify a time window and a number of matching events. When that number of matching events is detected within the time window, an alert event is triggered.

OLI resets the count after detecting 100 matching events. Therefore, all events that occur in the time window will not necessarily be recorded in an alert. For example, if you configure the alert to be sent when there are 20 matching events in 2 minutes, and 152 events occur within two minutes, you will get 7 alerts, and 12 matching events will not be included in any alert. In this situation, the following alert events are triggered:

- Alert 1 has 20 matching events.
- Alert 2 has 40 matching events.
- Alert 3 has 60 matching events.
- Alert 4 has 80 matching events.
- Alert 5 has 100 matching events (1-100).
- Alert 6 has 20 matching events (101-120).
- Alert 7 has 40 matching events (101-140).

The remaining 12 events are being held, waiting to meet the threshold of 20 more events in a 2 minute interval.

Receiving Alert Notifications

In order to receive notification of an alert, set up the alert to be sent to a previously configured destination, such as an e-mail address, SNMP server, Syslog server, and ArcSight Manager.

By default, only alerts to e-mail destinations include all matched events that triggered the alert. You can configure your OLI to include matched events for SNMP and Syslog. However, such a configuration is only possible through the command-line interface of the OLI; therefore, please contact customer support for instructions.

Sending Notifications to E-mail Destinations

When you send notifications for an alert via e-mail, the e-mail message contains both the trigger alert information and the matched (base) events.

The following is an example of the trigger alert information:

```
Alert event match count [1], threshold [10] sec
```

And the matched event:

```
Event Time [Tue Nov 11 16:46:49 PST 2008]
Event Receipt Time [Tue Nov 11 16:46:50 PST 2008]
Event Device Address [192.168.35.50]
Event Content [Dec 11 10:31:20 localhost
CEF:0|NetScreen|Firewall/VPN||traffic:1|Permit|Low| eventId=590
msg=start_time\= "2004-07-28 15:25:02" duration\=15 policy_id\=0
service\=SSH proto\=6 src zone\=Trust dst zone\=Untrust
action\=Permit sent\=656 rcvd\=680 src\=10.0.111.46
dst\=10.0.113.50 src_port\=54759 dst_port\=22 translated
ip\=192.91.254.2 port\=54759 app=SSH proto=TCP in=680 out=656
categorySignificance=/Normal categoryBehavior=/Access
categoryDeviceGroup=/Firewall categoryOutcome=/Success
categoryObject=/Host/Application/Service art=1165861874880
cat=Traffic Log deviceSeverity=notification act=Permit
rt=1165861874880 shost=n111-h046.qa.hp.com src=10.0.111.46
sourceZoneURI=/All Zones/System Zones/Private Address
Space/RFC1918: 10.0.0.0-10.255.255.255
sourceTranslatedAddress=192.91.254.2 sourceTranslatedZoneURI=/All
Zones/System Zones/Public Address Space/192.0.3.0-192.167.255.255
spt=54759 sourceTranslatedPort=54759 dst=10.0.113.50
destinationZoneURI=/All Zones/System Zones/Private Address
Space/RFC1918: 10.0.0.0-10.255.255.255 dp]
```

Sending Notifications to Syslog and SNMP Destinations

When configuring OLI to send alerts to SNMP and Syslog destinations, you should be familiar with this information:

- OLI supports SNMP 2.0.
- Unlike an e-mail alert, a trigger alert is sent separately from the alert that contains the matched (base) events that triggered the alert.
- All SNMP alerts are sent as SNMP traps; therefore, trigger alerts and their associated matched (base) events are received as SNMP traps on an SNMP destination. The SNMP trap includes the trigger event, but it does not include the events that caused the alert to trigger (matched events). The trigger event does include the event IDs of all the matched events. You can use the event IDs in the trigger alert to identify the associated matched events.

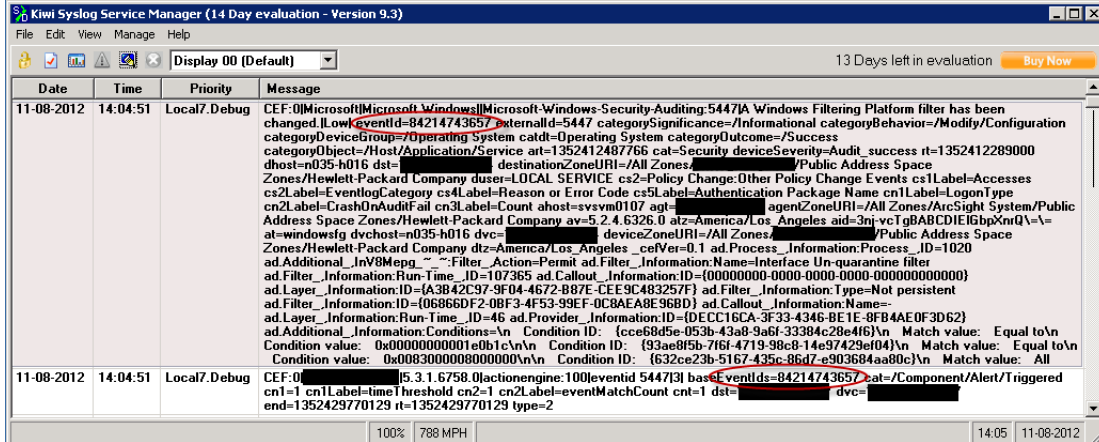


Figure 5-8 A triggered alert event and matching base event shown in Kiwi Syslog Service Manager



Non-CEF events do not contain event IDs. If you need to associate such base events with their trigger alert, send such events to OLI through a connector.

- SNMP uses UDP to send packets. As a result, the order in which alerts arrive at an SNMP destination is not guaranteed.
- When Syslog events are sent using UDP, the order in which the trigger alert and matched events arrive is not guaranteed.

Configuring and Managing Real Time Alerts

This section describes ways to configure and manage real time alerts. For information on Saved Search Alerts, see [“Creating and Managing Saved Search Alerts” on page 179](#).

Creating a Real Time Alert

To create an Alert, you will need to specify a query or filter, event aggregation values (Match count and Threshold, described below), and (optional) one or more notification destinations. If the new Alert will send notification using an SNMP or Syslog, set up those destinations before creating the Alert. To configure the e-mail destination, see [“SMTP” on page 236](#). See also [“Sending Notifications to SNMP Destinations” on page 185](#) and [“Sending Notifications to Syslog Destinations” on page 186](#), .

When you create an alert, it is in disabled state. You can enable it using instructions in [“To Enable or Disable a Real Time Alert:” on page 178](#).

To create a Real Time Alert:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click **Add**. The page shown in [Figure 5-9 on page 178](#) is displayed.

- 4 Enter a name for the new Alert, specify a query or select an available Filter from the list. Events that match this query are candidates for the Alert. Alphanumeric characters and spaces are acceptable, however, some special characters such as % and & are not. For more information on Filters, see [“Filters” on page 189](#).

You can only specify regular expression queries for real time Alerts.

However, a query expression for a scheduled saved alert can contain multiple search commands in a pipeline format, including regular expressions. Aggregation operators such as chart and top cannot be included in the search query. For more information about specifying a regular expression query, see [“The Need to Search Events” on page 67](#).



Tip

To test the validity of an alert query, use the Search user interface. Enter the query in the Search text box in the following format:

Real time Alert: |regex "regex expression"

Scheduled saved alert: _deviceGroup IN ["192.168.22.120 [TCPC]"] name="*[4924TestAlert]*" AND ("192.168.*" OR categoryBehavior CONTAINS Stop)

If the query is valid, cut and paste the regular expression between the double quotes (" ") in the Query text box on the Add Alert page.

- 5 Enter Match count and Threshold values. If the number of candidate events equals or exceeds the Match count within the Threshold number of seconds, the Alert will be triggered.

If you want to be notified when any event matches the filter (for example, for an internal event such as High CPU Temperature), enter a Match count of 1 and a Threshold of 1.



Note

To maintain an optimal size of an alert event, the event does not contain event IDs of all the triggering events if you specify **Match count of 101 or higher**. As a result, the `baseEventCount` field in the event does not reflect the true number of matching events for such alert events.

Triggering events are truncated in multiples of 100. Therefore, if you specify a Match count of 101, only one event is included in the alert event and the `baseEventCount` field value is 1. Similarly, if you specify a Match count of 720, only 20 events are included and the `baseEventCount` field value is 20.

- 6 Enter notification destinations. Enter any combination of:
 - ◆ One or more e-mail addresses, separated by commas
 - ◆ An SNMP Destination—for more information, see [“Sending Notifications to SNMP Destinations” on page 185](#).
 - ◆ A Syslog Destination—for more information, see [“Sending Notifications to Syslog Destinations” on page 186](#).
- 7 Click **Save**.

Figure 5-9 Add Alert dialog

To Enable or Disable a Real Time Alert:

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

Locate the Alert that you want to disable or enable. Click the associated icon (🚫 or ✅) to enable or disable the Alert.



Note

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

To Edit a Real Time Alert:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Locate the Alert that you want to edit and click the Edit icon (✎) on that row.

A screen similar to that shown in [Figure 5-9 on page 178](#) appears. Remember that only alphanumeric characters can be used in an Alert name.

To Remove a Real Time Alert:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Locate the Alert that you want to remove and click the remove icon (✖) on that row.

- 4 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Alert.

To view Real Time Alerts:

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

Click **Alerts** in the left panel. The Alerts list is displayed, as shown in [Figure 5-10](#).

Name	Email Destination(s)	SNMP Destination	Syslog Destination	Query
IT Governance - Access Rights Removed		NONE	NONE	:NOT: storageGroup/Internal Event Storage Group :AND: categoryOutcome=/Success :AND: categoryBehavior=/Authentication/Deletel /Authorizatio...

Figure 5-10 Alert list

Creating and Managing Saved Search Alerts

This section describes ways to configure and manage Saved Search alerts. For information on real time alerts, see [“Configuring and Managing Real Time Alerts” on page 176](#).

Saved Search Alerts are based on the search queries that you have saved on OLI. For detailed information about Saved Search queries, see [“Saved Searches” on page 192](#). For each Saved Search Alert, you configure a match count, threshold, destination, and a schedule at which the alert will be triggered (if specified number of matches occurs within the specified threshold).



Note

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

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


Creating a Saved Search Alert

You can create a Saved Search Alert in two ways:

- From the search results page (Analyze > Search)
- From the Scheduled Searches page (**Configuration > Saved Search > Scheduled Searches**)

To create a Saved Search Alert from the search results page:

- 1 Run a search, as described in [“Searching for Events on OLI” on page 96](#).

- 2 Click the Save icon () and enter the following settings.

Parameter	Description
Name	A name for the query you are saving.
Save as	Whether to save the query as a filter or as a Saved Search. To save the query as a Saved Search Alert, select “Saved search”.
Schedule it	Whether to schedule the alert right now or later. Click to schedule now, or leave blank to schedule later.
Schedule type	Whether the query is being saved as a scheduled search or as a scheduled alert. Scheduled searches run on a predetermined schedule and export results to a pre-specified location. Scheduled alerts run a search on a predetermined schedule and generate an alert if the specified number of events within the specified threshold is found.
Overwrite	If a query with the same name exists, whether that query should be overwritten. If you check this setting and a query with the same name exists, the existing query is overwritten with the one you are currently saving. If you do not check this setting, a warning message is displayed that prompts you to enter another name for the query.

- 3 Click **Save**.

If you checked the “Schedule it” setting in the previous step, you are prompted to choose if you want to edit the schedule, as follows. If you click **OK**, the Edit Scheduled Search page is displayed, as shown in the next step. If you click **Cancel**, the search is saved as a Saved Search but it is not scheduled to run.

- 4 If you checked the Schedule it setting previously, the Edit Scheduled Search page is displayed. This page enables you to define a schedule for the Saved Search job and alert options.

Saved Searches | Scheduled Searches/Alerts | Saved Search Files (OLI)

Edit Scheduled Search

Name:

Schedule:

Every Minutes

Saved Searches:

- dikka test oli bug**
- dikka test 8
- Unit - Rebooted Hosts [Import]
- Unix - Configuration - Configuration Changes (Un
- Unix - Configuration - Configuration Changes (Un
- Unix - Configuration - Configuration Changes (Un
- Unix - Error or Critical or Failure by Host
- Unix - Errors by Name
- Unix - Failure Outcome by Host
- VMware - Errors (table)

 Use ctrl-click to select or deselect items

Job type:

Search Result Export Options

Export Options:

- Save to Operations Log Intelligence
- File format:
- Export directory name:
- Fields:
 - Event Time, Receipt Time, Device, OLI Name, Version, Device Vendor, Device Product, Device Version, Signature ID, Severity
 - All fields
- Include summary:
- Include only CEF events:

To create a Saved Search Alert from the Scheduled Searches page:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.
- 4 Click **Add**.

Saved Searches | **Scheduled Searches/Alerts** | Saved Search Files (OLI)

Add Scheduled Search

Name

Schedule **Everyday**
Hour of day Hours (24 hour format)

Saved Searches
Unit - Configuration - Configuration Changes (Uni
Unix - Configuration - Configuration Changes (Uni
Unix - Configuration - Configuration Changes (Uni
Unix - Error or Critical or Failure by Host
Unix - Errors by Name
Unix - Failure Outcome by Host
VMware - Errors (table)
VMware - Number of Errors per Module and Host
VMware - Number of Errors per Service Name and

Job type **Alert**

Alert Options

Match count

Threshold (sec)

Email address(es)

SNMP destination **NONE**

Syslog destination **NONE**



5 Enter the following information.

Parameter	Description
Name	A name for the Saved Search you are saving.
Schedule	<p>Choose Everyday or Days of Week from the first pulldown menu.</p> <p>If Everyday:</p> <ul style="list-style-type: none"> • EITHER select Hour of Day to specify the hour of the day in 24-hour format • OR select Every to specify the number of hours or minutes after which a Saved Search is performed. Select from the pulldown on the right side of Every to specify Hours or Minutes. By default, the number of hours and number of minutes is set to 15. <p>If Days of Week, enter the days (day 1 is Sunday) in the text box. Then choose Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the second text box.</p> <p>For example, to perform the Saved Search every day at 2 a.m., select Everyday in the first pulldown menu, then choose Hour of Day from the second pulldown menu and enter 2 in the text box. To perform the Saved Search every day at 2 a.m. and 3 p.m., enter 2,15 in the text box.</p> <p>For another example, to perform the Saved Search Tuesdays and Thursdays at 10 p.m., select Days of Week from the first pulldown menu and enter 3,5 for days. Then choose Hour of Day from the second pulldown menu and enter 22 in the text box.</p>
Saved Searches	<p>Select from the list of Saved Searches. If a Saved Search to suit your needs does not exist in the list, click the Saved Searches tab (to the left of Scheduled Searches tab) to define it. For more information about defining a Saved Search query, see “Saved Searches” on page 192.</p> <p>Note: You can only select one Saved Search for each Alert you configure.</p>
Job Type	Select Alert for a Saved Search Alert.
Alert Options	
Match count	Number of events that should be matched in Threshold number of seconds for an alert to be triggered.
Threshold (sec)	Number of seconds within which the “Match count” events should be matched for an alert to be triggered.
Notification destinations are optional. If none is specified, a notification is not sent.	
Email address(es)	(Optional) A comma-separated list of email addresses to which the alert will be sent


Parameter	Description
SNMP destination	(Optional) An SNMP destination to which the alert will be sent. For more information, see “Sending Notifications to SNMP Destinations” on page 185.
Syslog destination	(Optional) A syslog server address to which the alert will be sent. For more information, see “Sending Notifications to Syslog Destinations” on page 186.

- 6 Click **Save**.
- 7 Once a Saved Search Alert is created, you need to enable it. See [“To Enable or Disable a Saved Search Alert:” on page 184.](#)


To Enable or Disable a Saved Search Alert:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.
- 4 Identify the Saved Search Alert that you want to enable.
- 5 Click the associated icon ( or ) to enable or disable the alert.

To edit a Saved Search Alert:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.
- 4 Locate the Saved Search Alert that you want to edit.
- 5 Click the Edit icon () and edit the information. For details about the settings, see [“To create a Saved Search Alert from the Scheduled Searches page:” on page 181.](#)
- 6 Click **Save**.

To remove a Saved Search Alert:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.
- 4 Identify the Saved Search Alert that you want to remove.
- 5 Click the remove icon ().
- 6 Click **OK** to confirm the removal, or click **Cancel** to keep the alert.

To view Saved Search Alerts:

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

- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.

A list of the currently configured Saved Search Alerts is displayed.

Sending Notifications to SNMP Destinations

SNMP Destinations describe how Alert notifications should be sent using Simple Network Management Protocol (SNMP). Set up SNMP Destinations before creating Alerts that will use them. Before configuring an SNMP destinations, you should be familiar with the information in [“Sending Notifications to Syslog and SNMP Destinations” on page 175](#).

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 OLI’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 that you want to remove and click the remove icon (✖) on that row.

- 5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the SNMP Destination.

Sending Notifications to Syslog Destinations

Syslog Destinations describe how Alert notifications should be sent using the comparatively simple syslog protocol. You need to set up Syslog Destinations before creating Alerts that will use them. Before configuring an Syslog destination, you should be familiar with the information in [“Sending Notifications to Syslog and SNMP Destinations” on page 175](#).

To Add a Syslog Destination:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **Syslog Destinations** tab in the right panel.
- 4 Click the **Add** button.
- 5 Enter parameters:

Parameter	Description
Name	A name for this destination.
Type	UDP or TCP Syslog. This choice cannot be edited later.


- 6 Click **Next**. Enter the secondary parameters:

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


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

To Edit a Syslog Destination:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **Syslog Destinations** tab in the right panel.

- 4 Click the Edit icon (). You can edit the parameters of the Syslog Destination except its type.
- 5 Click **Save** to make the changes, or **Cancel** to return to the Syslog Destination table.

To Remove a Syslog Destination:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **Syslog Destinations** tab in the right panel.
- 4 Locate the Syslog Destination that you want to remove and click the remove icon () on that row.
- 5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Syslog Destination.

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 236](#) that can impact a scheduled task.

Maximum number of scheduled tasks that can be defined on OLI: No limit.

Scheduled Tasks

The Scheduled Tasks page, shown in [Figure 5-11](#), 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: <input type="text" value="All"/>			
Task	Type	Schedule	Next Run Time
all-fields	Scheduled search	Every 15 minutes	Dec 24, 2013 11:00:00 PM IST
Savelt Job	Scheduled search	Daily at 3:00	Disabled
some-fields	Scheduled search	Every 15 minutes	Dec 24, 2013 11:00:00 PM IST

Figure 5-11 Scheduled Tasks page

Scheduled Tasks can be created for:

- Saved Search (See [“Scheduled Saved Search”](#) on page 193)
- File Receivers and File Transfer Receivers (See [“Receivers”](#) on page 140)
- Event Archive (See [“Archiving Events”](#) on page 133)
- Configuration Backup (See [“Configuration Backup and Restore”](#) on page 207)

To delete a Scheduled Task:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Scheduled Tasks** in the left panel.
- 3 Click the **Scheduled Tasks** tab in the right panel.
- 4 Locate the Scheduled Task that you want to delete and click the delete icon (✖) on that row.
- 5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Scheduled Task.

Currently Running Tasks

The Currently Running Tasks page displays the Scheduled Tasks that are running at the present time. The table shows task name, type, and the date and time that the task started.

To view Currently Running Tasks:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Scheduled Tasks** in the left panel.
- 3 Click the **Currently Running Tasks** tab in the right panel.
- 4 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 **Configuration** from the top-level menu bar.

- 2 Click **Scheduled Tasks** in the left panel.
- 3 Click the **Finished Tasks** tab in the right panel.
- 4 Filter the list by selecting a specific type of Scheduled Task from the drop-down list, or select **All**.

Filters

The Filters section has two tabs: Filters and Search Group Filters.

Filters

The Filters page provides a convenient place to manage filters. There are two types of filters:

- **Shared**
A shared filter is visible to all OLI users. Once created, any OLI user can use it to search for events. The query you specify for a shared filter can be a Unified query (that uses keywords, indexed, and non-indexed fields) or a Regex query (that specifies a regular expression). Creating Regex Query shared filters are useful for creating real time alerts, which accept only regex queries.
- **Search Group**
Search group filters provide an access control mechanism to limit the events that users in a particular user group can see. The query for these filters can only contain regular expressions. Only users with administrative privileges can create these filters.

A set of pre-defined filters, also known as system filters, exist on your OLI as well. For more information about system filters, see [“System Filters/Predefined Filters” on page 119](#).

To create a filter:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Filters** in the left panel.
- 3 Click the **Filters** tab in the right panel to create a shared filter, or click the **Search Group Filters** tab to create a search group filter. (See [“Filters” on page 189](#) for information about shared and search group filters.)
- 4 Click **Add** to display the following page.
- 5 Enter a name for the new filter in the Name field.
Filter names are case-sensitive.
- 6 If you are creating a shared filter, select **Unified** or **Regex Query**.
For Search Group filters, select **Search Group**.



Non-administrator users cannot create Search Group filters.

- 7 Click **Next**.
- 8 If you selected Unified or Regex Query method in the previous step, enter the query for the new filter.

For Unified queries:

When you type a query, OLI's Search Helper enables you to quickly build a query expression by automatically providing suggestions, possible matches, and applicable operators. See ["Search Helper" on page 93](#) for more information.

OR

Click **Advanced Search** to use the Search Builder Tool to create the query. For details about using the Search Builder Tool, see ["Using the Search Builder Tool" on page 86](#).

For Regex queries:


Enter the regular expression in the Query text box.

- 9 Click **Save**.




If you created a Search Group filter, make sure that you associate it to a user group, as described in ["Search Group Filters" on page 191](#).


To create a filter by copying an existing filter:

- 1 Click the **Configuration** tab, click **Filters** in the sub-menu, and then click the **Filters** tab.
- 2 Locate the filter to copy from the list of filters on the Filters 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, and then click the **Filters** tab.
- 2 Find the filter that you want to edit and click the Edit icon () on that row.
- 3 Change the information in the form and click **Save**.

To delete a filter:

- 1 Click the **Configuration** tab, then click **Filters** in the sub-menu.
- 2 Find the filter that you want to delete in the table.
- 3 Click the Delete icon (). Confirm the delete.

Search Group Filters

Search Group Filters can be used to restrict events visible by members of a user group. A Search Group Filter can be associated with a user group (of type OLI 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.

Name	Filter	Description
Default OLI Search Group	NONE	The default search group allows both local and distributed searches.

Figure 5-12 Search Group Filters Page



In the Search Group Filters page (shown in [Figure 5-12](#)), 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 189](#). To create, edit, or delete User Groups, see [“Users/Groups” on page 251](#).



Only users that are members of a System Admin group can assign Search Group Filters. For more information, see [“Users/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 [“Users/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 189](#). When creating the filter, from the **Type** pull-down menu select the **Search Group** option.
- 3 Click the **Configuration** tab, click **Filters** in the sub-menu, and then click the **Search Group Filters** tab. The page shown in [Figure 5-12](#) is displayed.

- 4 Find the User Group to which to apply a Search Group Filter. Click the edit icon ().
- 5 Select a filter from the pulldown list. (Only Search Group type filters are listed.)
- 6 Click **Save**.

Saved Searches

A Saved Search, like a Saved Filter, recalls a specific query. A Saved Search includes a time range. Also, a saved filter does not include the field set information that determines the fields that are displayed for each event in the search results.

For information about Saved Search Alerts, see [“Alerts” on page 172](#).

You can schedule a saved search to run at a specific interval. For more information, see [“Scheduled Saved Search” on page 193](#).

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 236](#) before adding a Saved Search.

Saved Searches

The Saved Searches tab displays all Saved Searches and supports Adding, Editing, and Deleting Saved Searches.

To add a Saved Search:

- 1 Click the **Configuration > Saved Search**.
- 2 Click **Add** and enter the following parameters:

Parameter	Description
Name	A name for this Saved Search. This name will be used for exported output files, with the Saved Search date and time appended.
Start Time	Absolute date and time of earliest possible event. Alternatively, check Dynamic to specify the start time relative to the time when the Saved Search job is run.
End Time	Absolute or Dynamic date and time of latest possible event, as described above.
Query	Enter the query in the text field or select one or more Filters from the list below the text field. When you type a query, OLI's Search Helper enables you to quickly build a query expression by automatically providing suggestions, possible matches, and applicable operators. See “Search Helper” on page 93 for more information.


Parameter	Description
Local Search	Check this box to limit the Saved Search to the local OLI box. If the Local Search box is left unchecked, the Saved Search will include all Peer OLIs as well as the local OLI.

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


Name	Start	End	Type	Query	Creator
SaveIt Job	\$Now - 10m	\$Now	Unified Query	operations	admin
test	12/22/2013 23:32:28	\$Now	Unified Query	_deviceGroup in ("Operations Log Intelligence Internal Event Device")	admin
Unik - Rebooted Hosts	\$Now - 1d	\$Now	Unified Query	deviceVendor=Unik and name contains "reboot" transaction deviceHostName, startTime chart count by transactionId, deviceHostName chart co...	admin
Unik - Configuration - Configuration Changes (Unified) by Host	\$Now - 1d	\$Now	Unified Query	deviceVendor=Unik and categoryBehavior = "Modify /Configuration" chart count as "Num of Events" by deviceHostName sort -"Num of Events", de...	admin
Unik - Configuration - Configuration Changes (Unified) by Host and Action	\$Now - 1d	\$Now	Unified Query	deviceVendor=Unik and categoryBehavior = "Modify /Configuration" chart count as "Num of Events" by deviceHostName, deviceAction sort devik...	admin
Unik - Configuration - Configuration Changes (Unified) by Host and Significance	\$Now - 1d	\$Now	Unified Query	deviceVendor=Unik and categoryBehavior = "Modify /Configuration" chart count as "Num of Events" by deviceHostName, categorySignificance s...	admin
Unik - Created Users	\$Now - 1d	\$Now	Unified Query	deviceVendor = "Unik" and categoryBehavior = "Authentication/Add" and categoryOutcome = /Success chart count as "Num of Users" by deviceHo...	admin
Unik - Error or Critical or Failure by Host	\$Now - 1d	\$Now	Unified Query	deviceVendor=Unik and ("error" or "critical" or "failure") chart count as "Num of Erros" by deviceHostName sort -"Num of Erros", deviceHostName	admin

Figure 5-13 Saved Search page

To edit a Saved Search:

- Click the **Configuration > Saved Search**.
- Find the Saved Search that you want to edit and click the Edit icon () on that row.
- Change the information in the form and click **Save**.

To delete a Saved Search:

- Click the **Configuration > Saved Search**.
- Find the Saved Search that you want to delete in the table.
- Click the Delete icon (). Confirm the delete.

Scheduled Saved Search

A scheduled Saved Search schedules a Saved Search to be run at a later time. Before you schedule a Saved Search, you must have created or saved at least one Saved Search. A scheduled Saved Search can be also configured to generate an alert. For more information about scheduled Saved Search Alerts, see [“Creating a Saved Search Alert” on page 179](#).

The results of a scheduled Saved Search are written to a file, as described in “[Saved Search Files](#)” on page 197.

To add a scheduled Saved Search:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.
- 4 Click **Add**. The screen shown in [Figure 5-14](#) is displayed.

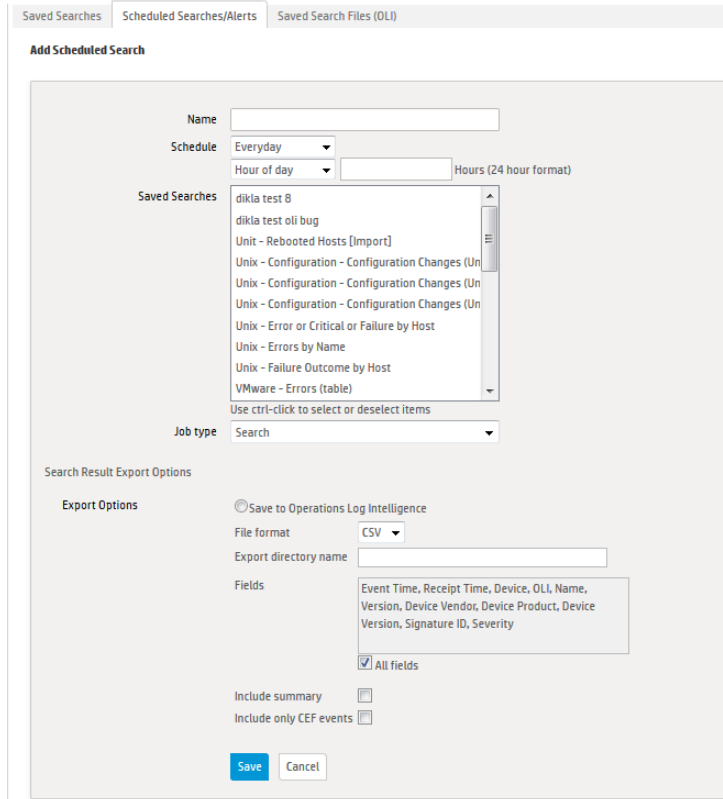


Figure 5-14 Saved Search Jobs page

- 5 Enter the following parameters:


Parameter	Description
Name	A name for this Scheduled Saved Search Job.

Parameter	Description
Schedule	<p>Choose Everyday or Days of Week from the first pulldown menu.</p> <p>If Everyday, select Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the text box.</p> <p>If Days of Week, enter the days (day 1 is Sunday) in the text box. Then choose Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the second text box.</p> <p>For example, to perform the Saved Search every day at 2 a.m., select Everyday in the first pulldown menu, then choose Hour of Day from the second pulldown menu and enter 2 in the text box. To perform the Saved Search every day at 2 a.m. and 3 p.m., enter 2,15 in the text box.</p> <p>For another example, to perform the Saved Search Tuesdays and Thursdays at 10 p.m., select Days of Week from the first pulldown menu and enter 3,5 for days. Then choose Hour of Day from the second pulldown menu and enter 22 in the text box.</p>
Saved Searches	<p>Select from the list of Saved Searches. If a Saved Search to suit your needs does not exist in the list, click the Saved Searches tab (to the left of Scheduled Searches tab) to define it. For more information about defining a Saved Search query, see “Saved Searches” on page 192. When <i>multiple</i> Saved Searches are specified in one scheduled saved search job, the resulting file contains the number of hits for each Saved Search and not the actual events.</p>
Job Type	Select Search for a scheduled Saved Search.
Export Options	The only applicable option is “Save to Logger”, which is preselected for you.
File Format	<p>CSV, for comma-separated values file.</p> <p>PDF, for a report-style file that contains search results as charts and in tables. You must specify a title for the report in the Title field. If the search query contains an operator that creates charts such as chart, top, and so on, charts are included in the PDF file. In that case, you can also set the Chart Type and Chart Result Limit fields. These fields are described later in this table.</p>
Export directory name	<p>Enter the directory path in this field, which can be a path to a local directory or to a mount point on the machine on which the OLI is installed.</p> <p>If a directory of the specified name does not exist, it is created. If a directory of the specified name exists and the Overwrite box is not checked, an error is generated. If the Overwrite box is checked, the existing directory contents are overwritten.</p>
Title (for PDF only)	(Optional) A meaningful name that appears on top of the PDF file. If no title is specified, “Untitled” is included.
Fields	<p>A list of event fields that will be included in the exported file.</p> <p>By default, all fields are included.</p> <p>You can enter fields or edit the displayed fields by deselecting All Fields.</p>

Parameter	Description
Chart Type (for PDF only)	<p>(If the search query includes an operator that creates a chart, this field is meaningful; otherwise, it is ignored.)</p> <p>Type of chart to include in the PDF file. You can select from: Column, Bar, Pie, Area, Line, Stacked Column, Stacked Bar.</p> <p>Note: If the Chart Type is different from the chart displayed on the Search Results screen, the value selected for this option overrides the one shown in the screen. Therefore, the exported PDF contains the chart you specify for this option and not the one shown on the screen.</p>
Chart Result Limit (for PDF only)	<p>(If the search query includes an operator that creates a chart, this field is meaningful; otherwise, it is ignored.)</p> <p>Number of unique values to plot. Default: 10</p> <p>If the configured Chart Result Limit is less than the number of unique values for a query, the top values equal to the Chart Result Limit are plotted. That is, if the Chart Result Limit is 5 and 7 unique values are found, the top 5 values will be plotted.</p>
Include Summary	<p>Check this box to include an event count with the Saved Search, or a total when more than one Saved Search is specified.</p>
Include only CEF Events	<p>Check this box to include only CEF events. Uncheck the box to include all events in the output.</p> <p>For more information about the Common Event Format (CEF), refer to Implementing ArcSight CEF. For a downloadable a copy of this guide, search for "ArcSight Common Event Format (CEF) Guide" on the Protect 724 Community at https://protect724.arcsight.com.</p>

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

To edit a scheduled Saved Search:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.
- 4 Locate the Saved Search Job that you want to edit and click the Edit icon () on that row.
- 5 Change the parameters of the Saved Search Job.
- 6 Click **Save** to update the Saved Search Job, or **Cancel** to abandon your changes.

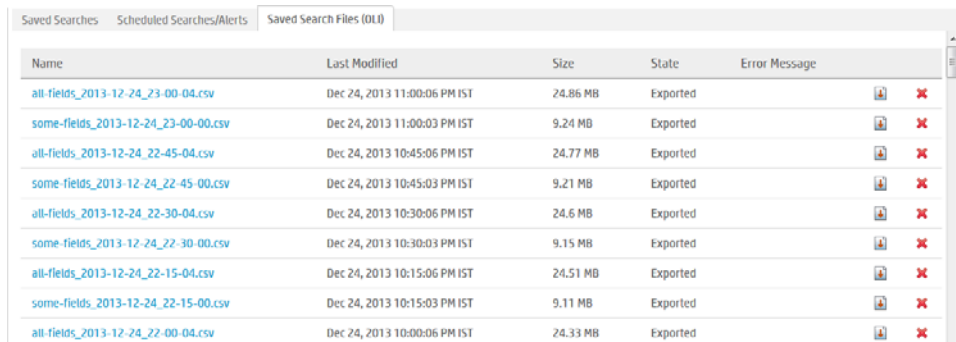
To delete a scheduled Saved Search:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.

- 4 Locate the Saved Search Job that you want to delete and click the delete icon (✖) on that row.
- 5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Saved Search Job.

Saved Search Files

Access Saved Search results that were saved to OLI with the Saved Search Files command. Saved Search Files can be retrieved (streamed to the browser) or deleted.



Name	Last Modified	Size	State	Error Message
all-fields_2013-12-24_23-00-04.csv	Dec 24, 2013 11:00:06 PM IST	24.86 MB	Exported	
some-fields_2013-12-24_23-00-00.csv	Dec 24, 2013 11:00:03 PM IST	9.24 MB	Exported	
all-fields_2013-12-24_22-45-04.csv	Dec 24, 2013 10:45:06 PM IST	24.77 MB	Exported	
some-fields_2013-12-24_22-45-00.csv	Dec 24, 2013 10:45:03 PM IST	9.21 MB	Exported	
all-fields_2013-12-24_22-30-04.csv	Dec 24, 2013 10:30:06 PM IST	24.6 MB	Exported	
some-fields_2013-12-24_22-30-00.csv	Dec 24, 2013 10:30:03 PM IST	9.15 MB	Exported	
all-fields_2013-12-24_22-15-04.csv	Dec 24, 2013 10:15:06 PM IST	24.51 MB	Exported	
some-fields_2013-12-24_22-15-00.csv	Dec 24, 2013 10:15:03 PM IST	9.11 MB	Exported	
all-fields_2013-12-24_22-00-04.csv	Dec 24, 2013 10:00:06 PM IST	24.33 MB	Exported	

Figure 5-15 Saved Search Files page

Search

The search screen enables you to:

- Add search indexes for field query search operations
- Tune advanced search options
- View and delete custom field sets
- View default schema
- View custom schema fields
- View and end currently running search tasks
- View and add parsers for specific log types

For general search information, including how to search, see [“Searching and Analyzing Events” on page 67](#).

Adding Search Indexes

See [“Indexing” on page 111](#) for more information.

Tuning Advanced Search Options

The advanced search options support internationalization (i18n) choices. To change these options, click **Configuration > Search > Search Options**.

The following table lists the advanced search options you can view and configure. Several of the options on this screen will require you to restart your OLI.

Option	Description
Field Search Option	
Case sensitive	<p>Default: Yes</p> <p>Controls whether to differentiate between upper- and lower-case characters during a search. When this option is set to No, searching for "login" will find "login," "Login," and "LOGIN".</p> <p>Notes:</p> <ol style="list-style-type: none"> 1 Case-sensitive search only applies to the local OLI. Peer OLIs will continue to use case-insensitive search. 2 Full-text search (keyword search) is case insensitive. You cannot change its case sensitivity. 3 Set this option to Yes to increase local query performance.

Option	Description
Full-text Search Options	
Use primary delimiters	<p>Default: Yes</p> <p>Controls whether primary delimiters are applied to an event to tokenize it for indexing.</p> <p>A primary delimiter tokenizes an event for indexing. For example, an event "john doe the first" is tokenized into "john" "doe" "the" "first" using the "space" primary delimiter.</p> <p>Supported primary delimiters: space, tab, newline, comma, semi-colon, (,), [,], {, }, ", , *, >, <, !</p>
Use secondary delimiters	<p>Default: No</p> <p>Controls whether secondary delimiters are applied to an event to further tokenize a token created by a primary delimiter thus enabling searches that can match a part of a primary token.</p> <p>For example, you can search for "hp.com" in http://www.hp.com.</p> <p>Supported secondary delimiters: =, ., :, /, \, @, -, ?, #, \$, &, _, %</p>
Regular Expression Search Options	
Case sensitive	<p>Default: No</p> <p>See "Case sensitive" on page 198.</p>
Unicode case sensitive	<p>Default: No</p> <p>Controls whether events in languages other than English should be compared in a case-sensitive way.</p> <p>Caution: HP strongly recommends that you do not change this option. You must restart the OLI for this change to take place.</p>
Check for canonical equality	<p>Default: No</p> <p>Controls whether events in languages other than English should be compared using locale-specific algorithms.</p> <p>Caution: HP strongly recommends that you do not change this option. You must restart the OLI for this change to take place.</p>

Option	Description
Search Display Options	
Populate rawEvent field for syslog events	<p>Default: No</p> <p>Controls whether raw events are displayed in a formatted column called rawEvent using the Raw Event field set. This option applies to syslog events only. If you want to view the raw events associated with CEF events, you do not need to configure this setting. Instead, configure the connector that is sending events to OLI to populate the rawEvent field with the raw event.</p> <p>Note: Even though the rawEvent column displays the raw event, this column is not added to the OLI database and is not indexed. Therefore, you can only run a keyword (full-text) or regular expression search on the event.</p>
Show Source and SourceType fields	<p>Default: No</p> <p>Controls whether the Source and SourceType fields are included in the Field Summary and query results.</p> <p>You must restart the OLI for this change to take place.</p> <p>Note: Setting this option to Yes can impact query performance.</p>
Field Summary Options	
Use Field Summary	<p>Default: Yes</p> <p>Controls the whether the Field Summary panel is included in the search results by default. Regardless of the default, you can change the setting on-the-fly by using the Fields Summary check box on the Search screen.</p>
Discover Fields	<p>Default: No</p> <p>Controls whether the Field Summary feature automatically detects non-CEF fields in raw events. Regardless of the default, you can change the setting on-the-fly by using the Discover Fields check box on the Search screen.</p> <p>This field is hidden if Use Field Summary is set to No.</p>

Viewing and Deleting Field Sets

You can view the field sets you have created and the predefined field sets on the Fieldsets tab (Configuration > Search > Fieldsets). You can also delete the field sets you created.



Note

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

To delete a custom field set:

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

- 2 Click **Search** from the left panel and then click Fieldsets.
- 3 In the Fieldsets tab, identify the field set you want to delete and click the delete (✖) icon.
- 4 Confirm the deletion.

Viewing Default Fields

The OLI schema comes with a set of predefined fields. Some of these fields are already indexed for improved search speed and efficiency. You can add custom fields to the OLI schema and index them for field-based search. A field-based search can only use fields in OLI's schema.



Note

The size of each field in the OLI schema is predetermined. If the string you are searching for is longer than the field-length, you should use a STARTSWITH rather than an = search, and include no more than the number of characters in the field size. For more information, see [“Field-based Search” on page 70](#).

The OLI Default Fields tab (**Configuration > Search > Default Fields**) displays the predefined fields included in the OLI schema. It includes the Display Name, Type, Length, and Field Name for each default field. To view information on existing custom fields, see [“Viewing Custom Fields” on page 201](#).

To view the default schema fields:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Search** from the left panel and then click **Default Fields**.

Search Indexes	Search Options	Fieldsets	Default Fields	Custom Fields	Running Tasks	Parsers
Display Name	Type	Length	Field Name	Indexed		
agentAddress	TEXT	16	agt			
agentHostName	TEXT	40	ahost			
agentNtdomain	TEXT	40	agentNtdomain			
agentSeverity	TEXT	-	agentSeverity	✓		
agentType	TEXT	16	at	✓		
agentZone	TEXT	200	agentZone			
agentZoneName	TEXT	50	agentZoneName			
agentZoneResource	TEXT	100	agentZoneResource			
agentZoneURI	TEXT	2048	agentZoneURI			
applicationProtocol	TEXT	40	app	✓		

- 3 The Default Fields tab displays the default fields. You can sort the fields by clicking the column headers.

Viewing Custom Fields

You can view the custom fields that have been added to the OLI schema under the Custom Fields tab (**Configuration > Search > Custom Fields**).

Display Name	Type	Length	Field Name	Actual Field Name	Creator	Created
Age	BIGINT	-	age	ad.age.l	admin	Dec 29, 2013 12:02:59 PM IST
Emp_Name	TEXT	255	name	ad.name	admin	Dec 29, 2013 12:03:21 PM IST
Salary	DOUBLE	-	salary	ad.salary.r	admin	Dec 29, 2013 12:03:36 PM IST
SocialSecurityNumber	TEXT	7	SSNs	ad.SSN.s.s	admin	Dec 29, 2013 12:04:05 PM IST

This page lists all custom schema fields that have been saved. You can view the alphabetical list of fields, but cannot edit or delete them.

For detailed information about custom schema fields, see [“Adding or Importing Schema Fields” on page 221](#).

Running Search Tasks

The Running Tasks page displays the search tasks that are running at the present time. The table shows the session ID, user who started the tasks, the date and time that the task started, the number of hits, the number of scanned events, the elapsed time, and the query.

Session ID	User	Start	Hits	Scanned	Elapsed	Query
104857871	admin	Jan 6, 2014 4:30:00 PM IST	180,115	205,786	01:01.571	receiver = "SmartMessage Receiver" csvexport name="Smart Message Receiver Events job" timestamp="true" download="false" hasSummary="true" allfields="true" baseevents="false" ceOnly="true" start="2014/01/05 16:30:00 IST" end="2014/01/06 16:30:00 IST" chartType="2" chartLimit="10" limit="1000000" ❌
104857872	admin	Jan 6, 2014 4:31:00 PM IST	73	205,786	00:01.560	_loggerStartTime="2014/01/05 16:30:00:000 IST" _loggerEndTime="2014/01/06 16:30:00:000 IST" (deviceEventClassid is null OR deviceEventClassid="actionengine:100" OR name is null OR name != "VMware Error job") (deviceVendor = "VMware") and deviceSeverity = "error" where message="HTTP Transaction failed on stream TCP(error:transport endpoint is not connected) with error N7Vmware15SystemException(Connection reset by peer)" alert name="VMware Error job" eventMatchL... ❌

Once a task finishes, the task's entry on the Running Tasks page is removed. (The task entry is removed upon page refresh, either when you refresh the browser page or when you navigate away from this page and come back to it.)

To view Running Tasks:


- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Search** in the left panel.
- 3 Click the **Running Tasks** tab in the right panel.
The Running tasks are displayed.

Ending Currently Running Tasks

You might need to end a currently running search task when it is taking too long to run, or appears to be stuck and slowing the overall OLI performance. When a search initiated as a result of any of the following operations is **in-progress**, the Running Tasks page (**Configuration > Search > Running Tasks**) displays the currently running process.

- A manual search on local or peer OLI (**Analyze > Search**)
- A scheduled search (**Configuration > Saved Search > Scheduled Searches/Alerts**)

- A saved search alert (**Configuration > Saved Search > Scheduled Searches/Alerts**)
- A search export, with the “Rerun query” option checked (**Analyze > Search > Export Results**)

To end a process, click the  icon. You must have admin user privileges to end a running search process.

View and Add Parsers for Specific Log Types

The Parsers tab on this screen is the same as the one documented in the Event Input section. See [“Parsers” on page 159](#) for information on parsers and how to add them.

Peer OLIs

An OLI can establish peer relationships with one or more OLIs to enable distributed event searches.

When two OLIs peer with each other, one OLI initiates the relationship. The initiator OLI sends the credentials to authenticate itself to the other OLI, called the remote OLI from here on. If the authentication succeeds, a peer relationship is established between the two OLIs. (The remote OLI must have the credentials for the initiator OLI configured on it for authentication to succeed.)

Adding a peer on an OLI is a bi-directional process. That is, when OLI A adds peer access for OLI B, OLI B automatically adds peer access for OLI A. Similarly, if you delete the peer access for B on A, the peer access for A is automatically deleted on B.



On an OLI 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 OLI, authorization ID and code is the only way to authenticate a peer. You cannot use a user name and password.

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

Peer OLIs can authenticate using any of these methods:

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

Guidelines

You should be aware of these guidelines when peering OLIs.

- You can peer an OLI to one or more remote OLIs.
- The time and date on the system on which the OLI is installed must be set correctly with respect to its timezone to peer with other OLIs. HP recommends that you configure the OLI system to synchronize its time with an NTP server regularly.
- If the remote OLI is configured for SSL Client authentication (CAC), you must configure an authorization ID and code on the initiator OLI.

There are no special authentication requirements for FIPS-enabled OLIs. Such OLIs can use any of the allowed authentication methods.

- Peer OLIs cannot be edited, however you can delete and re-add a peer.
- If you are running distributed searches (searches across peers), follow these additional guidelines:
 - ◆ A user must belong to the OLI Search User Group with “Search for events on remote peers” privilege set to Yes and the OLI Rights Group with “View registered peers” privilege set to Yes. See [“Searching Peer OLIs \(Distributed Search\)” on page 97](#).
 - ◆ Users performing search operations on peers have the same privileges on the peer that they have on the OLI they are logged in.
 - ◆ For example, User A is restricted by a search group filter to only search for events in which deviceVendor is set to “Cisco”. When User A performs a search operation across OLI A'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 OLI, 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 OLIs 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 OLIs.

OLI A	OLI B
1	1
1	2
2	2
	3
	4

1 Select the OLI that will initiate the establishment of the peering relationship.

In this example, OLI A will initiate the relationship.

2 If OLI B is configured to use user name and password authentication, go to [Step 3](#).

If OLI B is configured to use SSL Client Authentication (CAC), go to [Step 4](#).

OLI A	OLI B
<p>5 Add OLI B's information, as described in "To add a peer OLI:" on page 205:</p> <p>If OLI B uses user name and password, use the user name and password you configured in Step 3.</p> <p>If OLI B uses SSL Client Authentication, use the Authorization ID and Code you generated in Step 4.</p>	<p>3 Set up a user name and password that OLI A will use to authenticate itself when peering with this OLI, as described in "Users/Groups" on page 251.</p> <p>4 Generate an Authorization ID and Code that OLI A will use for authenticating to OLI B, as described in "To generate Authorization ID and Code for configuring a peer relationship:" on page 207.</p>

Peer OLIs
Peer Authorizations

Add Peer OLI

Peer Host Name

Peer Port

Peer Login Credentials
 Peer Authorization Credentials

Peer User Name

Peer Password

In most cases, the fields below will be pre-populated for you, and you do not need to change them. In the event that you need to change these fields, please consult the Operations Log Intelligence Administrator's Guide for specific instructions.

External IP Address

Local Port

To add a peer OLI:

- 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 OLI's hostname or IP address.
Peer Port	Port you configured when installing OLI. See "Guidelines" on page 204.
Peer Login Credentials	Select Peer Login Credentials for password-based authentication with the remote OLI.
Peer Authorization Credentials	Select Peer Authorization Credentials for SSL client authentication with the remote OLI. (See "SSL Client Authentication" on page 244.)

If you selected Peer Login Credentials...

Peer User Name	The user name to use when connecting to the remote OLI.
Peer Password	The password for the user on the remote OLI.

If you selected Peer Authorization Credentials...

Peer Authorization ID	Enter the authorization ID of the other OLI to which this OLI is initiating a peering relationship. (See "To generate Authorization ID and Code for configuring a peer relationship:" on page 207 for more information.)
Peer Authorization Code	Enter the authorization code of the other OLI to which this OLI is initiating a peering relationship. (See "To generate Authorization ID and Code for configuring a peer relationship:" on page 207 for more information.)

These fields need to be updated in rare circumstances. For more information, read the description of each field in this table.

External IP Address	In most cases, the value in this field matches the IP address you use to connect to this OLI from your browser, and you do not need to do anything. However, if the IP address does not match (for example, when the OLI is behind a VPN concentrator), change the value to match the IP address with which you connect to this OLI.
Local Port	Make sure the value of this field is set to 443.

- 4 Click **Save** to add the new OLI, or **Cancel** to quit.

To delete a peer OLI:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Peer Loggers** from the left panel.

- 3 Locate the Peer that you want to delete and click the delete icon (✖) on that row.
- 4 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Peer.

To view peers of an OLI:

A list of remote OLIs that an OLI peers with is displayed on the Peer OLIs page (**Configuration > Peer Loggers**).

Authorizing Peers

Use the following procedure to generate the authorization ID and code on the OLI to which you are establishing a peer relationship. (OLI B in the example described earlier in this section.) This ID and Code is then configured on the OLI that initiates the peer relationship. (OLI A in the earlier example.)

To generate Authorization ID and Code for configuring a peer relationship:

- 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 OLI 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 OLI that is configured to use SSL client authentication.

Configuration Backup and Restore

By default, OLI does not back up any content. However, you can configure it to backup all non-event data to a remote system.

You can back up this content on ad-hoc basis or schedule it to run periodically. The content is saved to the backup location in a single .tar.gz format file.



Ensure that both Configuration Backups (for configuration settings) and Event Archives (for data) run on a regular basis and are stored in a remote location. In the event of catastrophic failure, you will need to restore the most recent Configuration Backup and Event Archive. For information on Event Archives, see [“Event Archives” on page 131](#).

The following lists the information included in the backup when you back up all non-event data:

- System Information
- Logs
- Global Settings
- User and Group Information

- All Configuration Settings
- Existing Filters and Saved Searches
- OLI Monitor settings

You can use the backed up content to:

- Restore an OLI that is not functioning as expected or that has been reset to its factory defaults
- Copy content from one OLI to another



When you restore content to an OLI, the existing content on it is deleted.

Running a Configuration Backup (Ad-hoc or Scheduled)

Configuration Backup

Edit Configuration Backup

Protocol

Port

IP/Host

User

Password

Remote directory

Schedule One time only

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 OLI 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)

Parameter	Description
Password	<p>Password for the user</p> <p>Note: The password cannot contain these characters: % = ; " ' <></p>
Remote Directory	The folder on the remote system in which to save the configuration backup files
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 236.</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 OLI:

- When you restore content to an OLI, the existing content on it is deleted. OLI restores the specific environment settings that were current at the time the backup was taken. Any configuration settings that were updated between the time of the backup and the time of the restore are lost.
- The operating system that OLI is running must be the same as the one used to create the backup file.

To restore from a configuration backup:

- 1 Click the **Configuration > 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 208.

System Maintenance

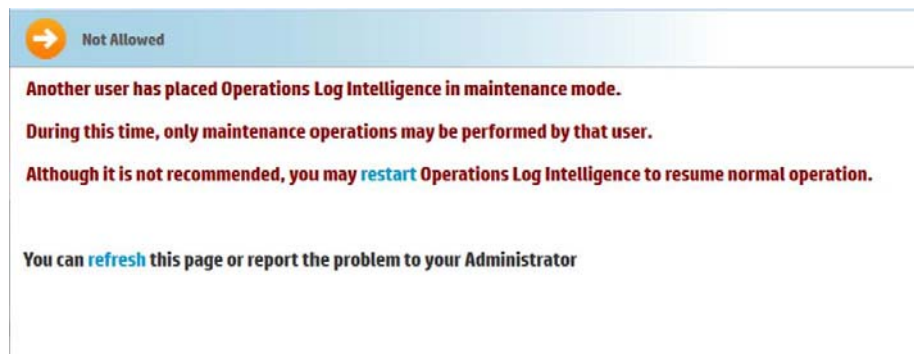
Certain operations on OLI, such as database defragmentation, extending the storage volume size, adding storage groups, and adding additional schema fields, require that OLI be in a maintenance state—a state in which operations related to data on the OLI are not running. Maintenance mode enables you to place the OLI in such a state. When an OLI is in maintenance mode:

- Events are not processed
- Search cannot run
- Scheduled jobs do not run

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



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



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

➔
Not Allowed

Another user has placed Operations Log Intelligence in maintenance mode.

During this time, only maintenance operations may be performed by that user.

You can [refresh](#) this page or report the problem to your Administrator

Entering Maintenance Mode

You cannot place an OLI in maintenance mode directly. An OLI 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 OLI in maintenance mode, as illustrated in “[Database Defragmentation](#)” on page 211.

Exiting Maintenance Mode

To exit maintenance mode, restart the OLI.

Checking Status of a Maintenance Operation

You can check the status of a maintenance operation on the Maintenance Results page. To access the Maintenance Results page (as shown in the example below), click **Configuration > System Maintenance > Maintenance Results**.

Status	Operation	Start	End	Message	Creator
Success	Add Storage Groups	Jan 5, 2014 4:51:28 PM IST		Added Storage Group [Dev]	admin
Failure	Add Storage Groups	Jan 5, 2014 4:51:18 PM IST		Failed attempting to add Storage Group [] with error [Missing Name]	admin
Failure	Add Storage Groups	Jan 5, 2014 4:51:10 PM IST		Failed attempting to add Storage Group [Dev] with error [Maximum Size (GB) value is not valid (3 - 4 GB)]	admin
Success	Database Defragmentation	Jan 5, 2014 2:05:16 PM IST	Jan 5, 2014 2:05:16 PM IST	Defragmentation complete	admin
Success	Database Defragmentation	Jan 5, 2014 12:00:43 PM IST	Jan 5, 2014 12:00:48 PM IST	Defragmentation complete	admin

Database Defragmentation

OLI’s database can get fragmented over time. Frequent retention tasks can exacerbate this issue. The following symptoms are observed on an OLI when the database should be fragmented:

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

You can defragment an OLI that exhibits the above listed symptoms. Make sure that you have read the following guidelines before starting the defragmentation process.

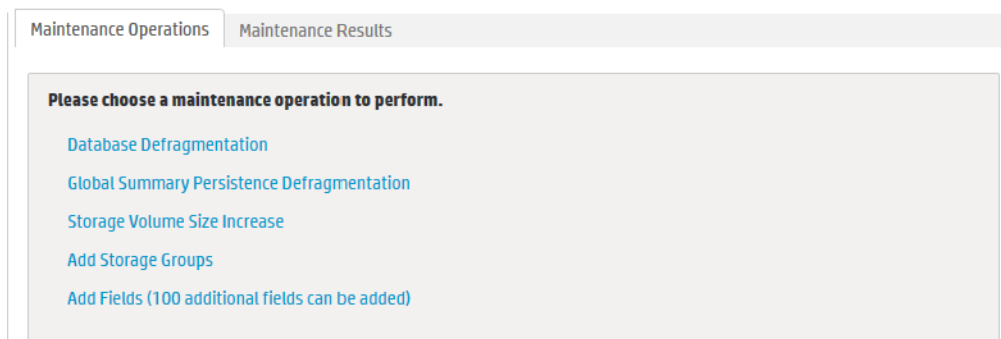
Guidelines for Database Defragmentation

- Ascertain that the OLI symptoms are not due to issues related to network infrastructure such as network latency or unexpected load on the OLI.
- The OLI system needs to be placed in maintenance mode before defragmentation can begin. As a result, most processes on the OLI 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 210](#).
- A minimum amount of free disk space is required on your system to run database defragmentation. The utility automatically checks for the required free space and displays a message if sufficient disk space is not found.
- Although you can defragment as needed, if you are using this utility too often (such as on a system that was defragmented over the last few days), contact customer support for guidance.
- If the defragmentation process fails at any point, the OLI returns to the same state that it was in before you started defragmentation. Restart the OLI process as described in [“Process Status” on page 237](#).
- You can perform this process only if you have the “Enable Maintenance Mode” privilege set to Yes (System Admin > User/Groups > Manage Groups > System Admin Group).

Defragmenting an OLI

To defragment an OLI:

- 1 Click **Configuration** > **System Maintenance**. The Maintenance Operations panel displays the available options.



- 2 Click **Database Defragmentation**.
- 3 Click **Enter Maintenance** so that the OLI can enter maintenance mode. For more information about maintenance mode, see [“System Maintenance” on page 210](#).

→ Database Defragmentation

Before performing any maintenance operation, Operations Log Intelligence must first enter maintenance mode to safeguard event data. During this time, Operations Log Intelligence won't receive or forward events. Once in maintenance mode, Operations Log Intelligence will need to be restarted to resume normal operations.

This will take about two minutes to complete.


Please check the Operations Log Intelligence release notes for additional information.

Press **Enter Maintenance** to enter maintenance mode now.

Enter Maintenance

4 A minimum amount of free storage is required for the database defragmentation process to proceed. Therefore, OLI performs a check to determine free storage when entering maintenance mode.

- ◆ If the required storage is not found, follow the instructions found in ["Freeing storage space for defragmentation" on page 214](#).
- ◆ If the required amount of free storage is found and OLI successfully enters maintenance mode, the following screen is displayed. Click **Begin Defragmentation** to start the defragmentation process.



Note

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

→ Database Defragmentation

Operations Log Intelligence is ready to perform the database defragmentation. There is sufficient free storage to perform this operation. (Required free storage: 86.88 MB, available free storage: 11.49 GB)

Please check the Operations Log Intelligence Release Notes for additional information.

This should take approximately 17 seconds.

Press **Begin Defragmentation** to begin.

Begin Defragmentation

Operations Log Intelligence is in Maintenance Mode. You may [restart](#) Operations Log Intelligence at any time to resume normal operation.

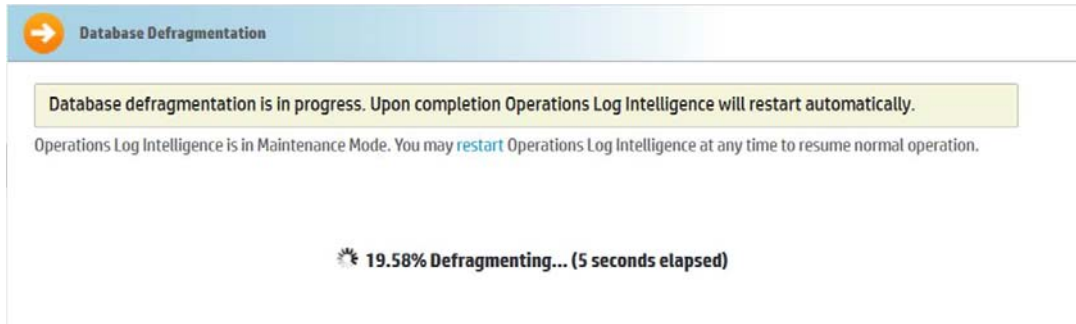
Figure 5-16 Begin Database Defragmentation

5 The defragmentation process starts. A progress indicator shows the status of defragmentation, as shown in the example below. HP recommends that you do not attempt any operation on the OLI until defragmentation has completed.

Once defragmentation is complete, the OLI restarts automatically. This exits maintenance mode.

Confidential

HP Operations Log Intelligence 1.0.0 Administrator's Guide **213**



Freeing storage space for defragmentation

If the required storage is not found, OLI prompts you to free sufficient space, as shown in the following example:

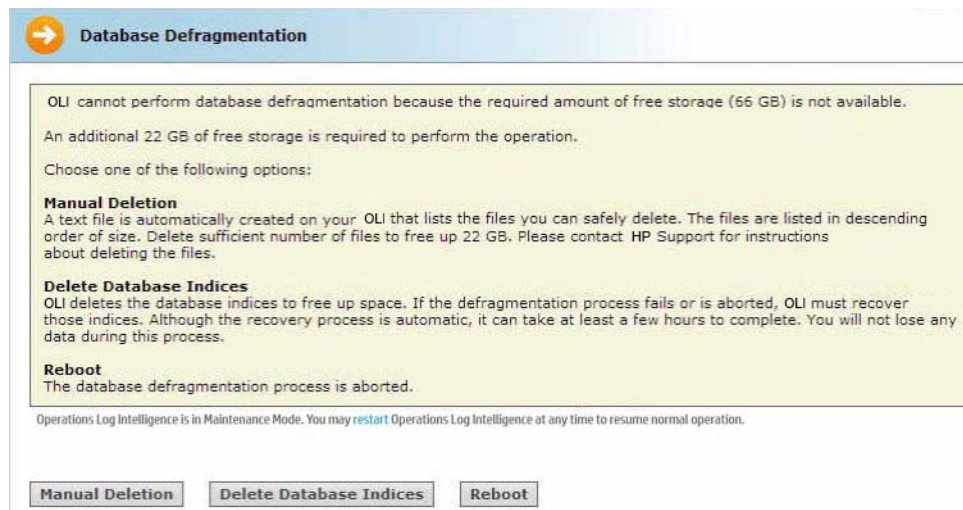


Figure 5-17 Required storage for Database Defragmentation is not available

You can choose from one of the following options:

◆ Manual Deletion

A text file is automatically created on your OLI that lists the files you can safely delete. This file is located in `<install_dir>/current/arcsight/logger/user/logger/defragmentation/filelist.txt`.

The files are listed in descending order of size in the text file. You can delete sufficient number of files to free up storage. However, **do not** delete the files before contacting customer support for instructions and guidance.

Follow these steps to proceed:

- i Leave the message screen without taking any action.
- ii Contact customer support for instructions on deleting files listed in the text file.

- iii After deleting sufficient number of files, resume the Database Defragmentation process from the message screen in [Step i on page 214](#). 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 5-16 on page 213](#) is displayed. Click **Begin Defragmentation** to proceed further.

If sufficient storage is still not found, the screen in [Figure 5-17 on page 214](#) is displayed. Choose from the listed options to create additional space. See [“You can choose from one of the following options:” on page 214](#) for more information.



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

◆ Delete Database Indices

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

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

The screenshot shows a window titled "Database Defragmentation" with a yellow background. The text inside reads:

OLI cannot perform database defragmentation because the required amount of free storage (66 GB) is not available. An additional 22 GB of free storage is required to perform the operation.

Choose one of the following options:

Manual Deletion
A text file is automatically created on your OLI that lists the files you can safely delete. The files are listed in descending order of size. Delete sufficient number of files to free up 22 GB. Please contact HP Support for instructions about deleting the files.

Reboot
The database defragmentation process is aborted.

Note: All available database indices have already been deleted.

Operations Log Intelligence is in Maintenance Mode. You may **restart** Operations Log Intelligence at any time to resume normal operation.

At the bottom, there are two buttons: "Manual Deletion" and "Reboot".

Follow these steps to proceed:

- i Click **Manual Deletion**.

A text file is created on your OLI that lists the files you can safely delete. The files are listed in descending order of size in a text file.

- ii Click **Reboot**.

OLI exits the maintenance mode.

- iii Contact customer support for instructions on manually deleting the files.

You can delete sufficient number of files to free up storage.

- iv After deleting the files, restart the defragmentation process from [Step 1 on page 212](#).



If the defragmentation process fails or is aborted at any time, OLI 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 OLI returns to the state it was in before you started the defragmentation utility.

Global Summary Persistence Defragmentation

There is a known issue with the new Global Summary Persistence functionality in version 1.0.0 of OLI. This feature is designed to persist the statistics reported in the global summary section of OLI through a reboot. In some environments, disk space may be affected due to this feature.

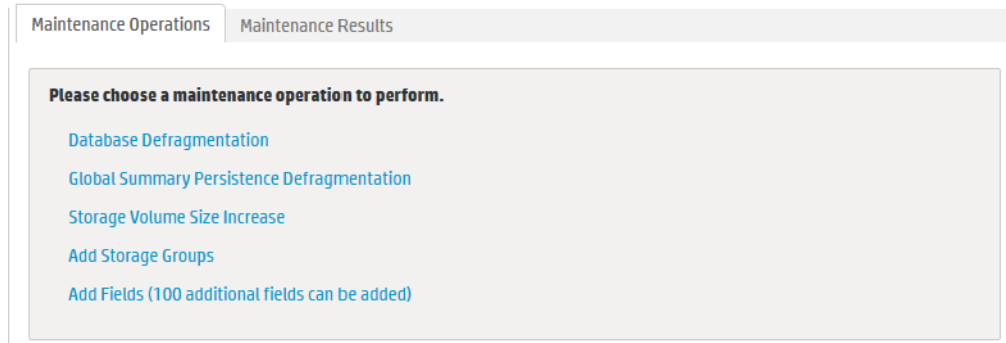
This release turns off the Global Summary Persistence functionality. As soon as possible after installing OLI version 1.0.0, enter System Maintenance mode and defragment the Global Summary table. Make sure that you have read the following guidelines before starting the defragmentation process.

Guidelines for Global Summary Persistence Defragmentation

- The OLI system needs to be placed in maintenance mode before Global Summary Persistence defragmentation can begin. As a result, most processes on the OLI 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 210](#).
- A minimum amount of free disk space is required on your system to run Global Summary Persistence defragmentation. The utility automatically checks for the required free space and displays a message if sufficient disk space is not found.
- If the defragmentation process fails at any point, the OLI returns to the same state that it was in before you started defragmentation. You can safely restart the OLI process as described in [“Process Status” on page 237](#) and try again.
- You can perform this process only if you have the “Enable Maintenance Mode” privilege set to Yes (System Admin > User/Groups > Manage Groups > System Admin Group).

To defragment for the Global Summary Persistence issue:

- 1 Click **Configuration > System Maintenance**. The Maintenance Operations panel displays the available options.



- 2 Click **Global Summary Persistence Defragmentation**.
- 3 Click **Enter Maintenance** so that OLI can enter maintenance mode. For more information about maintenance mode, see [“System Maintenance” on page 210](#). The Global Summary Persistence Panel displays information about the operation.

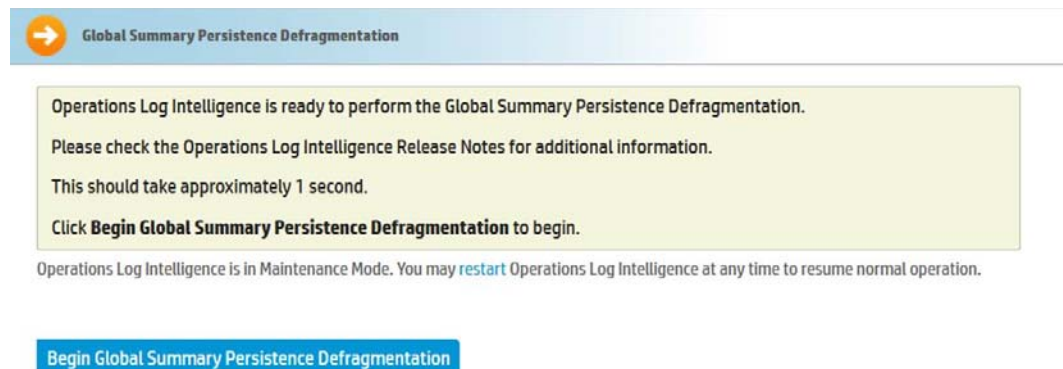


Figure 5-18 Begin Global Summary Persistence Defragmentation

- 4 Click **Begin Global Summary Persistence Defragmentation** to start the defragmentation process.
- 5 The defragmentation process starts. A progress indicator shows the status of defragmentation. HP recommends that you do not attempt any operation on the OLI until defragmentation has completed.

Once defragmentation is complete, the OLI reboots or restarts. This automatically exits maintenance mode.



Only the OLI service and its related processes are restarted.

Note

Storage Volume Size Increase

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

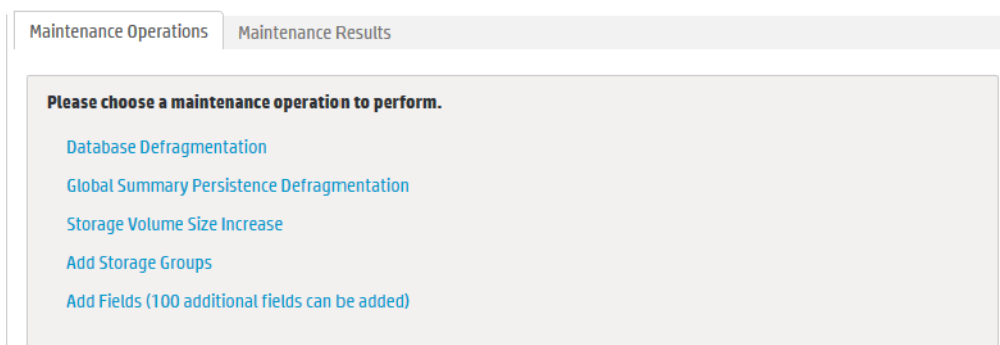
guides you about current and the maximum value to which you can increase the size.



For the “Storage Volume Size Increase” operation to show as an option under the System Maintenance operations (Configuration > System Maintenance), you need to belong to the System Admin group (with “Enable Maintenance Mode” privilege enabled) and the Operations Log Intelligence Rights group.

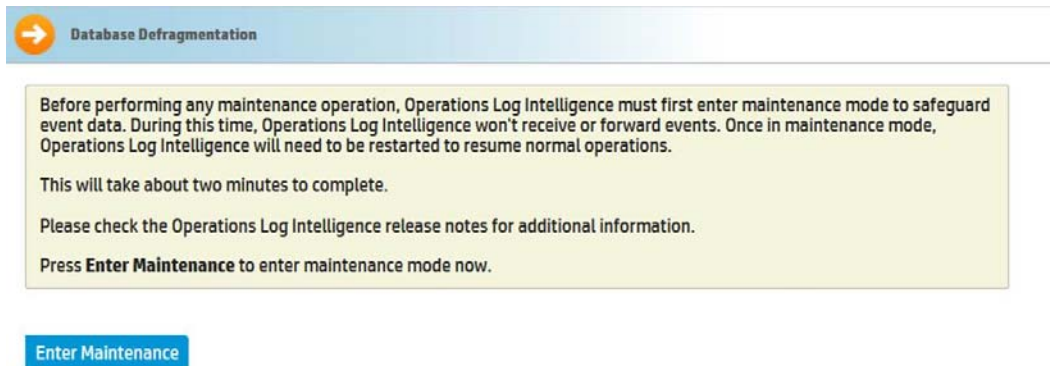
To increase the size of a storage volume:

- 1 Click **Configuration > System Maintenance**. The Maintenance Operations panel displays the available options.



- 2 Click **Storage Volume Size Increase**.
- 3 Click **Enter Maintenance** so that the OLI can enter maintenance mode.

For more information about maintenance mode, see [“System Maintenance” on page 210](#).



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

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



Note

The following Storage Volume Size Increase screens instruct you to click **restart** to resume normal operation when OLI is in maintenance mode. When you click restart, only the OLI service and its related processes are restarted.

→ **Storage Volume Size Increase**

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

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

New size (GB)

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

→ **Storage Volume Size Increase**

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

Operations Log Intelligence is in Maintenance Mode. You may **restart** Operations Log Intelligence at any time to resume normal operation.

Adding Storage Groups

In addition to the two storage groups that exist on your OLI by default, you can add up to four additional storage groups. You can add storage groups at any time if the following conditions are met:

- The maximum allowed six storage groups do not exist on your OLI already.
- The storage volume contains spare storage space that can be allocated to the storage groups you will add.



Tip

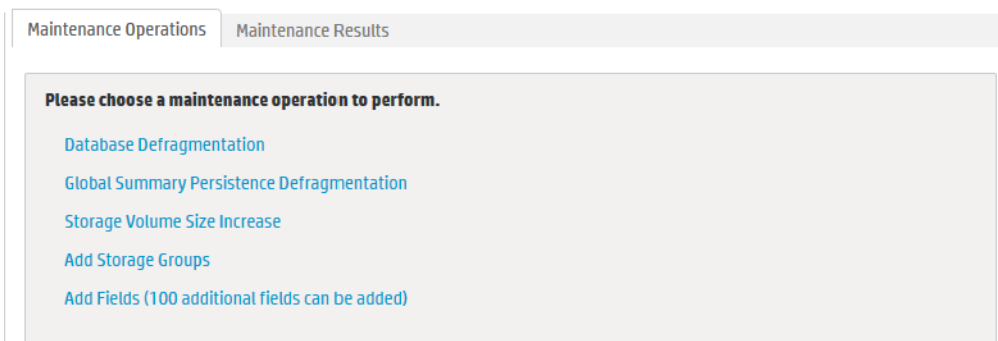
If you do not have sufficient space in the storage volume to add another storage group and the existing groups have free space, consider reducing the size of existing storage groups to make space available for the storage groups you want to add. Alternatively, increase the size of your existing storage volume, as described in [“Storage Volume Size Increase” on page 217](#).

The OLI must be in maintenance mode when adding storage groups. When you add a storage group, OLI automatically checks to ensure that the storage group size you specified is greater than the minimum size required (5 GB) and less than the amount of space available in the storage volume.

Once you have added storage groups and rebooted your OLI to exit the maintenance mode, remember to configure the Archive Storage Settings for the groups you just added so that event archives are created for them.

To add a storage group:

- 1 Click **Configuration > System Maintenance**. The Maintenance Operations panel displays the available options.



- 2 Click **Add Storage Groups**.

A maximum of six storage groups can exist on OLI. Therefore, you can add up to four storage groups in addition to the two that exist by default on OLI.

If the maximum number of allowed storage groups **do not** exist on OLI, a screen prompts you to enter maintenance mode, as described in the next step.

If all six storage groups exist on OLI or sufficient space does not exist in the storage volume to add additional group, a message is displayed on your screen and the OLI cannot enter maintenance mode.

- 3 Click **Enter Maintenance** so that the OLI can enter maintenance mode.

For more information about maintenance mode, see [“System Maintenance” on page 210](#).

- 4 Once OLI enters maintenance mode, the following Add Storage Groups page is displayed.

Name	Maximum Age (Days)	Maximum Size (GB)	Creator	Last Editor
Default Storage Group	180	7	admin	admin
Internal Event Storage Group	365	3	System	System

This screen also lists information about the existing storage groups and the amount of space remaining in the storage volume.

- 5 Enter the following information.

Parameter	Description
Name	Choose a name for the storage group
Maximum Age (Days)	Specify the number of days to retain events. Events older than this number of days are deleted.
Maximum Size (GB)	Enter a maximum event data size, in GB.

- 6 Click **Add**.

The storage group is added to your OLI. If your OLI has not reached the maximum allowed six storage groups, you can click **Add** to add more storage groups. However, if the maximum number has been reached, the Add button is not displayed. If you do not want to add more storage group, go to the next step.

- 7 Restart your OLI for changes to take effect.

Adding or Importing Schema Fields

The OLI schema contains a predefined set of fields. A field-based query can contain only these fields. Additionally, you can index only these fields for faster search operations. For instructions on how to view the default OLI Schema fields, see [“Viewing Default Fields” on page 201](#).

You can add additional fields to the OLI schema. That is, you can insert fields in your OLI schema that are relevant to the events you collect on your OLI, thus enabling you to search using these fields. Additionally, you can index the fields you add so that the search queries that use these fields run faster. For example, a financial institution might want to add credit card numbers or social security numbers to the schema.

You can add up to 100 custom schema fields on OLI. You can also import custom fields from a peer OLI. However, the total number of added and imported fields cannot exceed the maximum allowed 100 fields.

You can index up to 123 fields on OLI. Therefore, the number of custom schema fields you can index will depend on the number of default fields you currently have indexed on your OLI.

The events that contain custom fields must be in CEF format (key-value pairs) for OLI to process them. Therefore, you will need to either use a SmartConnector that generates additional data or define an ArcSight FlexConnector to collect and parse events containing custom fields from the event source, convert them into CEF format, and forward them to the OLI.

OLI can only process events from FlexConnectors written using connector build 5.0.0.5560 or later. For details about designing FlexConnectors, see the ArcSight FlexConnector Developer's Guide.



Note

OLI cannot process the additional fields data received in CEF version 0 from a FlexConnector, and assumes a NULL value for such fields when they are present in a CEF version 0 event. As a result, you cannot search on these fields or index them. However, these fields are displayed in the UI display when you select "*" in the fieldset because the interface displays information contained in the raw event. Therefore, if OLI receives "ad.callnumber=5678", the OLI UI will display a column, ad.callnumber, with value 5678. However, a search on "5678" will not return this event in the search results.

You need to be in maintenance mode to add or import custom schema fields. The process of adding or importing schema fields involves an add or import operation followed by a save operation. The add or import operation adds the specified fields but does not write them to the OLI schema. You can edit or delete the added or imported fields at this point. Once you save these fields, the fields are written to the schema. From this point on, these fields cannot be edited or deleted. Therefore, carefully review the fields you are adding to the schema before saving them.



Note

For the "Add Fields" operation to show as an option under the System Maintenance operations (Configuration > System Maintenance), you need to belong to the System Admin group (with "Enable Maintenance Mode" privilege enabled) and the OLI Rights group.

You need to specify the following information to add a custom schema field:

- Display name
A meaningful name for the field. This name is displayed as the column header name for the field and is the one you specify in a search query. For example, SocialSecurityNumber.
- Type

The type of data this field will contain. The available options are Double, BigInt, DateTime, Text.

The following table describes each data type.

Type	Description
Double	Use to store decimal numbers or fractions. Numbers from -1.79769313486231570E+308 through -4.94065645841246544E-324 for negative values and 4.94065645841246544E-324 through 1.79769313486231570E+308 for positive values.
BigInt	Use to store whole numbers. Numbers from -2^{63} (-9,223,372,036,854,775,808) through $2^{63}-1$ (9,223,372,036,854,775,807)
DateTime	Use to store both dates and time or only dates.
Text	Use to store any characters. You can store a maximum of 255 characters per field.

- Length

This field is only relevant when the Type specified is Text. This field specifies the maximum number of characters allowed in the value of the field when the data type is Text.

- Field name

The field name that you want to add to the OLI schema. Typically, this is an abbreviated version of the Display name. For example, SSN.

Importing Schema Fields from Peers

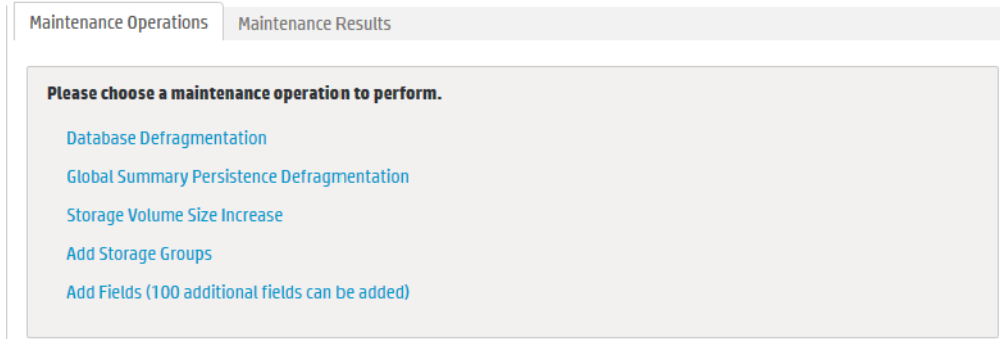
If your OLI is a peer of another OLI, you can import the custom fields added to the peer's schema. You specify the peer from which you want to import fields in the user interface screen. Fields can be imported if the following conditions are met:

- A field of the same Display name and Field name does not exist on the OLI to which you are importing schema fields. If conflicting fields exist, they are still imported but are flagged in the user interface screen. You cannot save the imported fields to schema until you resolve the conflicts.
- A maximum of 100 custom fields has not been reached on the importing OLI. If there are more fields than can be imported, only the first N until the allowed maximum is reached will be imported.

The custom schema fields contained in a search query must exist on all peers on which the query is run. Otherwise, the query will not run and return an error.

To add or import custom schema fields:

- 1 Click **Configuration > System Maintenance**. The Maintenance Operations panel displays the available options.



- 2 Click **Add Fields (100 additional fields can be added)**.

You can add a maximum of 100 custom fields to OLI schema. The number in the “Add Fields” link reflects the number of custom fields you can add. This number decreases as you add fields to OLI schema.

- 3 Click **Enter Maintenance** so that the OLI can enter maintenance mode.

For more information about maintenance mode, see [“System Maintenance” on page 210](#).

- 4 Once OLI enters maintenance mode, the following Add Fields page is displayed.

You can add fields manually or import them from a peer OLI.

To add fields manually see

To manually add fields:

- 1 Click “Add a New Field”, if it is not selected.

- 2 Enter a meaningful name in the Display Name field.

This name is the one you specify in a search query and is displayed as the column header name for the field in search results. For example, SocialSecurityNumber. This name is not added to the OLI schema. Follow these guidelines when specifying a display name:

- ◆ The name can contain up to 100 characters.
- ◆ The name can contain alphanumeric characters, hyphens (“-”), and underscores (“_”). However, a hyphen (“-”) or an underscore (“_”) cannot

be the first character in the name. Additionally, the name cannot begin with “arc_”.

- ◆ The name must be unique; that is, another field (custom or OLI schema) of the same display name must not already exist on the OLI.
 - ◆ Only ASCII characters are allowed. That is, no native Chinese or Japanese characters are accepted in this field.
- 3 Select a data type for the field from the Type drop-down menu.
 - 4 The available options are Double, BigInt, DateTime, Text. See [“Type” on page 222](#) for more information.
 - 5 In the Length field, enter the maximum number of characters allowed in the value of the field *when the data type is Text*. This field is only available when the Type specified is Text. You can specify from 1 to 255 characters in this field.
 - 6 Enter a name in the Field name field.

This is the name that will be added to the OLI schema. Typically, this is an abbreviated version of the Display name. For example, SSN. Follow these guidelines when specifying a Field name:

- ◆ This is a required field.
- ◆ The name can contain up to 40 characters and can contain alphanumeric, hyphen (“-”), and underscore (“_”) characters. Underscore (“_”) is used as an escape character for the actual field name. Therefore, the underscore (“_”) you specify in the field name is converted to a double underscore (“__”) in the actual field name.
- ◆ The name must be unique; that is, a custom field of the same Field name must not already exist on the OLI.
- ◆ Only ASCII characters are allowed. That is, no native Chinese or Japanese characters are accepted in this field.

Once you enter a name in this field, a prefix and a suffix is automatically added to it, and the resulting name is displayed in the Actual Field Name field, as shown in the following figure. This field displays the way the field name you entered earlier will be stored on OLI. The prefix, “ad.” signifies “additional data” and the suffix signifies the data type of the field. The Actual Field Name field is a non-editable field and is displayed on the user interface only for your reference.

- 7 Click **OK**.

The field you added is displayed in the upper section of the Add Fields form, as shown in the following figure. This field is not saved yet (in “Ready to Save” state) and you can edit or delete it. Once you click Save, the field is added to the schema and cannot be changed or deleted.

5 Configuration

Add Fields

Operations Log Intelligence is ready for adding new fields. You can add up to **96** additional fields. The fields in "Ready to save" status are not in OLI schema yet. Click Save to write these fields to the schema.

Operations Log Intelligence is in Maintenance Mode. You may **restart** Operations Log Intelligence at any time to resume normal operation.

Save

Display Name	Type	Length	Field Name	Actual Field Name	Creator	Created	Status	
Age	BIGINT	--	age	ad.age.i	admin	Jan 5, 2014 5:36:49 PM IST	Ready for save	
Emp_Name	TEXT	255	name	ad.name	admin	Jan 5, 2014 5:37:19 PM IST	Ready for save	
Salary	DOUBLE	--	salary	ad.salary.r	admin	Jan 5, 2014 5:37:34 PM IST	Ready for save	
SocialSecurityNumber	TEXT	7	SSNs	ad.SSNs.s	admin	Jan 5, 2014 5:38:36 PM IST	Ready for save	

Add a New Field
 Import Fields From Peers

Display Name:
Type:
Field Name:

OK

- Follow [Step 1](#) through [Step 7](#) to add additional fields.
- Review the added fields and make any edits () or deletions () , if necessary.



Caution

- The next step commits the added fields to OLI's schema. This process is irreversible; that is, once the fields are written to OLI's schema, they cannot be edited or deleted.
- If you exit this process without saving, the fields you were adding are not remembered and your changes are lost.

- Click **Save** to commit the added fields and write them to your OLI's schema.

To import fields from a peer:

- Click "Import Fields From Peers".
- Select the peer from which you want to import the fields from the Peer Host Name drop-down list.

Add a New Field
 Import Fields From Peers



Peer Host Name:

OK

- Click **OK** in the bottom right corner of the screen.

If there are no conflicting fields, all fields from the peer are imported successfully.

If there are conflicts, the conflicting fields are displayed ahead of the ones that were imported successfully. The Status column describes the reason for the conflict. You must fix the listed issues before you can save these fields to the

schema. Use the edit () or delete () icon to make changes or delete the added fields.




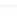


Add Fields

Operations Log Intelligence is ready for adding new fields. You can add up to **95** additional fields. The fields in "Ready to save" status are not in OLI schema yet. Click Save to write these fields to the schema.

Another field of the same display name, "Age", exists. Enter another display name.

Operations Log Intelligence is in Maintenance Mode. You may restart Operations Log Intelligence at any time to resume normal operation.

Save

Display Name	Type	Length	Field Name	Actual Field Name	Creator	Created	Status
string1	TEXT	20	string1	ad.string1	admin	Jan 5, 2014 5:43:18 PM IST	Ready for save  
Age	BIGINT	-	age	ad.age.i	admin	Jan 5, 2014 5:36:49 PM IST	 Saved
Emp_Name	TEXT	255	name	ad.name	admin	Jan 5, 2014 5:37:19 PM IST	 Saved
Salary	DOUBLE	-	salary	ad.salary.r	admin	Jan 5, 2014 5:37:34 PM IST	 Saved
SocialSecurityNumber	TEXT	7	SSNs	ad.SSN.s.s	admin	Jan 5, 2014 5:38:36 PM IST	 Saved

If there are more fields than can be imported, only the first N until the allowed maximum (100) is reached will be imported.



Caution

The imported fields are not committed to OLI's schema yet. The next step commits them. This process is irreversible; that is, once the fields are written to OLI's schema, they cannot be edited or deleted.

If you exit this process without saving, the fields you were adding are not remembered and your changes are lost.

- 4 Click **Save** to commit the added fields and write them to your OLI's schema.

To view existing custom schema fields:

See ["Viewing Custom Fields" on page 201](#).

License Information

The License Information page (**Configuration > License Information**) provides information about the currently applied license, as shown in the following figure.

License Information
Data Volume Restrictions

License

Customer: HP OLI Trial Customer
 Expiration date: 2015/01/08
 Activation date: 2014/01/08
 Creation date: 2014/01/08
 Operations Log Intelligence features: Enabled
 Connector appliance features: Disabled

Operations Log Intelligence features

Alerting: Enabled
 Local storage: Enabled
 SAN storage: Disabled
 Peering Enabled: Enabled

Operations Log Intelligence limits

Device Event Sources: 25
 EPS incoming: 1,000,000
 Daily data: 25GB
 Maximum capacity: 4300GB
 Maximum violations: 5
 Violation days: 30

To upload a new license, open **System Admin** from the top-level menu bar, and then click **License & Update** in the **System** section. For details, see “[License & Update](#)” on page 236 for OLIs.

Data Volume Restrictions

The Data Volume Restrictions page lists the data stored on your OLI on day-by-day basis in the last 30 days. It also indicates the days on which data limits were exceeded, as shown in the following figure.

The screenshot shows the HP Operations Log Intelligence interface. The top navigation bar includes Summary, Analyze, Dashboards, Configuration, and System Admin. The Configuration section is active, and the Data Volume Restrictions page is displayed. The page features a table with the following data:

Date	Data Stored (MB)	Limit Exceeded
11/22/13	0	false
11/23/13	0	false
11/24/13	0	false
11/25/13	0	false
11/26/13	0	false
11/27/13	0	false
11/28/13	0	false
11/29/13	0	false
11/30/13	0	false
12/1/13	0	false
12/2/13	0	false
12/3/13	0	false
12/4/13	0	false
12/5/13	0	false
12/6/13	0	false
12/7/13	0	false
12/8/13	33844	true

Retrieve Logs

OLI records some audit and debug information, including details of any issues that occur. These system logs (not be confused with the event logs that OLI was designed to process), are like the “black box” on an airliner. If something goes wrong, the logs can be helpful.

Customer support may ask you to retrieve logs as part of an incident investigation. If so, follow the steps below and upload the resulting .zip file to customer support.

To retrieve OLI system logs:

- 1 Click the **Configuration > Retrieve Logs**.

The page shown in [Figure 5-19](#) appears.

- 2 When the Summary Status is Completed, click **Download** to retrieve the system log files are compressed into a single zip file.

Retrieve Snapshot Status		
Summary		
Name	Thread-66	
Request ID	u3RYPkMBABCAAVSP-DsGjg	
Processing Time	4 sec 105 ms	
Status	Processing...	
<hr/>		
Action	Start Time	Time to Complete
Database content	12/29/13 2:33 PM	909 ms
Retrieving logs	12/29/13 2:33 PM	Processing...
<hr/>		
<input type="button" value="Download"/>		

Figure 5-19 Retrieve Logs provides snapshot status.

Content Management

Depending on their rights, users can export Alerts, Dashboards, Filters, Parsers, Saved Searches, and Source Types from an OLI to a file, and then import that content onto another OLI or re-import it onto that same OLI, as a backup. For information on the user rights necessary to import or export a particular type a content, instructions, and guidelines for importing and exporting OLI content, see [“Exporting Content” on page 231](#) and [“Importing Content” on page 230](#).

Content import and export is useful in these situations:

- When you want to make a backup of OLI content. If your OLI becomes unavailable or is reset to its factory defaults, you can quickly restore its content by importing the saved content.
- When multiple OLIs with the same content need to be installed in your network, you need to configure only one OLI. Subsequent OLIs can be deployed by importing the first OLI’s content on them, thus reducing deployment time.
- When you want to add content from one OLI to the content on another.

Using the Export function, you save the content from an OLI to a storage location on your network or to the local disk of the computer from which you connect to the OLI. When you need to use that content for any of the situations described previously, simply import the saved content.

Importing Content

The content you are able to import depends on your user rights. If you have any of the following rights, the Import tab displayed in [Figure 5-20](#) is available:

- Operations Log Intelligence Rights > Filters: Edit, save, and remove shared filters.
- Operations Log Intelligence Rights > Forwarders and Alerts: Edit, save, and remove forwarders and alerts.



Note

While this OLI right enables you to edit, save, and remove both forwarders and alerts, you can only import alerts, but not forwarders.

- Operations Log Intelligence Rights > Dashboards: Edit, save, and remove dashboards.
If the user has the dashboard save right but does not have the saved search save right, then the dashboards using search results panels will not be imported (A warning message will indicate which dashboards are skipped).
- Operations Log Intelligence Rights > Saved Search: Edit, save, and remove saved search.
- System Admin: For parsers and source types, the user can be assigned to any System Admin Group. If the user is not an admin, then Parsers and Source Types are not importable.

Even if you see the import tab, you may not be able to import all of the content types. If you do not have one of the above user rights, then you cannot import that type of content and will get a warning message instead.

Importing Guidelines

Make sure you are familiar with these guidelines before importing OLI content:

- If an object with the same name exists on the importing system, object being imported is named *<ObjectName> [import]*. For example, an imported alert is named *AlertName [import]* and an imported filter is named *FilterName [import]*.
If an object with the name *<ObjectName> [import]* already exists on the importing OLI (from a previous import procedure), the object being imported is named *<ObjectName> [import] [import]*.
- Be sure to set the alert destinations (SNMP, Syslog, and SMTP servers) for alerts you import because this information is not included in the exported content.

To import content from another OLI:

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

- 2 Click **Content Management** in the left panel.
- 3 Click the **Import** tab.

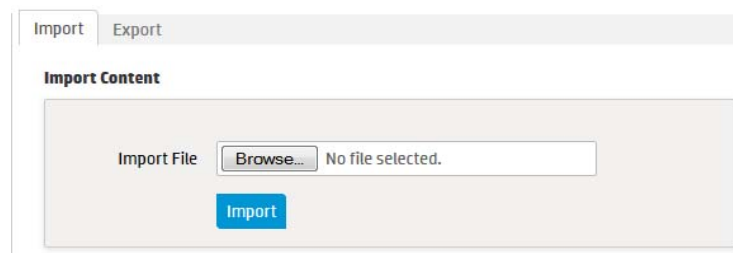


Figure 5-20 Import tab

- 4 Click **Browse** to locate the file

The file must reside on a local or remote drive accessible to the system whose browser you are using to access OLI's user interface.

- 5 Click **Import**.

Exporting Content

The content you are able to export depends on your user rights. If you have any of the following rights, the Export tab displayed in [Figure 5-21](#) is available:

- Operations Log Intelligence Rights > Filters: Use and view shared filters.
- Operations Log Intelligence Rights > Forwarders and Alerts: View forwarders and alerts.



Note While this Operations Log Intelligence right enables you to view both forwarders and alerts, you can only export alerts, but not forwarders.

- Operations Log Intelligence Rights > Dashboards: Use and view dashboards. If the user has the dashboard read right but does not have the saved search read right, then dashboards having search results panels are not available for selection from the Content to Export dialog box.
- Operations Log Intelligence Rights > Saved Search: View saved search.
- System Admin: For parsers and source types, the user can be assigned to any System Admin Group. If the user is not an admin, then Parsers and Source Types are not exportable.

Even if you see the Export tab, you may not be able to export all of the content types. If you do not have one of the above user rights, then the corresponding type of content type is not available in the Content to Export dialog box.

Exporting Guidelines

Make sure you are familiar with these guidelines before exporting OLI content:

- The exported content is in XML format in a gzip file. For example, allfilters.xml.gz.
- The folder on the remote file system to which you are exporting OLI content needs to exist before you can export content to it.
- When exporting alerts, the query associated with the alert, match count, threshold, and status are included in the export. The export does not include e-mail, SNMP or syslog destination information. Since alert destination (SNMP, Syslog, and SMTP servers) information is not exported, you will need to set this information for alerts you import.
- When exporting dashboards, the content of any saved searches used in the exported dashboards is also exported.
- When exporting source types, the content of the parsers used in the exported source types is also exported.

To export OLI content:

- 1 Click **Configuration** on the top-level menu bar.
- 2 Click **Content Management** in the left panel.
- 3 Click the **Export** tab.

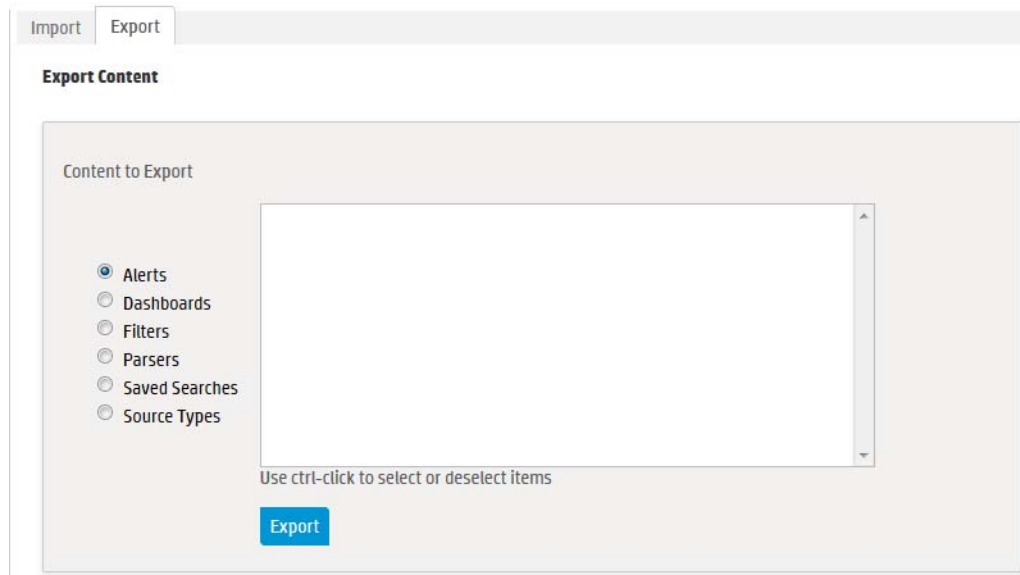


Figure 5-21 Export tab, with content type Filters selected

- 4 Click the button to select the type of content that you want to export. The displayed content changes with the type of content you select.
- 5 Select the objects to export in the **Choose Items to Export** field.
To select one object, click its name. To select multiple objects, hold the Ctrl key down and click the names.
- 6 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**.

Chapter 6: System Admin

This chapter describes the System Administration tools that enable you to create and manage users and user groups, and to configure SMTP and other system settings.

This chapter includes information on the following areas of system administration:

- [“System Locale” on page 235](#)
- [“SMTP” on page 236](#)
- [“License & Update” on page 236](#)
- [“Process Status” on page 237](#)
- [“System Settings” on page 238](#)
- [“Audit Logs” on page 238](#)
- [“SSL Server Certificate” on page 239](#)
- [“SSL Client Authentication” on page 244](#)
- [“FIPS 140-2” on page 247](#)
- [“Authentication” on page 251](#)
- [“Login Banner” on page 260](#)
- [“User Management” on page 261](#)
- [“Change Password” on page 266](#)
- [“Monitoring System Health” on page 266](#)

System

From the System tab, you can configure system specific settings such as network settings (if applicable) and SMTP.

System Locale

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

The System Locale is configured during the OLI installation process. Once configured it cannot be changed.

To view the System Locale:

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

- 2 Click **System Locale** in the **System** section.
The System Locale Setting dialog box displays the Locale.

Impact of Daylight Savings Time Change on Logger Operations

Scheduled operations on OLI such as revent archives and file transfers are impacted when system time is adjusted on the OLI at the start and end of the daylight saving time period (DST). The operations scheduled for the hour lost at the start of DST (for example, on March 1, 2012) are not run on the day of time adjustment. Similarly, operations scheduled for the hour gained at the end of the DST (for example, on November 4, 2012) are run at standard time instead of DST time.

SMTP

Your system uses the Simple Mail Transfer Protocol (SMTP) setting to send email notifications such as alerts and password reset emails.

To add or change SMTP settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SMTP** in the **System** section and enter these settings.

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

- 3 Click **Save**.

License & Update

This page displays license information and the version of the components. To view details of your license, open **Configuration** from the top-level menu bar, and then click **License Information**. For details, see [“License Information” on page 227](#).

Updating the License File

To update a license file:

- 1 Download the update file from the HP Customer Support site (SSO) at <https://support.openview.hp.com> to the computer from which you can connect to the OLI with your browser.

For more information, see [“Acquiring a License for OLI” on page 27](#).

- 2 From the computer to which you downloaded the update file, log in to the OLI user interface using an account with administrator (upgrade) privileges.

- 3 Click **System Admin** from the top-level menu bar.
- 4 Click **License & Update** in the **System** section.
- 5 Browse to the license file you downloaded earlier, and click **Upload Update**.

An “Update In Progress” page displays the update progress.

Once the update has completed, the Update Results page displays the update result (success/failure). If you are only updating a license, a restart is not required.

Process Status

The **Process Status** page lists all processes related to your system and enables you to view the details of those processes and start, stop, or restart them.

To view the Process Status page:

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

Process Status


[Refresh Status](#)

System						
System	Status	Load	CPU Usage	Memory Usage	Data Collected	
LABM3AMRND51	running	[0.19] [0.27] [0.25]	2.4%us 0.3%sy 0.1%wa	28.9% [4760240 kB]	12/18/2013 09:27:04	

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

Processes					
Process	Status	Uptime	CPU Usage	Memory Usage	
Start Stop Restart					
apache	running	17h 32m	0.0%	0.0% [3840 kB]	
aps	running	17h 34m	0.0%	3.5% [579252 kB]	
connector	running	17h 35m	0.0%	0.0% [600 kB]	
mysqld	running	17h 36m	0.0%	0.7% [122580 kB]	
postgresql	running	17h 34m	0.0%	0.0% [9272 kB]	
processors	running	17h 31m	0.0%	1.1% [190280 kB]	
receivers	running	17h 32m	0.0%	1.8% [311088 kB]	
servers	running	17h 34m	0.1%	6.7% [1117336 kB]	
web	running	17h 31m	2.1%	2.1% [348140 kB]	

In the Processes > Process list on this screen, “processors” refers to forwarders.

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

Process Status

[Refresh Status](#)

System						
System	Status	Load	CPU Usage	Memory Usage	Data Collected	
LABM3AMRND51	running	[0.19] [0.21] [0.22]	2.7%us 0.4%sy 0.1%wa	29.0% [4767320 kB]	12/18/2013 09:33:34	

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

Processes						
Start Stop Restart						
Process	Status	Uptime	CPU Usage	Memory Usage		
<div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;"> <ul style="list-style-type: none"> apache Children 0 CPU Percent 0.0% CPU Percent Total 0.0% Data Collected 12/18/2013 09:33:41 Memory Kilobytes 3840 Memory Kilobytes Total 51128 Memory Percent 0.0% Memory Percent Total 0.3% Monitoring Status monitored Parent PID 1 PID 1935 Status running Uptime 17h 39m </div> <div> <ul style="list-style-type: none"> running 17h 39m 0.0% 0.0% [3840 kB] </div> </div>	<ul style="list-style-type: none"> aps connector mysqld postgresql processors receivers servers web 	<ul style="list-style-type: none"> running running running running running running running running 	<ul style="list-style-type: none"> 17h 40m 17h 41m 17h 43m 17h 40m 17h 38m 17h 38m 17h 40m 17h 38m 	<ul style="list-style-type: none"> 0.1% 0.0% 0.0% 0.0% 0.0% 0.0% 0.1% 2.0% 	<ul style="list-style-type: none"> 3.5% [583876 kB] 0.0% [600 kB] 0.7% [122580 kB] 0.0% [9272 kB] 1.1% [190280 kB] 1.8% [311792 kB] 6.7% [1117344 kB] 2.1% [346052 kB] 	

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

System Settings

If you did not select OLI to start as service during the installation process, you can do so using the **System Settings** page. When you select this option OLI will use a service called `arcsight_logger`, enabled to run at levels 2, 3, 4, and 5.

To configure OLI to start as a service:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **System Settings** in the left panel.
- 3 From under **Service Settings**, choose the appropriate option:
 - ◆ Start as a Service
 - ◆ Do not start as a Service
- 4 Click **Save**.

Logs

Your system can generate audit logs at the application and platform levels. Use the Logs sub-menu to search audit logs.

Audit Logs

Your system's audit logs are available for viewing.

Audit Logs

Search Audit Logs

Timestamp --

Description

User

Search Results

User	Description	Timestamp ▾
admin	Successful login	2013/12/18 09:15:09
admin	Failed login attempt	2013/12/18 09:15:09
admin	Session expired	2013/12/17 20:31:27
admin	Saved Search [Unix - Logins by Host] has been updated	2013/12/17 20:13:23
admin	Saved Search [Unix - Logins by Host] has been updated	2013/12/17 20:12:17
admin	Saved Search [Unix - Errors by Name] has been updated	2013/12/17 20:09:24
admin	Saved Search [Unix - Failure Outcome by Host] has been updated	2013/12/17 20:08:08
admin	Saved Search [Unix - Error or Critical or Failure by Host] has been updated	2013/12/17 20:06:32

To view audit logs:

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

Security

Security settings enable you to configure SSL server certificates, enable and disable FIPS (Federal Information Processing Standards) mode on your system, and configure SSL client authentication for client certificate and Common Access Card (CAC) support.



For steps on how to create a user DN, see [“Users” on page 261](#), and refer to the section “Use Client DN” in the parameters table.

SSL Server Certificate

Your system uses Secure Sockets Layer (SSL) technology to communicate securely over an encrypted channel with its clients, such as SmartConnectors, when using the SmartMessaging technology and other ArcSight systems. Your system ships with a self-signed certificate so that an SSL session can be established the first time you use the application. For more information on this option, see [“Generating a Self-Signed Certificate” on page 240](#).

Although a self-signed certificate is provided for your use, HP strongly recommends using a certificate authority (CA) signed certificate. Even if FIPS is

not enabled on your system, it must use a **CA-signed certificate** if it is a destination of a FIPS-enabled container (on a Connector Appliance) or a software-based SmartConnector. Additionally, ensure that the root certificate of the CA that signed your system's certificate is trusted on the container or the SmartConnector. If the CA's root certificate is not trusted, load it on the container by following instructions in the "Managing Certificates on a Container" section in the Connector Appliance Administrator's Guide.

To facilitate obtaining a CA-signed certificate, your system can generate a Certificate Signing Request. Once a signed certificate file is available from the CA, it can be uploaded to your system for use in a subsequent authentication. For detailed instructions, see ["Generating a Certificate Signing Request \(CSR\)" on page 242](#).

Your system generates an audit event when the installed SSL certificate is going to expire in less than 30 days or has already expired. The event with Device Event Class ID "platform:407" is generated periodically until you replace the certificate with one that is not due to expire within 30 days.

Generating a Self-Signed Certificate

Your application has a self-signed certificate so that an SSL session can be established the first time you connect. This type of certificate does not require signing from another entity and can be used immediately.

To generate a self-signed certificate:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Server Certificate** from the **Security** section in the left panel to display the **Generate Certificate/Certificate Signing Request** page.
- 3 Click the **Generate Certificate** tab.

The screenshot shows the 'SSL Settings' interface. At the top, there are two tabs: 'Generate Certificate' (selected) and 'Import Certificate'. Below the tabs is the section title 'Generate Certificate/Certificate Signing Request'. Underneath, there is a form titled 'Enter Certificate Settings' with the following fields:

Country (2-letter code)	US
State/Province	California
City/Locality	Sunnyvale
Organization Name	Hewlett-Packard
Organizational Unit	Support Team
Hostname	[Redacted]
Email Address	arst-support@hp.com
Private Key Length	1024

- 4 From the **Enter Certificate Settings** field, enter new values for the following fields:

Parameter	Description
Country	A two-letter country code, such as 'US' for the United States.
State/Province	State or province name, such as 'California.'
City/Locality	City name, such as 'Sunnyvale'.
Organization Name	Company name, governmental entity, or similar overall organization.
Organizational Unit	Division or department within the organization.
Hostname	The host name or IP address of this system. When specifying the host name, make sure that this name matches the name registered in the Domain Name Service (DNS) server for the system. Note: If the host name or IP address of this system changes in the future, you must generate a new self-signed certificate or CSR. Once a new certificate is obtained, you must upload it to ensure that the connectors (in FIPS mode) which communicate with the system are able to validate the host name.
Email Address	The email address of the administrator or contact person for this CSR.
Private Key Length	Select the length (in bits) of the private key: 1024 , 2048 , 4096 , or 8192 .

Use the first two buttons to generate a CSR or a self-signed certificate. The **View Certificate** button is only used to view the resulting certificate.

Button	Description
Generate CSR	Click to generate a Certificate Signing Request (CSR).
Generate Certificate	Click to generate a self-signed certificate.
View Certificate	Click to view the generated certificate.

- 5 Click the **Generate Certificate** button to generate the self-signed certificate.
- 6 Click **Ok** after the confirmation message appears.
- 7 Click the **View Certificate** button to view the PEM encoded self-signed certificate.

Generating a Certificate Signing Request (CSR)

The first step in obtaining a CA-signed certificate is to generate a Certificate Signing Request (CSR). The CSR must be generated on the system for which you are requesting a certificate. That is, you cannot generate a CSR for System A on System B or use a third-party utility for generation.

The resulting CSR must be sent to a CA, such as VeriSign, which responds with a signed certificate file.

To generate a certificate signing request:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Server Certificate** from the **Security** section in the left panel to display the **Generate Certificate/Certificate Signing Request** page.
- 3 Click the **Generate Certificate** tab.

The screenshot shows the 'SSL Settings' page with the 'Generate Certificate' tab selected. Below the tabs is the 'Generate Certificate/Certificate Signing Request' section. A form titled 'Enter Certificate Settings' contains the following fields:

- Country (2-letter code): US
- State/Province: California
- City/Locality: Sunnyvale
- Organization Name: Hewlett-Packard
- Organizational Unit: Support Team
- Hostname: [Redacted]
- Email Address: arst-support@hp.com
- Private Key Length: 1024

- 4 From the **Enter Certificate Settings** field, enter new values for the following fields:

Parameter	Description
Country	A two-letter country code, such as 'US' for the United States.
State / Province	State or province name, such as 'California.'
City / Locality	City name, such as 'Sunnyvale'.
Organization Name	Company name, governmental entity, or similar overall organization.
Organizational Unit	Division or department within the organization.

Parameter	Description
Hostname	The host name or IP address of this system. When specifying the host name, make sure that this name matches the name registered in the Domain Name Service (DNS) server for the system. Note: If the host name or IP address of this system changes in the future, you must generate a new self-signed certificate or CSR. Once a new certificate is obtained, you must upload it to ensure that the connectors (in FIPS mode) which communicate with the system are able to validate the host name.
Email Address	The email address of the administrator or contact person for this CSR.
Private Key Length	Select the length (in bits) of the private key: 1024 , 2048 , 4096 , or 8192 .

Use the first two buttons to generate a CSR or a self-signed certificate. The **View Certificate** button is only used to view the resulting certificate.

Button	Description
Generate CSR	Click to generate a Certificate Signing Request (CSR).
Generate Certificate	Click to generate a self-signed certificate.
View Certificate	Click to view the generated certificate.

- Choose **Generate CSR** to generate a certificate signing request.
- If the CSR was successfully generated, a pop-up window appears, allowing you to either download the CSR file or to cut-and-paste its content.

To do so, copy all the lines from `-----BEGIN CERTIFICATE REQUEST-----` to `-----END CERTIFICATE REQUEST-----`.
- Send the CSR file to your certificate authority to obtain the CA-signed certificate.
- Once the CA-signed certificate file is obtained, continue on to [Importing a Certificate](#) below.

Importing a Certificate

If you have obtained a certificate from your certificate authority (CA), follow the steps below to import it onto your system.

- Click **System Admin** from the top-level menu bar.
- Click **SSL Server Certificate** under the **Security** section in the left panel.
- Select the **Import Certificate** tab.

SSL Settings

Generate Certificate Import Certificate

Import Signed Certificate

Note: After uploading the new certificate, close and re-open the browser.

Select the Signed Certificate File

Browse...

Import and Install

Last Import Status

Last SSL Certificate Import Status

--- No Status Exists ---

- 4 Click the **Browse** button to locate the signed certificate file on your local file system.



The imported certificate must be in **Privacy Enhanced Mail (PEM)** format.

Note

- 5 Click **Import and Install** to import the specified certificate.
- 6 If using **HTTPS** and depending on your browser, you may need to close and restart the browser for the new certificate to take effect. If you are unsure of your browser's requirements, close and restart it.

SSL Client Authentication

Your system supports client authentication using SSL certificates. SSL client authentication is a form of two-factor authentication that can be used *as an alternate* or *in addition to* local password authentication. As a result, your system can be configured for SmartCards, such as Common Access Card (CAC) based authentication. CAC is a standard identification card for active duty members of the Uniformed Services, Selected Reserve, DOD civilian employees, and eligible contractor personnel.



CAC is a form of client certificate authentication. Information on client certificate authentication applies to CAC.

Note

Your system also supports LDAPS authentication. The SSL certificate for the LDAPS server must be uploaded into the trusted store. After uploading the SSL certificate, the aps process must be restarted (**System Admin > Process Status > aps > Restart**).

Configuring OLI to Support SSL Client Authentication

Perform the following steps to configure OLI to support SSL client authentication.

On the OLI:

- 1 If the OLI uses the default signed certificate, replace it with a *FIPS-compliant*, signed SSL server certificate. Follow instructions at [“Uploading Trusted Certificates” on page 245](#) to load the certificate.

**Caution**

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

- 2 Enable client certificate authentication, as described in [“Client Certificate Authentication” on page 255](#).
- 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 245](#).

If the client certificates used to authenticate with OLI are signed by different CAs, make sure you upload root certificates of **all** CAs.

If the client certificates are **self-signed**, upload the public portion of the client certificate.
- 4 Configure a user name for each user who will be connecting to the OLI using a client certificate, as described in [“User Management” on page 261](#).
- 5 (Optional) Upload a certificate revocation list (CRL), as described in [“Uploading a Certificate Revocation List” on page 246](#).
- 6 (Optional) If this OLI is configured to use **only** SSL Client Authentication, make sure this OLI’s Authorization ID and Code are appropriately configured on other OLIs that with it. For more information, see [“Peer OLIs” on page 203](#).

On the Client (Web browser):

Configure your browser to provide the SSL client certificate when accessing OLI. (Upload the private key in PKCS 12 format in your browser.)

Uploading Trusted Certificates

A trusted certificate is used to authenticate users that log in to your system. Uploading a trusted certificate is required if you are using LDAPS authentication. The trusted certificate is used to authenticate the remote LDAPS server. The certificate needs to be in Privacy Enhanced Mail (PEM) format.

To upload a trusted certificate:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Client Authentication** in the **Security** section in the left panel.
- 3 On the **Trusted Certificates** tab, click **Browse** to find the trusted certificate on your local file system.
- 4 Click **Upload**.

The trusted certificate is uploaded and listed in the “Certificates in Repository” list on the same page where you uploaded it.

Trusted Certificates | Certificate Revocation List

Trusted Certificates

Upload Certificate

The certificate file needs to be in the PEM format.

Upload File:

Certificates in Repository

<input checked="" type="checkbox"/>	Certificate Name	Start date	Expiration Date

To view details about a trusted certificate, click the link displayed in the Certificate Name column.

To delete a trusted certificate, select the certificate and click **Delete**.

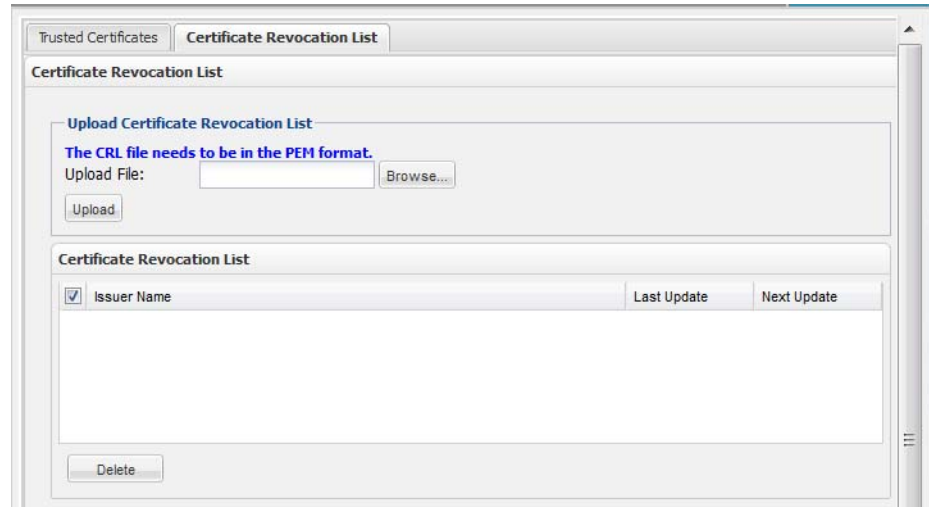
Uploading a Certificate Revocation List

A certificate revocation list (CRL) is a computer-generated record that identifies certificates that have been revoked or suspended before their expiration dates. To support CAC, you need to upload a CRL file to your HP system. The CRL file needs to be in PEM format.

To upload a CRL file:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Client Authentication** in the **Security** section in the left panel.
- 3 In the **Certificate Revocation List** tab, click **Browse** to find the CRL file on your local file system.
- 4 Click **Upload**.

The CRL is uploaded and listed in the Certificate Revocation List.



To view details about a CRL, click the link displayed in the Issuer Name column.

To delete a CRL file, select it and click the **Delete** button.

Enabling Client Certificate Authentication

To enable client certificate authentication, see [“Client Certificate Authentication” on page 255](#).

FIPS 140-2

Your system supports the Federal Information Processing Standard 140-2 (FIPS 140-2). FIPS 140-2 is a standard published by the National Institute of Standards and Technology (NIST) and is used to accredit cryptographic modules in software components. The US Federal government requires that all IT products dealing with Sensitive, but Unclassified (SBU) information meet these standards.

If your system needs to be FIPS 140-2 compliant, you can enable FIPS. Once you do so, the system uses the cryptographic algorithms defined by the NIST for FIPS 140-2 for all encrypted communication between its internal and external components.



To be fully FIPS 140-2 compliant, all components of your OLI deployment need to be in FIPS 140-2 mode. For example, if you enable FIPS 140-2 on your OLI 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 OLI 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 249 to ensure that your connector is FIPS compliant.

- OLI forwarders, such as ArcSight Managers to which OLI forwards events and alerts

The system to which your FIPS-compliant OLI forwards events should be FIPS compliant as well. Additionally, you need to import that system’s SSL server certificate on the OLI so that OLI can communicate with it.

- OLIs

OLIs running 45.0 or later automatically use FIPS 140-2 compliant algorithms. Therefore, no action is required on such OLIs, except enabling FIPS as described in this section.

When enabling FIPS on a OLI, make sure that the machine on which OLI is installed is used exclusively for OLI.



Note

Enabling FIPS 140-2 on OLI does not make the system on which it is installed FIPS 140-2 compliant. Consult your system’s documentation to determine the requirements for making the entire system FIPS 140-2 compliant.

You can enable or disable FIPS mode on OLI to suit your needs; however, you will need to restart the OLI before the new mode will be effective.

Things to be Aware of When Enabling FIPS Mode on OLI

- Your OLI must be set up with a CA-signed SSL certificate. For more information, see [“SSL Server Certificate”](#) on page 239.
- An OLI, even when in non-FIPS mode, must use a CA-signed certificate if it is a destination of a FIPS-enabled software-based SmartConnector. Additionally, ensure that the root certificate of the CA that signed OLI’s certificate is trusted on the SmartConnector. If the CA’s root certificate is not trusted on the SmartConnector, follow instructions in [“Installing or Updating a SmartConnector to be FIPS compliant”](#) on page 249.
- Once FIPS is enabled on your OLI, the SmartMessage receiver (if configured) stops receiving events from non-FIPS connectors if those connectors are not running version 4.7.5.5372 or later. Make sure you have the correct connectors.

To enable or disable FIPS mode:



Note

Make sure you are familiar with the configuration requirements on your OLI as described in [“Things to be Aware of When Enabling FIPS Mode on OLI”](#) on page 248.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **FIPS 140-2** in the Security section in the left panel.
- 3 Click **Enable** or **Disable** for the Select FIPS Mode option.
- 4 Click **Save**.

- 5 Use the following command to restart your OLI:

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

The FIPS Status Table shows which processes and components of the OLI are FIPS enabled.

Installing or Updating a SmartConnector to be FIPS compliant

FIPS mode is supported on SmartConnectors running version 4.7.5.5372 or later.

If you are...	Then...
Installing a new SmartConnector to send events to an OLI in FIPS-compliant mode	<ol style="list-style-type: none"> 1 Download a FIPS-supported SmartConnector version from the HP Customer Support site (SSO). 2 Go to Step 1 on page 249.
Updating a SmartConnector to be FIPS compliant and the SmartConnector is not running version 4.7.5.5372 or later.	<ol style="list-style-type: none"> 1 Upgrade the SmartConnector to a FIPS-supported version. Follow instructions in the SmartConnector User's Guide to upgrade the SmartConnector. 2 Perform only Step 2a on page 249.
Updating a SmartConnector to be FIPS compliant and the SmartConnector is running version 4.7.5.5372 or later.	Perform only Step 2a on page 249 .

To make a SmartConnector FIPS compliant:

- 1 Follow device configuration steps provided in the SmartConnector's configuration guide (available from the HP Customer Support site (SSO) at <http://support.openview.hp.com>), then follow the installation procedure through installation of the core connector software (SmartConnector Installation step 2).

At Step 3 of the Connector setup, click **Cancel** to exit the setup. You must then configure the NSS DB, which is necessary for installing the connector in FIPS-compliant mode.

Once the NSS DB is configured, continue with [Step 2](#), below.

- 2 To enable FIPS Mode on the SmartConnector:
 - a Create an `agent.properties` file at the following location if it does not exist already:

```
$ARCSIGHT_HOME\current\user\agent
```

- b Enter the following property, then save and close the file.

```
fips.enabled=true
```

- 3 Import OLI's Certificate on the SmartConnector:

- a In a DOS prompt window on your SmartConnector machine, from `$ARCSIGHT_HOME\current\bin`, enter the following command to turn off FIPS mode.

```
arcsight runmodutil -fips false -dbdir
user/agent/nssdb.client
```

- b Export the OLI certificate file and import it to the SmartConnector's NSS DB as follows:

- i Export OLI's certificate file from the browser you use to connect to it. Refer to your browser's Help for instructions. Save the certificate file with a `.crt` or `.cer` extension.

- ii Copy the certificate file you exported in the previous step (in this example, **loggercert.crt**) to the `$ARCSIGHT_HOME\current\bin` directory on the SmartConnector. From `$ARCSIGHT_HOME\current\bin`, enter the following:

```
arcsight runcertutil -A -n mykey -t "CT,C,C" -d
user/agent/nssdb.client -i bin/loggercert.crt
```

- c Enter the following command to re-enable FIPS mode that you turned off in Step 1:

```
arcsight runmodutil -fips true -dbdir
user/agent/nssdb.client
```

- d Ensure that the SmartConnector can resolve the name specified in the CN value of the OLI certificate's *Subject* field. If the name is not resolvable, add it to SmartConnector system's Hosts file.

- e If you are installing a new SmartConnector, continue to the next step.

If you are updating your SmartConnector to be FIPS compliant, ensure that the connector's OLI destination host name is same as the CN value in the certificate's *Subject* field, and **exit this procedure**.

- 4 To return to the SmartConnector configuration wizard, enter the following from `$ARCSIGHT_HOME\current\bin`:

```
arcsight connectorsetup
```

- 5 When prompted whether you want to start in Wizard Mode, click **Yes**.

The Destination selection window is again displayed; return to your SmartConnector Configuration Guide, **SmartConnector Installation step 4** to continue the connector configuration.



When configuring the connector, ensure that the connector's OLI 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 HP events.

Users/Groups

Use the **Users/Groups** sub-menu to configure users and user groups, and to set authentication options.

Authentication

Authentication Settings enable you to specify the settings and policies for user login sessions, password rules and lockouts, and external authentication options.

Sessions

The **Session** tab lets you specify the maximum number of simultaneous sessions for a single user account, and the length of time after which a user session is automatically logged out or a user account disabled. By default, a single user account can have up to 15 simultaneous active sessions, and a user account is logged out after 15 minutes of inactivity.

To change session settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 On the **Sessions** tab, update the parameters described in the following table.

Parameters	Description
Max Simultaneous Logins/User	The maximum number of simultaneous sessions allowed for a single user account. The default is 15 sessions .
Logout Inactive Session After	The length of time, in minutes, after which an inactive session is automatically ended. The default is 15 minutes . This value does not apply to the user interface pages accessed through the Monitor menu. If a user is on any of the Monitor menu pages and the session has been inactive for the specified number of minutes, the user's session remains active.

Parameters	Description
Disable Inactive Account After	The number of days after which an inactive user account is disabled. The default is 0 , meaning the account is never disabled.

- 4 Click **Save** to make the changes, or click another tab to cancel.

Local Password

The **Local Password** tab enables you to set password policies, such as the minimum and maximum number of characters and other password requirements.

Authentication Settings

Sessions Local Password External Authentication

Local Password Settings

Lockout Account

Enable Account Lockout

Lockout Account After Failed Attempts

Remember Failed Attempts For hours minutes seconds

Lockout Account For hours minutes

Password Expiration

Enable Password Expiration

Password Expires in days

Notify User Days Before Expiration

[Users Exempted From Password Expiration Policy \(0\)](#)

Password Strength Rules

Enforce Password Strength

Minimum Length characters

Maximum Length characters

Password Character Rules

Password must have a minimum of the following characters

Numeric [0-9] Uppercase [A-Z]

Special [!\$^*...] Lowercase [a-z]

To change the password settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **Local Password** tab.

Use the parameters described in the following table to customize your password settings.

Table 6-1 Authentication Settings, Local Password tab

Parameter	Description
Lockout Account (policy)	
Enable Account Lockout	Select the checkbox to enable user accounts to be locked out as defined by the following settings. By default, the policy is disabled .
Lockout Account After	Number of failed login attempts after which a user account is locked out. The default is 3 .
Remember Failed Attempts For	The length of time, in minutes, for which a failed login attempt is remembered. The default is 1 .
Lockout Account For	The length of time, in minutes, for which a locked out account cannot be unlocked. The default is 15 .
Password Expiration (policy)	
Enable Password Expiration	Select the checkbox to enable user passwords to expire as defined by the following settings. By default, the policy is disabled .
Password Expires in	Number of days after which the password expires. The default is 90 .
Notify User	Number of days before expiration to notify the user. Select this option to allow users to update their password before expiration. The default is 5 .
Users Exempted From Password Expiration Policy	Click the link to set the number of users whose password should never expire. For information on how to use this feature, see “Users Exempted From Password Expiration” on page 254 .
Password Strength Rules (policy)	
Enforce Password Strength	Select the checkbox to enforce password policy as defined by the following settings. By default, the policy is disabled .
Minimum Length	Minimum number of characters that a password must contain. The default is 10 .
Maximum Length	Maximum number of characters that a password can contain. The default is 20 .
Password Character Rules	
Password character rules define additional character requirements to ensure password strength.	
Numeric	Minimum number of numeric characters (0-9) in a password. The default is 2 .

Table 6-1 Authentication Settings, Local Password tab (Continued)

Parameter	Description
Uppercase	Minimum number of uppercase characters (A-Z) in a password. The default is 0 .
Special	Minimum number of non-digit and non-letter characters that are required in a password. The default is 2 .
Lowercase	Minimum number of lowercase characters (a-z) in a password. The default is 0 .
Password Must be At Least N Characters Different From Old Password	Minimum number of characters by which the new password must differ by from the previous one. The default is 2 .

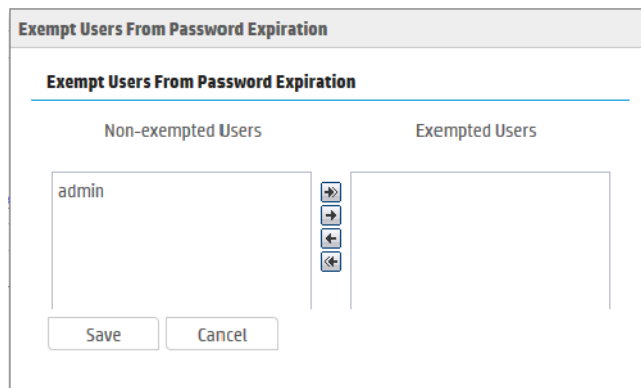
- 4 Click **Save** to save the changes, or click another tab to cancel.


Users Exempted From Password Expiration


Even though you have set a password expiration policy for most users, you may want to have a user whose password does not expire automatically.

To exempt a user from the password expiration policy:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **Local Password** tab, and then click **Users Exempted From Password Expiration Policy**.
- 4 The **Exempt Users From Password Expiration** page appears.



- 5 Select users from the **Non-exempted Users** list and click the right arrow icon  to move the selected users to the **Exempted Users** list. Do the reverse to remove users from the list of exempted users.

You can select multiple users at the same time and move them over. Or you can move all users by clicking the  icon.

- 6 Click **Save** to save the policy or **Cancel** to exit.

External Authentication

Besides providing a local password authentication method, your system supports Client Certificate/CAC, LDAP, and RADIUS authentication. It is not possible to enable all authentication methods simultaneously.



Note

CAC is a form of client certificate authentication. Information on client certificate authentication applies to CAC.

From the **External Authentication** tab, use the drop-down menu to choose one of the following authentication methods:

- [Local Password](#)
- [Client Certificate Authentication](#)
- [Client Certificate and Local Password Authentication](#)
- [LDAP/AD and LDAPS Authentication](#)
- [RADIUS Authentication](#)

Local Password

This option is the default method and implements the local password policies set in the **Local Password** tab. Leave this as the default, or click **Save** if changing from another option.

Client Certificate Authentication

This authentication method requires that users authenticate using a client certificate. For each client certificate, a user account with a Distinguished Name (DN) matching the one in the client certificate must exist on your system.



Caution

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

To configure client certificate authentication:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **External Authentication** tab.
- 4 From the drop-down menu, choose **Client Certificate**.
- 5 **Allow Local Password Fallback** provides two options:
 - ◆ **Allow Local Password Fallback for Default Admin Only**
Select this option to allow the default admin user to log in using only a username and password if the client certificate is not available or invalid. This privilege is restricted to the default admin user only—other users

must have a valid client certificate to gain access to the system. This option is enabled by default.

◆ **Allow Local Password Fallback for All Users**

Select this option to allow all users to log in using their local user name and password if their client certificate is invalid or unavailable.

For more information, see [“Local Password Fallback” on page 259](#).

6 Click **Save**.

Client Certificate and Local Password Authentication

This authentication method requires that users authenticate using an SSL client certificate and a valid local password. *Local Password* refers to the password associated with the user credentials created in **User Management** in the **Users/Groups** section. See [“User Management” on page 261](#) for details.

A user account on your system must be defined with a Distinguished Name (DN) that matches the one in the client certificate.

For instructions on how to create a user DN, see [“Users” on page 261](#) and refer to the section called “Use Client DN” in the parameters table.



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

To configure client certificate and password authentication:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **External Authentication** tab.
- 4 From the drop-down menu, choose **Client Certificate AND Local Password**.
- 5 **Allow Local Password Fallback** provides two options:

◆ **Allow Local Password Fallback for Default Admin Only**

This option, always enabled, allows the default admin user to log in using only a username and password.

◆ **Allow Local Password Fallback for All Users**

This option is always disabled. You cannot enable it when using the **Client Certificate AND Local Password** authentication method.

For more information, see [“Local Password Fallback” on page 259](#).

6 Click **Save**.

LDAP/AD and LDAPS Authentication

This authentication method authenticates users against an LDAP server. Even when LDAP is enabled, each user account must exist locally on your system. Although the user name specified locally can be different from the one specified

on the LDAP server, the Distinguished Name (DN) specified for each user account must match the one in the LDAP server.



For steps on how to create a user DN, see [“Users” on page 261](#), and the parameter [“Use Client DN” on page 262](#)”.

To set up LDAP authentication:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **External Authentication** tab.
- 4 From the drop-down menu, choose **LDAP**.
- 5 **Allow Local Password Fallback** provides two options:
 - ◆ **Allow Local Password Fallback for Default Admin Only**
Select this option to allow the default admin user to log in using only a username and password if LDAP authentication fails. This privilege is restricted to the default admin user only—all others must be authenticated by LDAP. This option is enabled by default.
 - ◆ **Allow Local Password Fallback for All Users**
Select this option to allow all users to log in using their local user name and password if LDAP authentication fails.

For more information, see [“Local Password Fallback” on page 259](#).

LDAP Server has the following parameters:

Parameter	Description
Server Hostname[:port] (optional)	(Optional) Enter the host name or IP address and port of the LDAP server in the following format: <code>ldap://<hostname or IP address>:<port></code> <code>ldaps://<hostname or IP address>:<port></code> Additional steps are required for the use of LDAPS. See Using the LDAP over SSL (LDAPS) Protocol below.
Backup Server Hostname[:Port] (optional)	(Optional) Enter the backup LDAP server to use if the primary server does not respond. If the server returns an authentication failure (bad password, unknown username, etc), then the backup server is not tried. The backup server is tried only when the primary server has a communication failure. Use the same format as the primary server to specify the host name and port.
Request Timeout	The length of time, in seconds, to wait for a response from the LDAP server. The default is 10 .

6 When finished, click **Save**.

Using the LDAP over SSL (LDAPS) Protocol

When choosing the LDAPS protocol to authenticate users, make sure the following conditions are true:

- The SSL certificate for the LDAPS server has been uploaded into the trusted store.
- The external authentication method is set to "LDAP".
- The URL for the LDAPS server(s) starts with "ldaps://".

After uploading the SSL certificate, the **aps** process must be restarted (**System Admin > Process Status > aps Restart**).



Caution

If the aps process is not restarted, attempts to authenticate via LDAPS will fail.

RADIUS Authentication

This authentication method allows users to authenticate against a RADIUS server. Even when RADIUS authentication is enabled, each user account must exist locally on your system. The username must match the one in the RADIUS server, although the password can be different. A user must present a valid username and (RADIUS) password to be successfully authenticated.

To configure RADIUS authentication settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **External Authentication** tab.
- 4 From the drop-down menu, choose **RADIUS**.
- 5 **Allow Local Password Fallback** provides two options:
 - ◆ **Allow Local Password Fallback for Default Admin Only**
Select this option to allow the default admin user to log in using only a username and password if RADIUS authentication fails. This privilege is restricted to the admin user only—all others must be authenticated by RADIUS. This option is enabled by default.
 - ◆ **Allow Local Password Fallback for All Users**
Select this option to allow all users to log in using their local user name and password, if RADIUS authentication fails. For more information, see ["Local Password Fallback" on page 259](#).

6 Update the RADIUS Server parameters as necessary:

Parameter	Description
Server Hostname[:port]	Enter the host name and port of the RADIUS server.
Backup Server hostname[:port] (optional)	(Optional) Enter the backup RADIUS server to use if the primary server does not respond. If the server returns an authentication failure (bad password, unknown username, etc), then the backup server is not tried. The backup server is tried only when the primary server has a communication failure. Use the same format as the primary server to specify the host name and port.
Shared Authentication Secret	Enter a RADIUS passphrase.
NAS IP Address	The IP address of the Network Access Server (NAS).
Request Timeout	The length of time, in seconds, to wait for a response from the RADIUS server (in seconds). The default is 10 .
Retry Request	Number of times to retry a RADIUS request. The default is 1 .
RADIUS Protocol:	Use the drop-down menu to choose a protocol option. The default is None .

7 Click **Save**.**Local Password Fallback**

You can use this feature to log in using your local user name and password if the external authentication (Certificate, LDAP, or RADIUS) fails, if you forgot your password to the authentication server, or if the authentication server is not available.

The Use Local Authentication allows the default admin to log in even when the remote authentication server is not available, by adding a **Use Local Authentication** checkbox to the login screen. Out-of-box, this option is enabled only for the default administrator. However, it is possible to allow local password fallback for all users. For example, you could configure the RADIUS authentication method to allow users to log in using local authentication instead of RADIUS should they fail to authenticate to the configured external RADIUS server(s).

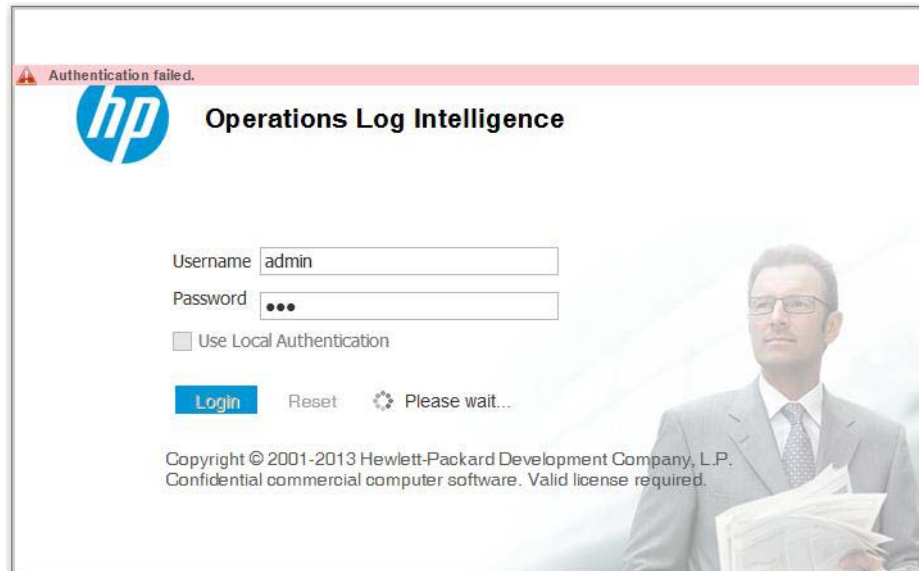
For information on how to allow local password fallback for all users for all users, see [“Client Certificate Authentication” on page 255](#), [“LDAP/AD and LDAPS Authentication” on page 256](#), or [“RADIUS Authentication” on page 258](#).

To log in when authentication fails:

- 1 Mark the **Use Local Authentication** checkbox if the login failure was caused by failure of the external authentication.



This option is only available to the default admin unless it has been enabled for other users.



- 2 Enter your login and password and click **Login**.

Login Banner

You can customize the message on the login screen to suit your needs. The text you enter in the Content field is displayed above the Username and Password fields on the login screen. In addition, you can enter a confirmation message that the user must click to enable the Username and Password fields.

The screenshot shows a web-based configuration window titled "Login Banner". The window contains two main text input fields. The first field, labeled "Content", contains the text "This is a restricted access system...". The second field, labeled "Confirmation", is currently empty. Below the "Confirmation" field is a "Save" button.

You must have the “Configure Login Settings” permission enabled for your user account to edit the login banner.

To customize the login banner:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Login Banner** in the **Users/Groups** section.
- 3 Enter the text you want to display as the login banner in the Content field.

You can enter only unformatted text in this field; however, you can apply standard HTML tags to display formatted text. Loading images in this field is not allowed.

- 4 (Optional) Enter text in the Confirmation field.

If you enter text in this field, the text will be accompanied by a check box that the user must click to enable the Username and Password fields. For example, if you enter “Are you sure?”, “Do you want to proceed?”, or “I agree.” in this field, the user must click the checkbox in order to log in.

- 5 Click **Save**.

User Management

The **Users** and **Groups** tabs enable you to manage users and user groups on your system. User groups are a way to enforce access control to various sections of your system.

Users

Open the **Users** tab to manage the users that can log in to your system. You can add a new user, edit user information, or delete a user at any time. You must have the appropriate System Admin group rights to perform these functions.

To add a new user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 In the **Users** tab, click **Add** from the top left side of the page.
- 4 Enter the following parameters.

Parameter	Description
<i>Credentials</i>	
Login	The user's login name.
Password	The user's password.
Confirm Password	Reenter the users' password.
<i>Contact Information</i>	
Use Client DN	<p>If you enabled SSL client certificate or LDAP authentication, click this link to enter user's the Distinguished Name (Certificate Subject) information. The Distinguished Name should be similar to this format:</p> <p>CN=UserA,OU=Engg Team,O=HP\, Inc.,L=Cupertino,C=US,ST=California</p> <p>To determine the DN, use this URL to display the certificate:</p> <p>https://<hostname or IP address>/platform-service/DisplayCertificate</p> <p>OR</p> <p>Obtain the DN information for a user from the browser that the user will open to connect to the system. For example, on Firefox, click Tools > Options > Advanced > Encryption > View Certificates > Your Certificates > Select the certificate > View.</p>
First Name	The user's first name.
Last Name	The user's last name.
Email	The user's email address.
Phone Number	(Optional) The user's phone number.
Title	(Optional) The user's title.
Department	(Optional) The user's department.
Fax	(Optional) The user's fax number.
Alternate Number	(Optional) The user's alternate phone number.

Parameter	Description
<i>Assign to Groups</i>	Select the groups to which this user belongs. This setting controls the privileges a user has on this OLI.
Notes	(Optional) Other information about the user.

- 5 Click **Save and Close**.

To edit a user:

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

To delete a user:

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

Reset Password

The Reset Password feature enables you to reset a user's password without knowing their password. If you are using an SMTP-configured server and have permissions to create and update users, you can reset a user's password by clicking the **Reset Password** button. An automated email is sent to the user with the new password string.

An SMTP server must be configured for the automated email containing the temporary password to be sent. If an SMTP server is not configured, the password will not be reset because an email cannot be sent.

To reset a user's password:

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

The user must use the temporary string to log in within the time specified in the email. If the user does not log in within the specified time, the account becomes deactivated. If the account has been deactivated, the admin must re-activate it before resetting the password.

To activate a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 In the **Users** tab, select the user (or users) that you want to activate.
- 4 Choose **Edit**.
- 5 Check the **Active** box.
- 6 **Save** the changes.

Groups

User groups define privileges to specific functions on your system and serve to enforce access control to these functions.

User groups are organized by the following types: System Admin, OLI Rights, and OLI Search. Each type has a pre-defined, default user group in which all privileges for the type are enabled. To authorize a subset of the privileges for a specific group type, create a new user group and enable only the privileges you want to provide for that group. Then, assign restricted users to the newly created group.

System Admin Groups

The System Admin Group controls the system administration operations for your system, such as configuring network information, setting storage mounts, installing SSL certificates, and user management.

Read Only System Admin Group

In addition to the default System Admin Group that enables all rights (privileges), a Read Only System Admin Group is available on your system. Users assigned to this group can view System Admin settings, but cannot change them.

Refer to your system's user interface for a complete list of privileges available to this group.

OLI Rights Group

The OLI Rights Group controls the OLI application operations for your system, such as viewing the OLI dashboards and configuring all the settings in the Configuration menu (including event archives, storage groups, alerts, filters, and scheduling tasks.)

Refer to your system's user interface for a complete list of privileges available to this group.

OLI Search

The OLI Search Group controls local and peer searches through the following privileges:


- Search for events

- Search for events on remote peers

If the group is configured to allow users to run local and peer searches, users assigned to this group can perform those operations. Conversely, if the group is configured to prevent users from running local and peer searches, users assigned to this group cannot perform those operations.

Managing a User Group

To create a new user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 Click the **Groups** tab.
- 4 Click **Add** from the top left side of the page.
- 5 Define the new group:
 - a In the **Group Name** field, provide a name for the group.
 - b In the **Description** field, provide a description for the group.
 - c From the Group Type drop-down box, select the group type.
 - d Click the down arrow icon () next to the group type name to view and select privileges that you want to assign to the users in this group.
- 6 Click **Save and Close** to save the settings of the group, or click **Save and Edit Membership** to add users to this group.

To edit a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 Click the **Groups** tab.
- 4 Select the group that you want to edit, and click **Edit** at the top left side of the page.
- 5 Update the user group information.

If you need to edit the group's membership:

- a Click **Save and Edit Membership** to display the Edit Group Membership page.
- b Click **Add** from the top left of the Edit Group Membership page.
- c Select users you want to add. By default, you can add only users who do not belong to other groups of the type that you are editing. To add such users, click **Show users that belong to other <group_type> groups**.

When you add a user who belongs to another group of the same group type as the one you are updating, that user is automatically removed from the previous group.

- d Click **OK**.

- e Click **Back to Group List**.
- 6 Click **Save and Close**.

To delete a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 Click the **Groups** tab.
- 4 Select the group (or groups) that you want to delete.
- 5 Click **Delete** at the top left side of the page.

Change Password

You can use the **Change Password** menu to change your password. Passwords are subject to the password policy specified by the Admin user.

To change your password:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Change Password** in the **Users/Groups** section in the left panel to display the Change Password for <User Name> page.
- 3 Enter the Old Password, the New Password, and enter the New Password a second time to confirm.
- 4 Click **Change Password**.

Other System Administration Information

This section contains information related to system administration that you will need to fully administer your OLI, including system health events and SNMP polling.

Monitoring System Health

You can monitor your OLI's health in these ways:

- By using a pre-defined system filter, as listed in [“System Filters/Predefined Filters” on page 119](#). The pre-defined system health filters are based on the system health events listed in [“System Health Events” on page 267](#).
- By searching for system health events in OLI's Internal Storage Group, as listed in [“System Health Events” on page 267](#). If a pre-defined system health filter does not suit your needs, you can create alerts based on the system health events.

To set up notification of system health events:

- 1 Configure the OLI's SMTP settings (see [“SMTP” on page 236](#)) or create an SNMP Destination (see [“Sending Notifications to SNMP Destinations” on page 185](#)) or Syslog Destination (see [“Sending Notifications to Syslog Destinations” on page 186](#)).

- 2 Create an Alert that uses one or more System Alert Filters or define a query that searches for the system health events in OLI's Internal Storage Group, and specify match count and threshold (see ["Alerts" on page 172](#)).
- 3 Enable the new Alert.

System Health Events

The following table lists the system health events that OLI generates. These events are also referred as OLI Internal Events because they are stored in OLI's Internal Storage Group. See [Appendix F, Examples of System Health Events on page 151](#) for examples of these events.

The pre-defined System Filters that provide system health status are based on some of these events. If a pre-defined filter does not suit your needs, create an alert using one of these events.

All hardware-related events are classified as `hardware:nnn` events, where `nnn` is a three-digit number that identifies the hardware component (for example, `hardware:13x` identifies the fan events.)

The following table lists the Device Event Class IDs.

Group	Device Event Category	Device Event Class ID
System Health Events for OLIs		
CPU	/Monitor/CPU/Usage	cpu:100
Disk	/Monitor/Disk/Read	disk:102
	/Monitor/Disk/Write	disk:103
EPS	/Monitor/Receiver/EPS/All	eps:100
	/Monitor/Receiver/EPS/Individual	eps:102
	/Monitor/Forwarder/EPS/All	eps:101
	/Monitor/Forwarder/EPS/Individual	eps:103
Memory	/Monitor/Memory/Usage/Platform	memory:100
Network	/Monitor/Network/Usage/In	network:100
	/Monitor/Network/Usage/Out	network:101
Search	/Monitor/Search/Performed	search:100
Storage Group	/Monitor/StorageGroup/Space/Used	storagegroup:100
	Note: The size of the storage group, indicated by the "fsize" field is in GB.	0

Appendix A: Search Operators

This appendix describes the operators you can use in search queries you specify in the Search box (Analyze > Search) and gives examples of their use.

This appendix provides information on the following search operators.

- [“chart” on page 269](#)
- [“dedup” on page 276](#)
- [“eval” on page 277](#)
- [“extract” on page 277](#)
- [“fields” on page 279](#)
- [“head” on page 280](#)
- [“keys” on page 280](#)
- [“parse” on page 281](#)
- [“rare” on page 282](#)
- [“regex” on page 283](#)
- [“rename” on page 283](#)
- [“replace” on page 284](#)
- [“rex” on page 286](#)
- [“sort” on page 288](#)
- [“tail” on page 289](#)
- [“top” on page 289](#)
- [“transaction” on page 290](#)
- [“where” on page 292](#)

chart

Displays search results in a chart form of the specified fields.

Usage

```
... | chart <field>
```

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

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

```
... | chart {<function> (<field>)} as <new_column_name> by <field>  
[span [<time_field>]=<time_bucket>]
```

where

`<field>`, `<field1>`, `<field2>` are the names of the field that you want to chart. The fields can be either event fields available in the OLI schema or a user-defined fields created using the `rex` or `eval` operator prior in the query.

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

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

`<new_column_name>` is the name you want to assign to the column in which the function's results are displayed. For example, `Total`.


Deprecated Usage

The following deprecated usage contains “`_count`”. The recommended usage, as shown above, is “`count`”.

```
...| chart _count by <field1> <field2> <field3> ...
```

Notes

By default, a column chart is displayed. Other chart types you can select from: bar chart, line chart, pie chart, area chart, stacked column, or stacked bar.

To change the chart settings (including its type), click  to the upper right corner of the Result Chart frame of the screen. You can change these settings:

- **Title:** Enter a meaningful title for the chart.
- **Type:** Column, Bar, Pie, Area, Line, Stacked column, Stacked Bar. The last two types create stacked charts in which multiple values are plotted in a stack form. These charts are an alternate way of representing multi-series charts, which are described below.
- **Display Limit:** Number of unique values to plot. Default: 10

If the configured Display Limit is less than the number of unique values for a query, the top values equal to the specified Display Limit are plotted. That is, if the Display Limit is 5 and 7 unique values are found, the top 5 values will be plotted.

All chart commands except “`count by`” accept only *one field* in the input. The specified field must contain numeric values.

If multiple fields are specified, separate the field names with a white space or a comma.

The `chart <field>` command does not aggregate field values. It simply lists and charts each occurrence of the values of the specified field. For example, `chart sourcePort`. However, when you use an **aggregation function** such as `count by`, `sum`, `avg` (or `mean`), and so on, an aggregation of the specified fields is performed and charted, as illustrated in “[Example 1](#)” on page 273.

You can click on a charted value to quickly filter down to events with specific field values. For more information, see [“Chart Drill Down” on page 104](#).

Aggregation Functions

If an aggregation function such as `count`, `sum`, or `avg` is specified, a chart of the aggregated results is displayed along with the tabular results of the aggregation operation in a Results Table. For example, for the aggregation function `sum(deviceCustomNumber1)`, the `sum_deviceCustomNumber1` column in the Results Table displays the sum of unique values of the `deviceCustomNumber1` field. If this field had two unique values 1 and 20, occurring 2 times each, the `sum_deviceCustomNumber1` column displays sum of those two values, as shown in the following figure:

Result Table	
Page 1 of 1	
deviceCustomNumber1	sum_deviceCustomNumber1
1	2
20	40

Aggregation functions can only be used on numeric fields.

The mathematical operators `avg` and `mean` are identical.

You can include multiple functions in the same `chart` command. When doing so, separate each function with a comma, as shown in this example:

```
... | chart count, sum(deviceCustomNumber3) by deviceEventClassId
```

When you include multiple functions, one column per function is displayed in the search Results Table. The Results Chart, however, plots the chart for the field specified in the “by” clause.

You can use the “as new_column_name” clause to name any column resulting from the aggregation functions, as shown in this example:

```
... | chart sum(deviceCustomNumber3) as TotalStorage,
avg(deviceCustomNumber3) as AverageStorage by deviceCustomNumber3
```

Once defined, the newly defined column can be used in the pipeline as any other field. For example,

```
... | chart sum(deviceCustomNumber3) as TotalStorage,
avg(deviceCustomNumber3) as AverageStorage by deviceCustomNumber3 |
eval UpdatedStorage = TotalStorage + 100
```

When you export the search results of a chart operator, the newly defined column name (using the `chart` function as `new_column_name` command) is preserved.

Multi-Series Charts

A multi-series chart can plot the values of multiple aggregation functions in a single chart.

If you include multiple aggregation functions in a `chart` command, OLI generates a multi-series chart that plots the values of the specified aggregation functions along the Y-axis, as illustrated in “[Example 2](#)” on page 274. Multi-series charts can be any of the chart types except Pie charts. For example, you can choose to plot a multi-series chart as a stacked chart—Stacked column or Stacked Bar—in which multiple values are plotted in a stack form, as illustrated in “[Example 3](#)” on page 275.


The span function

In addition to grouping events by the OLI schema fields (or the ones defined by the `rex` or `eval` operators), the `span` function provides an additional way to group events by a time field (such as `EventTime` or `deviceReceiptTime`) and a time bucket. In the following example, `deviceReceiptTime` is the time field and `5m` (5 minutes) is the time bucket:

```
...| chart count by deviceEventCategory span (deviceReceiptTime) = 5m
```

If a time field is not specified for the `span` function, `EventTime` is used as the default. For example, the following query uses `EventTime` by default:

```
...| chart count by deviceEventCategory span = 5m
```

By default, the `chart` command displays the first 10 unique values. If the `span` function creates more than 10 unique groups, not all of them will be displayed. If you want to view all of the unique groups, increase the Display Limit value under Chart Settings. (Click  to the upper right corner of the Result Chart frame of the screen.)

Grouping with `span` is useful in situations when you want to find out the number of occurrences in a specific time span.

If you want to find out the total number of incoming bytes every 5 minutes on a device, you can specify a span of `5m`, as shown in this example:

```
...| chart sum(deviceCustomNumber1) span=5m
```

The above example assumes that `deviceCustomNumber1` field provides the incoming bytes information for these events.

The `span` field can be used for grouping in conjunction with or without the event fields that exist in OLI schema or user-defined fields using the `rex` or `eval` operators. When a `span` field is specified in conjunction with an event field, the unique sets of all those fields is used for grouping. The following example uses `deviceCustomNumber3` and `deviceAddress` in conjunction with `span` to find out the number of events (using `deviceCustomNumber3`) from a specific source (using `deviceAddress`) in one hour:

```
...| chart sum (deviceCustomNumber3) by deviceAddress span=1h
```

When `span` is included in a query, search results are grouped by the specified time bucket. For example, if `span=5m`, the search results will contain one row for

each 5-minute span. If there are no events within a specific 5-minute span, that row will be empty.

Additionally, the `span` function assumes a 24-hour day, all year long. If `span=1d` or `24h`, on the day of daylight savings time change, the event time indicated by the `span_eventTime` field in the search results will be different from the previous day by one hour. On the day when there are 23 hours in a day (in March), the span bucket will still include events from the last 24 hours. Similarly, on the day when there are 25 hours in the day (in November), the span bucket will include events from the last 24 hours. The following example illustrates the `span_eventTime` field when the span time bucket is `1d` and the daylight savings times occurs on

March 14th, 2011 and November 7, 2011:

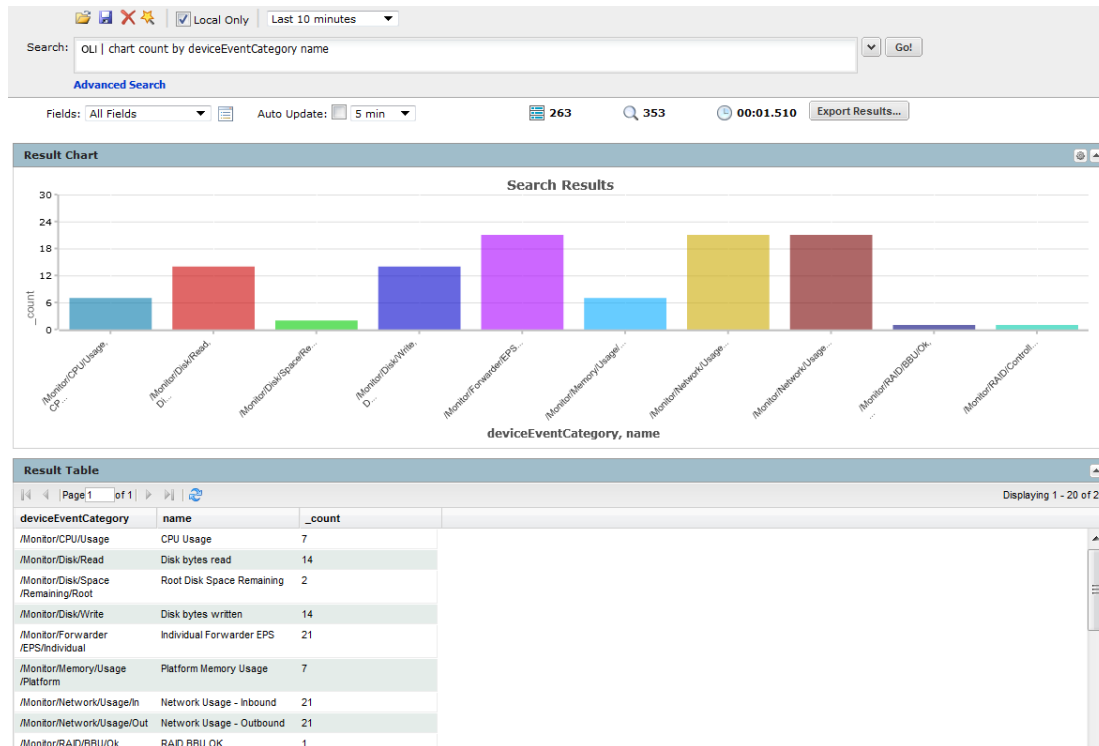
<code>span_eventTime</code>	<code>avg_logins</code>
3/11/2011 12am	8
3/12/2011 12am	10
3/13/2011 12am	4
3/14/2011 1am	6
3/15/2011 1am	7
....	
11/5/2011 1am	4
11/6/2011 1am	2
11/7/2011 12am	5
11/8/2011 12am	7
....	

Example 1

Use the default chart setting (Column Chart) to specify multiple fields. In this example, a count of unique groups of `deviceEventCategory` and `name` fields is displayed and plotted.

A Search Operators

... | chart count by deviceEventCategory name

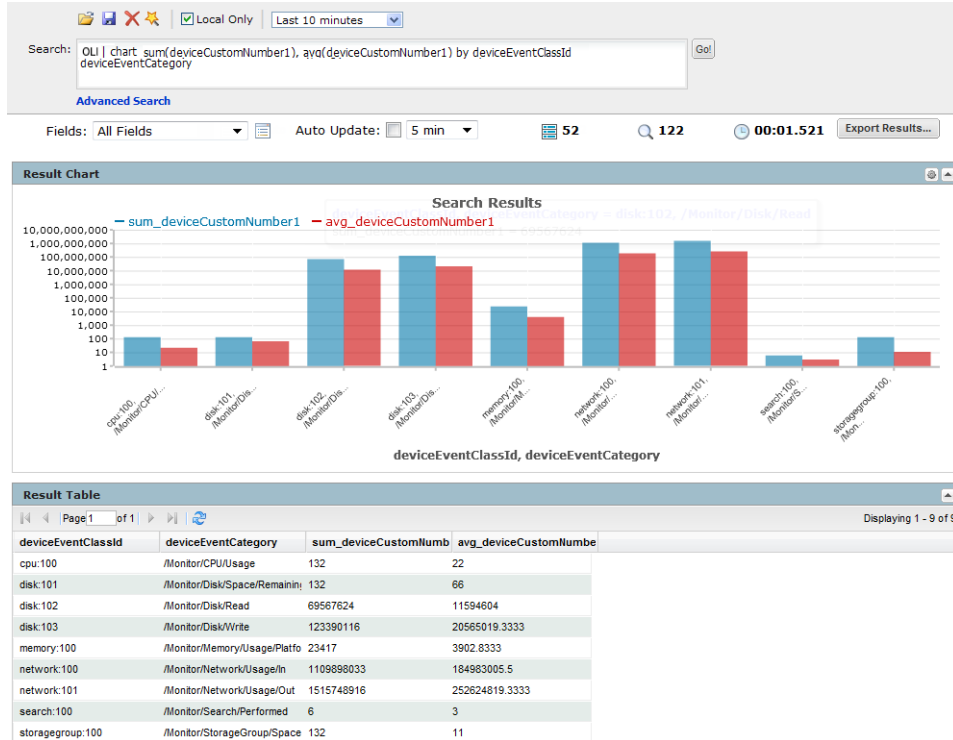


Example 2

Include average and sum in a chart command, to generate a multi-series chart that plots the values of these functions along the Y-axis in a single chart.

In the following query, unique groups of deviceEventClassId and deviceEventCategory are plotted along the X-axis, and the sum of deviceCustomNumber1 and average of deviceCustomNumber2 is plotted along the Y-axis.

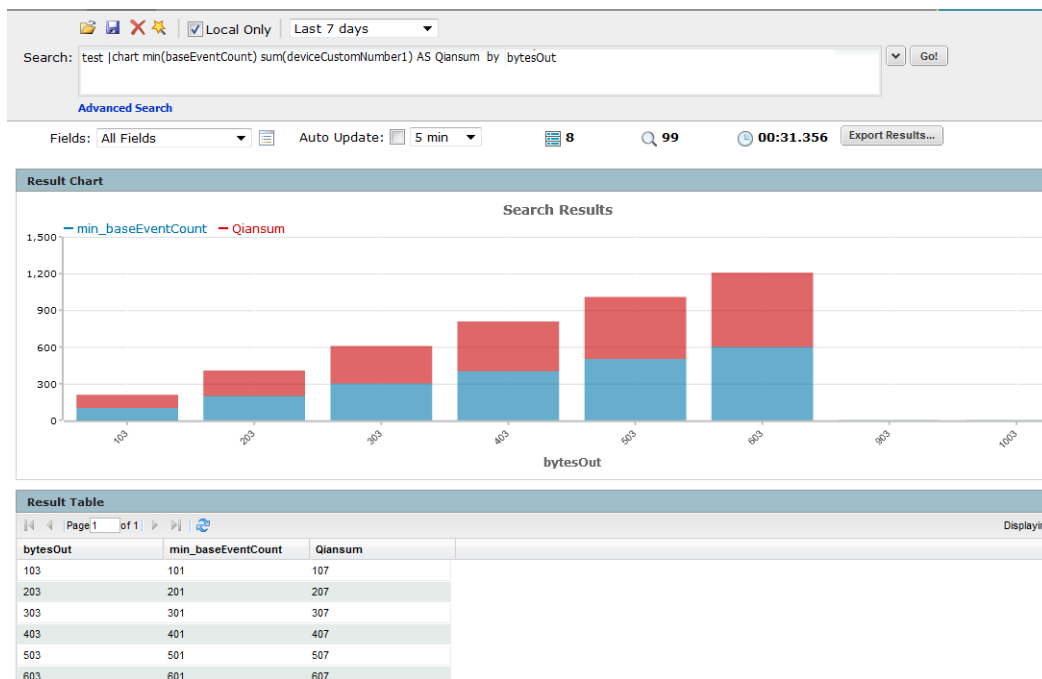
... | chart sum(deviceCustomNumber1), avg(deviceCustomNumber1) by deviceEventClassId deviceEventCategory



Example 3

Plot a multi-series chart as a stacked chart—Stacked column or Stacked Bar—in which multiple values are plotted in a stack form, as shown in the following figure.

... | chart min(baseEventCount) sum(deviceCustomerNumber1) AS Qiansum by bytesOut



dedup

Removes duplicate events from search results. That is, events that contain the same value in the specified field. The first matching event is kept, and the subsequent events with the same value in the specified field are removed.

Usage

```
... | dedup [N] <field1>,<field2>, ... [keepevents=(true|false)]
[keepempty=(true|false)]
```

`N` is an optional number that specifies the number of duplicate events to keep. For example, “dedup 5 deviceEventClassId” will keep the first five events containing the same deviceEventClassId values for each deviceEventClassId, and remove the events that match after the first five have been kept. Default: 1.

`field1`, `field2` is a field or a comma-separated field list whose values are compared to determine duplicate events. If a field list is specified, the values of the unique sets of all those fields are used to remove events. For example, if `name` and `deviceCustomNumber1` are specified, and two events contain “Network Usage - Outbound” and “2347896”, only the first event is kept in the search results.

`keepevents` specifies whether to set the fields specified in the field list to NULL or not. When this option is set to True, the values are set to NULL and events are not removed from search results. However, when this option is set to False, duplicate events are removed from the search results. Default: False.

`keepempty` specifies whether to keep events in the search results whose specified fields contain NULL values. When this option is set to True, events with NULL values are kept, however if this option is set to False, events with NULL values are removed. Default: False.

Example 1

To view events from unique devices:

```
... | dedup deviceAddress
```

Example 2

To view unique deviceEventClassId events from unique devices:

```
... | dedup deviceEventClassId deviceAddress
```

Example 3

To view the className in events with Java exceptions in the message field:

```
exception | <rex_expression> | dedup 5 className
```

In the above example, rex expression is not shown in detail however this expression extracts the class name in a field called className, which the dedup operator acts upon.

eval

Displays events that match the resultant of the specified expression. The expression can be a mathematical, string, or boolean operation and is evaluated when the query is run. The resulting value of the expression is assigned to a field name (as specified in the expression). Once a new field has been defined by the eval operator in a query, this field can be used in the query for further refining the search results (see Example #3 below, in which a new field “Plus” is defined by the eval operator; this field is then used by the sort operator.)

Usage

```
... | eval <expression>
```

<expression> is a mathematical, string, or boolean operation; for example, `total_bytes=bytesIn + bytesOut`.

Notes

Typically, a `cef` or `rex` operator (to extract fields from matching events) precedes the `eval` operator, as shown in the examples below. However, you can use the `eval` operator on a field that has been defined by a previous `eval` operator in a query.

Example 1

If the Category Behavior is “Communicate”, then assign the value “communicate” to a new field “cat”; otherwise, assign the value “notCommunicate” to it.

```
_storageGroup IN ["Default Storage Group"] | cef categoryBehavior |
eval cat=if(categoryBehavior== "/Communicate", "communicate",
"notCommunicate")
```

Example 2

Append the word, “END”, at the end of extracted event name. For example, if event name is “Operations Log Intelligence Internal Event”, after the eval operation it is “Operations Log Intelligence Internal EventEND” and is assigned to a new field, “fullname”.

```
logger | cef msg name | eval fullname=name + "END"
```

Example 3

Add 100 to the value of bytesIn and assign it to a new field, “Plus”. Then, sort the values assigned to “Plus” in ascending order.

```
_storageGroup IN ["Default Storage Group"] | cef bytesIn bytesOut
name | eval Plus=bytesIn +100 | sort Plus
```

extract

Extracts key value pairs from raw events.

Usage

```
... | extract [pairedlim="<delimiters>"] [kvdelim="<delimiters>"]
[maxchars=<n>] fields="key1,key2,key3..."
```

`pairdelim` is a delimiter (or a list of delimiters) that separates one key-value pair from another key-value pair in an event. By default, semi colon, pipe, and comma (; | ,) are used.

`kvdelim` is a delimiter (or a list of delimiters) that separates a key from its value. By default, "=".

`maxchars` is the maximum number of characters in an event that would be scanned for extracting key value pairs. By default, 10240.

`fields` is a key (or a list of comma-separated keys) whose values you want to display in the search results. For example, if you want to display the Name Age, and Location values from this event:

Name:Jane | Age:30 | Location:LA

Then, extract the "Name", "Age", and "Location" keys and list them in the `fields` list.

Understanding how the operator works:

The key represents a field in the raw event and its value consists of the characters that appear after the key until the next key in the event. The following raw event is used to illustrate the concept:

```
[Thu Jul 30 01:20:06 2009] [error] [client 69.63.180.245] PHP Warning:
memcache_pconnect() [<a
href='function.memcache-pconnect'>function.memcache-pconnect</a>]: Can't
connect to 10.4.31.4:11211
```

To extract the URL from the above event, you can define these key-pair delimiters, which separate the key-value pairs in the event:

- Greater than sign (">")
- Square bracket ("["

And, define this key delimiter, which separates the key from its value:

- Equal to sign ("=")

Thus, the following command will extract the URL

```
... | extract pairdelim= ">\" kvdelim= "=" fields="<a href"
```

The key value pairs in the event will be: [

The key in the event will be: <a href

The extracted URL will be: 'function.memcache-pconnect'

Notes

This operator only works on raw events. That is, you cannot extract key value pairs from CEF events or the fields defined by the `rex` operator.

You can specify the `pairdelim` and `kvdelim` delimiters in the `extract` operator command to extract keys and their values. However, if you want to determine the key names that these delimiters will generate, use the `keys` operator as described in [“keys” on page 280](#). The `keys` operator can only be used to determine keys; you cannot pipe those keys in the `extract` operator. That is, `...| keys | extract fields=field1` is incorrect.

The keys specified in the fields list can be used further in the pipeline operations. For example, `...| extract pairdelim= "|" kvdelim= ":" fields= "count" | top count`

If none of the specified `pairdelim` characters exist in an event, the event is not parsed into key value pairs. The whole event is skipped. Similarly, if the specified `kvdelim` does not exist, values are not separated from the keys.

To specify double quotes (") as the delimiter, enter it within the pair of double quotes with backslash(\) as the escape character. For example, `"\"`. Similarly, use two backslashes to treat a backslash character literally. For example, `\"`.

Example

```
... | extract pairdelim= "|" kvdelim= ":" fields=
>Name, Age, Location"
```

Extracts values from events in this format:

```
Name:Jane | Age:30 | Location:LA
```

fields

Includes or excludes specified fields from search results.

Usage

```
... | fields [(+ | -)] <field>+
```

+ includes only the specified field or fields in the search results. This is the default.

- excludes only the specified field or fields from the search results.

Notes

Typically, the `<field>` list contains event fields available in the OLI schema or user-defined fields created using the `rex` operator prior in the query, as shown in the examples below. However, fields might also be defined by other operators such as the `eval` operator.

The + and - can be used in the same expression when multiple fields are specified. For example, `| fields + name - agentType`

A complete field name must be specified for this operator; wildcard characters in a field name are not supported.

When this operator is included in a query, select **User Defined Fieldsets** from the System Fieldsets list to view the search results.

Example 1

```
... | fields - agentType + categorySignificance
```

Example 2

```
... | fields - name
```

head

Displays the first <N> lines of the search results.

Usage

```
... | head [<N>]
```

<N> is the number of lines to display. Default: 10, if <N> is not specified.

Notes

When this operator is included in a query, the search results are not previewable. That is, the query must finish running before search results are displayed.

Example

```
... | head
```

keys

Identifies keys in raw events based on the specified delimiters.

Usage

```
... | keys [pairdelim= "<delimiters>"] [kvdelim= "<delimiters>"]
[limit=<n>]
```

`pairdelim` is a delimiter (or a list of delimiters) that separates one key-value pair from another key-value pair in an event. By default, semi colon, pipe, and comma (; |,) are used.

`kvdelim` is a delimiter (or a list of delimiters) that separates a key from its value. By default, "=".

`limit` is the maximum number of key value pairs to find. There is no default or maximum number for this parameter.

Notes

This operator only works on raw events. That is, you cannot identify key value pairs from CEF events or fields defined by the `rex` operator.

Although this operator is not required to determine keys, it is recommended that you use it to first determine the keys whose values you want to obtain using the `extract` operator. This operator returns aggregated results. Therefore, the search results list the keys found in the matching events and their counts.

The `keys` operator can only be used to determine keys; you cannot pipe those keys in the `extract` operator. That is, `| keys | extract fields=field1` is incorrect.

If a key value is blank (or null), it is ignored and not counted toward the number of hits.

For example, for the following event data:

```
Date=3/24/2011 | Drink=Lemonade
Date=3/23/2011 | Drink=
Date=3/22/2011 | Drink=Coffee
```

Search Query: keys pairedelim="|" kvdelim="="

Search Result: Date, 3 hits and Drink, 2 hits

If none of the specified `pairedelim` characters exist in an event, the event is not parsed into key value pairs. The whole event is skipped. Similarly, if the specified `kvdelim` does not exist, values are not separated from the keys.

To specify double quotes (") as the delimiter, enter it within the pair of double quotes with backslash(\) as the escape character. For example, "\"". Similarly, use two backslashes to treat a backslash character literally. For example, "\\".

Example 1

```
... | keys pairedelim="|" kvdelim="="
```

Identifies keys (Date and Drink) in event of this format:

```
Date=3/24/2011 | Drink=Lemonade.
```

Example 2

```
... | keys pairedelim="," kvdelim=">="
```

Identifies keys (Path and IPAddress) in the event of this format:

```
Path>c:\usr\log, IPAddress=1.1.1.1
```

parse

Applies the named parser to the matching events of a search query. The parser definition for the specified parser name must exist before it can be used in a query.

Usage

```
... | parse <parser_name>
```

`<parser_name>` is the name of the parser to use.

The `parse` operator is useful in parsing the non-CEF (unstructured textual) data stored on OLI and parsing it into specific fields according to the parser's definition.

Once parsed into fields, this data can be used further in search operations. For example, the following `parse` operator parses the events using a user-defined parser "Web Server Access Logs" such that "username", "login_status", "num_attempts" fields are created. You can use these created fields further in a pipeline query to display the top 10 user names that resulted in the maximum failed login attempts and the number of attempts they made.

```
... | parse Web Server Access Log | where login_status = "failed" |
top username num_attempts
```

Because the parser definitions are `rex` or `extract` expressions, they create additional fields to contain values that match the specified expression. These fields are displayed in the Search Results just like the results of any `rex` or `extract` expression. Therefore, in the above example, three additional fields will be added to the Search Results—`username`, `login_status`, `num_attempts`.

An additional field called “parser” is also added to the Search Results when the `parse` operator is used in a search query.

This field contains the name of the parser when the parser is able to parse one or more fields specified in the definition for the matching events. If the event was not parsed successfully, if no parser is defined for the source type, or if there is no source type, this field displays, this field contains “Not parsed”. Similarly, the field contains the value “not parsed” when the parser definition is not able to parse any fields of the matching event.

You can also use this field to find out events that were successfully parsed or did not parse, as shown in the following example:

```
... | parse Apache Access Log | where parser = "not parsed"
```

When to use the parse operator: When non-CEF events are received through TCP or UDP receivers on OLI, they are not associated with a source type and thus a parser definition. Therefore, such events not parsed automatically. If you need such events parsed when they match a query, use the `parse` operator.

When an event for which a defined source type exists on OLI is parsed through the `parse` operator, it can result in the creation of multiple user-defined fields—through the parser associated with the source type and through the parser you specified in the `parser` pipeline command. If both parsers create unique field names, all those fields are created when a query that matches the event is run. If the parsers specify one or more same name fields, the field names specified in the `parse` operator parser take precedence as this parser is applied last.

Example:

```
... | parse Web Server Access Log | where url CONTAINS ".org" | top
url
```

rare

Lists the search results in a tabular form of the least common values for the specified field. That is, the values are listed from the lowest count value to the highest.

When multiple fields are specified, the count of unique sets of all those fields is listed from the lowest to highest count.

Usage

```
... | rare <field1> <field2> <field3> ...
```

Notes

Typically, the <field> list contains event fields available in the OLI schema or user-defined fields created using the `rex` or `eval` operators prior in the query, as shown in the examples below. However, fields might also be defined by other operators such as the `eval` operator.

A chart of the search results is automatically generated when this operator is included in a query. You can click on a charted value to quickly filter down to events with specific field values. For more information, see [“Chart Drill Down” on page 104](#).

If multiple fields are specified, separate the field names with a white space or a comma.

Example

```
... | rare deviceEventCategory
```

regex

Selects events that match the specified regular expression.

Usage

```
... | regex <regular_expression>
```

OR

```
... | regex <field> (=|!=) <regular_expression>
```

Notes

Regular expression pattern matching is case insensitive.

The first usage (without a field name) is applied to the raw event. While the second usage (with a field name), is applied to a specific field.

If you use the second usage (as shown above and in the Example #2 below), either specify an event field that is available in the OLI schema or a user-defined field created using the `rex` or `eval` operators.

Example 1

```
... | regex "failure"
```

Example 2

```
... | regex deviceEventCategory != "fan"
```

rename

Renames the specified field name.

Usage

```
... | rename <field> as <new_name>
```

<field> is the name of an event field that is available in the OLI schema or a user-defined field created using the `rex` or `eval` operator.

<new_name> is the new name you want to assign to the field.

Notes

An additional column is added to the search results for each renamed field. The field with the original name continues to be displayed in the search results in addition to the renamed field. For example, if you rename `deviceEventCategory` to `Category`, two columns are displayed in the search results: `deviceEventCategory` and `Category`.

You can include the wildcard character, `*`, in a field name. However, you must enclose the field that contains a wildcard character in double quotes (`" "`). For example:

```
...| rename "*IPAddress" as "*Address"
```

OR

```
...| rename "*IPAddress" as Address
```

If a field name includes a special character (such as `_`, a space, `#`, and so on), it should be included in double quotes (`" "`) in the `rename` operator expression. For example:

```
...| rename src_ip as "Source IP Address"
```

If the resulting field of a `rename` operation includes a special character, it must be enclosed in double quotes (`" "`) whenever you use it in the pipeline operator expression. For example,

```
...| rename src_ip as "Source IP Address" | top "Source IP Address"
```

The internal field names (that start with `"_raw"`) cannot be renamed.

The renamed fields are valid only for the duration of the query.

The resulting field of a `rename` operation is case sensitive. When using such a field in a search operation, make sure that you use the same case that was used to define the field.

When you export the search results of a search query that contains the `rename` expression, the resulting file contains the renamed fields.

Example 1

```
...| rename src_ip as IPAddress
```

Example 2

```
...| rename src_ip as "Source IP Address"
```

replace

Replaces the specified string in the specified fields with the specified new string.

Usage

`<orig_str> with <new_str> [in <field_list>]`

`<orig_str>` is the original string you want to replace. (See Notes for more details.)

`<new_str>` is the new string you want to replace with. (See Notes for more details.)

`<field_list>` is the optional, however highly recommended. See Notes for details.

Notes

Even though the field list is optional for this command, HP strongly recommends that you specify the fields on which the `replace` operator should act in this command.

If you skip the field list, the `replace` operator acts on the fields that have been either explicitly defined using the `cef`, `rex`, and `eval` operators preceding the `replace` command, or any fields that were used in other operator commands that preceded the `replace` operator command. For example, the `replace` command acts on `deviceEventCategory` in all of the following cases and replaces all instances of “EPS” with “Events”:

```
...| replace *EPS* with *Events* in deviceEventCategory
...| cef deviceEventCategory | replace *EPS* with *Events*
...| top deviceEventCategory | replace *EPS* with *Events*
```

An additional column of the same name is added to the search results for each field in which string is replaced. The column with the original value continues to be displayed in the search results in addition to the column with replaced values. For example, if you replace “err” with “Error” in the “message” column, an additional “message” column is added to the search results that contains the modified value.

If you want to replace the entire string, specify it in full (as it appears in the event). For example, “192.168.35.3”.

If you want to replace a part of the string, include wildcard character (*) for the part that is not going to change.

For example, if the original string (the string you want to replace) is “192.168*”, only the 192.168 part in an event is replaced. The remaining string is preserved. As a result, if an event contains 192.168.35.3, only the first two bytes are replaced. The rest (35.3) will be preserved. Similarly, if the event contains 192.168.DestIP, DestIP will be preserved. However, if the event contains the string 192.168, it will not be replaced.

If both, the original and the new strings contain wildcard characters, the number of wildcard characters in the *original* string must match the number of wildcard characters in the *new* string.

```
...| replace "*.168.*" with "*.XXX.*"
```

If the original or the new string includes a special character such as / or ?, enclose the string in double quotes (" "):

```
...| replace "/Monitor" with Error
```

You can replace multiple values for multiple fields in a single operation by separating each expression with a comma (.). Note that you must specify the field list after specifying the "with" expression for all values you want to replace, as shown in the following example:

```
...| replace "OLI*" with HP, "cpu:100" with EPS in deviceVendor,
deviceEventClassId
```

The original string is case-insensitive. Therefore, the string "err" will replace an event that contains "Err".

Example 1

Replace any occurrence of "a" with "b" but the characters preceding "a" and succeeding it are preserved.

```
...| replace *a* with *b*
```

Example 2

Replace any occurrence of "a" with "b" without retaining any characters preceding or succeeding "a".

```
...| replace *a* with b in name
```

rex

Extracts (or capture) a value based on the specified regular expression or extract and substitute a value based on the specified "sed" expression. The value can be from a previously specified field in the query or a raw event message.

Usage

```
... | rex <regular_expression containing a field name>
```

OR

```
... | rex field = <field> mode=sed "s/<string to be
substituted>/<substitution value>"
```

Understanding how extraction works:

When the value is extracted based on a regular expression, the extracted value is assigned to a field name, which is specified as part of the regular expression. The syntax for defining the field name is **?<fieldname>**, where *fieldname* is a string of alphanumeric characters. Using an underscore ("_") is not recommended.

We use the following event to illustrate the power of rex.

```
[Thu Jul 30 01:20:06 2009] [error] [client 69.63.180.245] PHP Warning: Can't
connect to 10.4.31.4:11211
```

If you want to extract any IP address from the above event and assign it to a field called "IP_Address", you can simply specify the following rex expression:

```
| rex "(?<IPAddress>\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
```

However, if you wanted to extract the IP address after the word "client" from the following event and assign it to a field called "SourceIP", you will need to specify a start and end point for IP address extraction so that the second IP address in the event is not captured. The starting point in this event can be "[client]" and the end point can be "]". Thus, the rex expression will be:

```
| rex "\[client (?<SourceIP>[^\]]*)"
```

In this rex expression **?<SourceIP>** is the field name defined to capture IP address and "client" specifies the text or point in the event AFTER which data will be extracted. The `[^\]]*` expression will match every character that is not a closing right bracket, therefore, for our example event, the expression will match until the end of the first IP address and not the second IP address that appears after the word "to".

Understanding how substitution works:

When the `rex` operator is used in `sed` mode, you can substitute the values of extracted fields with the values you specify.

The substitution only occurs in the search results. The actual event is not changed.

In the following example, the credit card numbers in the CCN field are substituted with "xxx", thus obfuscating sensitive data:

```
| rex field=CCN mode=sed "s/*/XXXX/g"
```

The `/g` at the end of the command indicates a global replace, that is, all occurrences of the specified pattern will be replaced in all matching events. If `/g` is omitted, only the first occurrence of the specified pattern in each event is replaced.

Multiple substitutions can be performed in a single command, as shown in the following example. In this example, the word "Authentication" is substituted with "xxx" globally (for all matching events), the first byte of the agent address that start with "192" is substituted with "xxx" and an IP address that starts with "10" is substituted with "xxx".

```
| rex field=msg mode=sed "s/Authentication/xxx/g" | rex field=agentAddress mode=sed "s/192/xxx/g" | rex field=dst mode=sed "s/10./xxx/g"
```

Notes

A detailed tutorial on the `rex` operator is available at [Appendix D, , on page 325](#).

A Regex Helper tool is available for formulating regular expressions of fields in which you are interested. The Regex Helper parses an event into fields. Then, you select the fields that you want to include in the rex expression. The regular expression for those fields is automatically inserted in the Search box. For

detailed information on the Regex Helper tool, see [“Regex Helper Tool” on page 91](#).

The extracted values are displayed as additional columns in the All Fields view (of the System FieldSets). To view only the extracted columns, select **User Defined Fieldsets** from the System Fieldsets list. In the above example, an additional column with heading “SourceIP” is added to the All Fields view; IP address values extracted from events are listed in this column.

If you want to use other search operators such as fields, sort, chart, and so on to refine your search results, you must first use this operator to extract those fields.

Example 1

The following example extracts name and social security number from an event that contains data in name:John ssn:123-45-6789 format and assigns them to Name and SSN fields:

```
... | rex "name: (?<Name>.* ) ssn: (?<SSN>.* )"
```

Example 2

The following example extracts URLs from events and displays the top 10 of the extracted URLs:

```
... | rex "http://(?<URL>[^\ ]*)" | top URL
```

Example 3

The following example substitutes the last four digits of social security numbers extracted in the first event with XXXX:

```
... | rex field=SSN mode=sed "s/-\d{4}/-XXXX/g"
```

sort

Sorts search results as specified by the sort criteria.

Usage

```
... | sort [<N>] ((+ | -) field)+
```

+ Sort the results by specified fields in ascending order. This is the default.

- Sort the results by specified fields in descending order.

<N> Keep the top N results, where N can be a number between 1 and 10,000. Default: 10,000.

Notes

Typically, the <field> list contains event fields available in the OLI schema or user-defined fields created using the `rex` operator prior in the query, as shown in the examples below. However, fields might also be defined by other operators such as the `eval` operator.

Sorting is based on the data type of the specified field.

When multiple fields are specified for a sort operation, the first field is used to sort the data. If there are multiple same values after the first sort, the second field is used to sort within the same values, followed by third field, and so on. For example, in the example below, first the matching events are sorted by “cat” (device event category). If multiple events have the same “cat”, those events are further sorted by “eventId”.

When multiple fields are specified, you can specify a different sort order for each field. For example, `| sort + deviceEventCategory - eventId`.

If multiple fields are specified, separate the field names with a white space or a comma.

Sorting is case-sensitive. Therefore, “Error:105” will precede “error:105” in the sorted list (when sorted in ascending order).

When a sort operator is included in a query, only the top 10,000 matches are displayed. This is a known limitation and will be addressed in a future OLI release.

When this operator is included in a query, the search results are not previewable. That is, the query must finish running before search results are displayed.

Example

```
... | sort deviceEventCategory eventId
```

tail

Displays the last <N> lines of the search results.

Usage

```
... | tail [<N>]
```

<N> is the number of lines to display. Default: 10, if <N> is not specified.

Notes

When this operator is included in a query, the search results are not previewable. That is, the query must finish running before search results are displayed.

Example

```
... | tail 5
```

top

Lists the search results in a tabular form of the most common values for the specified field. That is, the values are listed from the highest count value to the lowest.

Usage

```
... | top [<n>] <field1> <field2> <field3> ...
```

<n> limits the matches to the top *n* values for the specified fields. Default: 10, if <N> is not specified.

Notes

The fields can be either event fields available in the OLI schema or user-defined fields created using the `rex` or `eval` operators prior in the query. If multiple fields are specified, separate the field names with a white space or a comma.

When multiple fields are specified, the count of unique sets of all those fields is listed from the highest to lowest count.

A chart of the search results is automatically generated when this operator is included in a query. You can click on a charted value to quickly filter down to events with specific field values. For more information, see [“Chart Drill Down” on page 104](#).

To limit the matches to the top *n* values for the specified fields, specify a value for *n*. For example, `... | top 5 deviceEventCategory`

Example 1

```
... | top deviceEventCategory
```

Example 2

```
... | top 5 categories
```

transaction

Groups events that have the same values in the specified fields.

Usage

```
... | transaction <field1> <field2>... [maxevents=<number>]
[maxspan=<number> [s|m|h|d]] [maxpause=<number> [s|m|h|d]]
[startswith=<reg_exp>] [endswith=<reg_exp>]
```

`field1`, `field2` is a field or a comma-separated field list whose values are compared to determine events to group. If a field list is specified, the values of the unique sets of all those fields are used to determine events to group. For example, if `host` and `portNum` are specified, and two events contain “hostA” and “8080”, the events are grouped in a transaction.

`maxevents` specifies the maximum number of events that can be part of a single transaction. For example, if you specify 5, after 5 matching events have been found, additional events are not included in the transaction. Default: 1000

`maxspan` specifies the limit on the duration of the transaction. That is, the difference in time between the first event and all other events in a transaction will never be more than the specified `maxspan` limit. For example, if you specify `maxspan=30s`, the event time of all events within the transaction will be at most 30 seconds more than the event time of the first event in the transaction. Default: Unlimited

`maxpause` specifies the length of time by which consecutive events in a transaction can be apart. That is, this option ensures that events in a single transaction are never more than the `maxpause` value from the previous event in the transaction. Default: Unlimited

`startswith` specifies a regular expression that is used to recognize the beginning of a transaction. For example, if a transaction operator includes `startswith= "user [L|l]login"`, all events are scanned for this regular expression. When an event matches the regular expression, a transaction is created, and subsequent events with matching fields are added to the transaction.



The regular expression is applied to the raw event, not to a field in an event.

`endswith` specifies a regular expression that is used to recognize the end of an existing transaction. That is, an existing transaction is completed when an event matches the specified "endswith" regular expression. For example, if a transaction operator includes `endswith= "[L|l]logout"`, any event being added to a transaction is checked, and if the regular expression matches the event, the transaction is completed.



The regular expression is applied to the raw event, not to a field in an event.

Notes

Several of the above options specify "conditions to end" a transaction. Therefore, when multiple "end conditions" are specified in a transaction operator, the first end condition that occurs will end the transaction even if the other conditions have not been satisfied yet. For example, if `maxspan` is reached but `maxevents` has not been reached, or if the `endswith` regular expression is matched but `maxevents` has not been reached.

Understanding how the transaction operator works:

A transaction is a set of events that contain the same values in the specified fields. The events may be further filtered based on the options described above, such as `maxspan`, `maxpause`, and so on. In addition to grouping events, the transaction operator adds these fields to each event: `transactionid`, `duration`, and `eventcount`. These fields are displayed in the Search Results as separate columns.

A `transactionid` is assigned to each transaction when the transaction completes. Transaction IDs are integers, assigned starting from 1 for the transactions (set of events) found in the current query. All events in the same transaction will have the same transaction ID.

If an event does not belong to any transaction found in the current query, it is assigned the transaction ID 0. For example, in a `transaction` operator with a `startswith` regular expression, if the first event in the pipeline does not match the regular expression, that event is not part of the transaction, and is assigned transaction ID 0.

The `duration` is the time in milliseconds of the duration of a transaction, which is the difference between the event time of the last event in the transaction and the

first event in the transaction. The duration field for all events in a transaction is set to the duration value of the transaction.

The eventcount displays the number of events in a transaction.

Example 1

To view source addresses accessed within a 5-minute duration:

```
... | transaction sourceAddress maxspan=5m
```

Example 2

To group source addresses by source ports and view 5 events per group:

```
... | transaction sourceAddress sourcePort maxevents=5
```

Example 3

To group users and URLs they accessed within a 10-minute duration:

```
... | transaction username startswith= "http://" maxspan=10m
```

Example 4

To view login transactions from the same session ID and source address in a 1-hour duration:

```
... | transaction sessionID sourceAddress maxspan=1h startswith=
"user [L|l]ogin"
```

where

Displays events that match the criteria specified in the “where” expression.

Usage

```
... | where <expression>
```

<expression> can be any valid field-based query expression, as described in [“Field-based Search” on page 70](#).

Notes

<expression> can only be a valid field-based query expression. Arithmetic expressions or functions are not supported.

Example 1

```
... | where eventId is NULL
```

Example 2

```
... | where eventId=10006093313 OR deviceVersion CONTAINS
"4.0.6.4924.1"
```

Example 3

```
... | where eventId >=10005985569 OR categories= "/Agent/Started"
```

Appendix B: Audit Events

This appendix describes OLI's audit events in detail. It includes information on the following topics:

[“Types of Audit Events” on page 293](#)

[“Information in an Audit Event” on page 293](#)

[“Platform Events” on page 293](#)

[“OLI Application Events” on page 301](#)

Types of Audit Events

Two types of audit events are generated on OLI

- [“Platform Events” on page 293](#)—related to the OLI system
- [“OLI Application Events” on page 301](#)—related to OLI functions and configuration changes on it

Both types of events are stored in the OLI Internal Storage Group. As a result, these events can be searched using the OLI Search UI. For example, you can search for this platform event:

```
“/Platform/Authentication/Failure/Password”
```

Information in an Audit Event

An OLI audit event (in CEF format) contains information about the following prefix fields:

- Device Event Class ID
- Device Severity
- Message
- Device Event Category—(keyName for this CEF extension is “cat”)

For example:

```
Sep 19 08:26:10 zurich CEF:0|HP|OLI|3.5.0.13412.0|oli:500|Filter  
added|2| cat=/OLI/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 OLI platform. All events include the following fields.

B Audit Events

- duser—UserName
- duid—User ID
- src—IP address of client
- dst—IP address of appliance
- cat—Device Event Category
- cn1—Session number
- cn1label—Session

Additional fields (if applicable) are listed in the following table.

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform:200	5	/Platform/Authentication/PasswordChange/Failure	Failed password change	
platform:201	7	/Platform/Authentication/Failure	Failed login attempt	
platform:202	5	/Platform/Authentication/PasswordChange	Password changed	cs1: Affected User Id cs2: Affected User Login cs3: Affected User Full Name
platform:203	7	/Platform/Authentication/InactiveUser/Failure	Login attempt by inactive user	
platform:204	5	/Platform/Authentication/ScheduledHours/Failure	Login attempt by user outside scheduled work hours	
platform:220	5	/Platform/Certificate/Install	Installed certificate	cs1: Network Protocol
platform:221	7	/Platform/Certificate/Mismatch	Certificate mismatch failure	cs1: Network Protocol
platform:222	1	/Platform/Certificate/Request	Created certificate signing request	cs1: Certificate Signing Request cs2: Network Protocol
platform:224	5	/Platform/Certificate/Regenerate	Re-generate self-signed certificate	cs1: Certificate Signing Request cs2: Network Protocol
platform:226	7	/Platform/Update/Failure/CorruptPackage	Uploaded update file damaged or corrupt	cs1: Error cs2: fname cs3: fsize
platform:227	5	/Platform/Update/Applied	Update installation success	cs1: Update Name cs2: Is Reboot Required
platform:228	7	/Platform/Update/Failure/Installation	Update installation failure	cs1: Error cs2: Update Name

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform:230	3	/Platform/Authentication/Login	Successful login	
platform:234	7	/Platform/Authentication/Failure/LOCKED	Failed login attempt (LOCKED)	
platform:239	3	/Platform/Authentication/Logout	User logout	
platform:240	3	/Platform/Authorization/Groups/Add	Added user group	cn2: Current Number of Users cn3: Current Number of User Rights cs1: Affected Group Name cs2: Affected Group Id flexNumber1: Old Number of Users flexNumber2: Old Number of User Rights
platform:241	3	/Platform/Authorization/Groups/Update	Updated user group	cn2: Current Number of Users cn3: Current Number of User Rights cs1: Affected Group Name cs2: Affected Group Id flexNumber1: Old Number of Users flexNumber2: Old Number of User Rights
platform:242	5	/Platform/Authorization/Groups/Membership/Update/Clear	Removed all members from group	
platform:243	3	/Platform/Authorization/Groups/Membership/Update	Modified user group membership	
platform:244	3	/Platform/Authorization/Groups/Delete	Deleted user group	cs1: Affected Group Name cs2: Affected Group Id
platform:245	3	/Platform/Authorization/Users/Add	Added user	cs1: Affected User Id cs2: Affected User Login cs3: Affected User Full Name
platform:246	3	/Platform/Authorization/Users/Update	Updated user	cs1: Affected User Id cs2: Affected User Login cs3: Affected User Full Name

B Audit Events

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform:247	3	/Platform/Authorization/Users/Delete	Deleted user	cs1: Affected User Id cs2: Affected User Login cs3: Affected User Full Name
platform:248	3	/Platform/Authentication/Logout/SessionExpiration	Session expired	
platform:249	7	/Platform/Authentication/AccountLocked	Account locked	
platform:250	5	/Platform/Storage/RFS/Add	Added remote mount point	cs1: RFS Mount Name cs2: RFS Mount Host and Remote Path
platform:251	5	/Platform/Storage/RFS/Edit	Edited remote mount point	cs1: RFS Mount Name cs2: RFS Mount Host and Remote Path
platform:252	7	/Platform/Storage/RFS/Failure	Failed to create remote mount point	cs1: Server cs2: Remote Directory cs3: Mount Name cs4: Mount Type cs5: Username
platform:253	5	/Platform/Storage/RFS/Remove	Removed remote mount point	cs1: RFS Mount Name cs2: RFS Mount Host and Remote Path
platform:254	5	/Platform/Storage/SAN/Destroy	Destroyed SAN Logical Unit	cs1: Volume label
platform:255	5	/Platform/Storage/SAN/Attach	Attached SAN Logical Unit	cn2: Volume size (in MB) cs1: Volume label cs2: World-wide Name cs3: Filesystem type
platform:256	7	/Platform/Storage/SAN/Detach	Detached SAN Logical Unit	cs1: Storage unit details
platform:259	5	/Platform/Storage/SAN/Reattach	Reattached SAN Logical Unit	cs1: Volume label cs2: Filesystem type
platform:260	5	/Platform/Configuration/Network/Route/Update	Static route modified	cs1: Destination cs2: Subnet cs3: Gateway
platform:261	5	/Platform/Configuration/Network/Route/Remove	Static route removed	cs1: Destination cs2: Subnet cs3: Gateway
platform:262	5	/Platform/Configuration/Time	Appliance time modified	cs1: Old Date/Time cs2: New Date/Time cs3: Old Time Zone cs4: New Time Zone

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform:263	5	/Platform/Configuration/Network	NIC settings modified	cs1: NIC cs2: IP Address cs3: Netmask cs4: Speed
platform:264	5	/Platform/Configuration/Network/NTP	NTP server settings modified	cs1: NTP Servers cs2: Is Appliance NTP Server
platform:265	5	/Platform/Configuration/Network/DNS	DNS settings modified	
platform:266	5	/Platform/Configuration/Network/Hosts	Hosts file modified	cs1: Difference from previous hosts file
platform:267	5	/Platform/Configuration/SMTP	SMTP settings modified	cs1: EMail Address cs2: SMTP Server cs3: Backup SMTP Server
platform:268	5	/Platform/Configuration/Network/Route/Add	Static route added	cs1: Destination cs2: Subnet cs3: Gateway
platform:270	5	/Platform/Authorization/Users/Inactive/Disable	Inactive user disabled	cs1: User Login deviceCustomDate1: Date Last Active
platform:280	7	/Appliance/State/Reboot/Initiate	Appliance reboot initiated	
platform:281	3	/Appliance/State/Reboot/Cancel	Appliance reboot canceled	
platform:282	7	/Appliance/State/Shutdown	Appliance poweroff initiated	
platform:284	5	/Platform/Storage/Multipathing/Enable	Enabled SAN Multipathing	cs1: Multipath Configuration
platform:285	5	/Platform/Storage/Multipathing/Disable	Disabled SAN Multipathing	
platform:300	5	/Platform/Certificate/Install	Installed trusted certificate	cs1: Certificate details
platform:301	5	/Platform/Certificate/Revocation/Install	Installed certificate revocation list	cs1: CRL details
platform:302	5	/Platform/Certificate/Delete	Deleted trusted certificate	cs1: Certificate details
platform:303	5	/Platform/Certificate/Revocation/Delete	Deleted certificate revocation list	cs1: CRL details
platform:304	7	/Platform/Certificate/Install/Failure	Failed installing trusted certificate	cs1: Error cs2: File Size cs3: File Name

B Audit Events

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform:305	7	/Platform/Certificate/Revocation/Install/Failure	Failed installing certificate revocation list	cs1: Error cs2: File Size cs3: File Name
platform:306	5	/Platform/Process/Start	Start process	cs1: Process Name
platform:307	5	/Platform/Process/Stop	Stop process	cs1: Process Name
platform:308	5	/Platform/Process/Restart	Restart process	cs1: Process Name
platform:310	5	/Platform/Configuration/FIPS/Enable	Enabled FIPS mode	
platform:311	7	/Platform/Configuration/FIPS/Disable	Disabled FIPS mode	
platform:312	7	/Platform/Configuration/WebServer/CipherStrength	Web server cipher strength changed	cs1: New Value cs2: Old Value
platform:320	3	/Appliance/State/Shutdown/Cancel	Appliance poweroff canceled	
platform:371	5	/Platform/Service/Restart	Restarted OS service	cs1: Service Name
platform:400	2	/Platform/Diagnostics/Command	Ran diagnostic command	cs1: Diagnostic Command
platform:407	7	/Platform/Certificate/SSL/Expiration	SSL certificate expiration warning	cs1: Issuer cs2: Subject deviceCustomDate1: Expiration Date
platform:408	5	/Appliance/State/Startup	Appliance startup completed	deviceCustomDate1: Startup Date
platform:409	3	/Platform/Configuration/LoginBanner	Configure login warning banner	cs1: Acknowledgement Prompt cs2: Banner Text
platform:410	5	/Platform/Configuration/Network	Network settings modified	cs1: Gateway cs2: Multi-homing cs3: Hostname
platform:411	5	/Platform/Authentication/PasswordChange	Automated Password Reset	cn2: User ID cs1: User Login
platform:412	3	/Platform/Configuration/Locale	Set Locale	cs1: Locale
platform:440	3	/Platform/Configuration/SNMP	SNMP configuration modified	cn2: Port Number cn3: Refresh Interval cs1: SNMP Enabled cs2: Community String cs3: Listen Address(es)

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform:460	3	/Platform/Network/Alias/Add	NIC alias added	cs1: NIC cs2: IP Address cs3: Netmask
platform:462	3	/Platform/Network/Alias/Remove	NIC alias removed	cs1: NIC cs2: IP Address cs3: Netmask
platform:500	5	/Platform/Authorization/Groups/Membership/Remove	Remove member from group	cs1: Affected Group Name cs2: Affected User Login cs3: Affected Group Id cs4: Affected User Id
platform:501	5	/Platform/Authorization/Groups/Membership/Add	Group member added	cs1: Affected Group Name cs2: Affected User Login cs3: Affected Group Id cs4: Affected User Id
platform:502	5	/Platform/Authorization/Users/Groups/Remove	User removed from group	cs1: Affected Group Name cs2: Affected User Login cs3: Affected Group Id cs4: Affected User Id
platform:503	5	/Platform/Authorization/Users/Groups/Add	User added to group	cs1: Affected Group Name cs2: Affected User Login cs3: Affected Group Id cs4: Affected User Id
platform:530	5	/Platform/Configuration/Authentication/Sessions/Success	Authentication Session settings successfully changed.	cn2: New Value cn3: Old Value cs1: Parameter Changed
platform:540	5	/Platform/Configuration/Authentication/Password/Lockout/Success	Password Lockout settings successfully updated.	cn2: New Value cn3: Old Value cs1: Parameter Changed
platform:550	5	/Platform/Configuration/Authentication/Password/Expiration/Success	Password Expiration settings successfully updated.	cn2: New Value cn3: Old Value cs1: Parameter Changed
platform:560	5	/Platform/Configuration/Authentication/Password/Validation/Success	Password Validation settings successfully updated.	cn2: New Value cn3: Old Value cs1: Parameter Changed

B Audit Events

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform:570	5	/Platform/Configuration/Authentication/Password/AutomatedPassword Reset/Success	Password Automated Password Reset setting successfully updated.	cs1: Parameter Changed cs2: New Value cs3: Old Value
platform:580	5	/Platform/Configuration/Authentication/Certificate/Success	Client Certificate authentication settings successfully changed.	cs1: Parameter Changed cs2: New Value cs3: Old Value
platform:590	5	/Platform/Configuration/Authentication/RADIUS/Success	RADIUS authentication settings successfully changed.	cs1: Parameter Changed cs2: New Value cs3: Old Value
platform:600	5	/Platform/Configuration/Authentication/LDAP/Success	LDAP authentication settings successfully changed.	cs1: Parameter Changed cs2: New Value cs3: Old Value
platform:610	5	/Platform/Configuration/Authentication/Global/Success	Global Authentication settings successfully changed.	cs1: Parameter Changed cs2: New Value cs3: Old Value

OLI Application Events

The following table lists the information contained in audit events related to various OLI functions and configuration changes on it. The Severity for all OLI application events is 2.

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
Alerts			
logger:610	/Logger/Component/Alert/Configuration/Add	Alert [name] has been added	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmpIpAddr dvchost=syslogOrSnmp HostName cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
logger:611	/Logger/Component/Alert/Configuration/Delete	Alert [name] has been deleted	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmpIpAddr dvchost=syslogOrSnmp HostName cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:612	/Logger/Component/Alert/Configuration/Update	Alert [name] has been updated	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmplpAddr dvchost=syslogOrSnmpHostName cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
logger:613	/Logger/Component/Alert/Configuration/Enable	Alert [name] has been enabled	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmplpAddr dvchost=syslogOrSnmp HostName cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
logger:614	/Logger/Component/Alert/Configuration/Disable	Alert [name] has been disabled	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmplpAddr dvchost=syslogOrSnmp HostName cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:615	/Logger/Alert/Configuration/Sent	Alert [name] has been sent	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmpIpAddr dvchost=syslogOrSnmpOrEsmHostName cn1Label=SyslogoOrSnmpOrEsmDestination Port cn1=syslogOrSnmpOrEsmPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
Certificates			
logger:643	/Logger/Component/Certificate/Configuration/Add	Certificate [name] has been added	fname=alias duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Certificate
logger:650	/Logger/Component/Certificate/Configuration/Delete	Certificate [name] has been deleted	fname=alias duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Certificate
logger:651	/Logger/Component/Certificate/Configuration/Update	Certificate [name] has been updated	fname=alias duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Certificate
Configuration Backup			
logger:660	/Logger/Component/ConfigBackup/Configuration/Update	Configuration backup has been updated	fname=Configuration Backup duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Configuration Backup
logger:661	/Logger/Component/ConfigBackup/Configuration/Enable	Configuration backup has been enabled	fname=Configuration Backup duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Configuration Backup

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:662	/Logger/Component/ConfigBackup/Configuration/Disable	Configuration backup has been disabled	fname=Configuration Backup duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Configuration Backup
logger:665	/Logger/Component/ConfigBackup/Configuration/Backup	Configuration backup succeeded. Transfer process finished.	fname=Configuration Backup fileType=Configuration Backup fpath=pathToBackupFile fsize=fileSizeInByte
Forwarders			
logger:605	/Logger/Component/Forwarder/Configuration/Add	Forwarder [name] has been added	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger:606	/Logger/Component/Forwarder/Configuration/Delete	Forwarder [name] has been deleted	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger:607	/Logger/Component/Forwarder/Configuration/Update	Forwarder [name] has been updated	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:608	/Logger/Component/Forwarder/Configuration/Enable	Forwarder [name] has been enabled	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger:609	/Logger/Component/Forwarder/Configuration/Disable	Forwarder [name] has been disabled	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger:663	/Logger/Component/Forwarder/Configuration/Pause	Forwarder [name] has been paused	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger:664	/Logger/Component/Forwarder/Configuration/Resume	Forwarder [name] has been resumed	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
Receivers			
logger:600	/Logger/Component/Receiver/Configuration/Add	Receiver [name] has been added	fname=receiverName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort
logger:601	/Logger/Component/Receiver/Configuration/Delete	Receiver [name] has been deleted	fname=receiverName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort
logger:602	/Logger/Component/Receiver/Configuration/Update	Receiver [name] has been updated	fname=receiverName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort
logger:603	/Logger/Component/Receiver/Configuration/Enable	Receiver [name] has been enabled	fname=receiverName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort
logger:604	/Logger/Component/Receiver/Configuration/Disable	Receiver [name] has been disabled	fname=receiverName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
SNMP Destinations			
logger:644	/Logger/Component/SnmpDestination/Configuration/Add	SNMP destination [name] has been added	fname=snmpDestinationName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=SNMP Destination fileId=snmpDestinationId dvc=snmpDestinationIp dvchost=snmpDestinationHost cn1Label=SNMP Destination Port cn1=snmpDestinationPort cs1Label=ConnectorName cs1=connectorName cs2Label=Connector Location cs2=connectorLocation cs3Label=Logger Location cs3=loggerLocation
logger:645	/Logger/Component/SnmpDestination/Configuration/Delete	SNMP destination [name] has been deleted	fname=snmpDestinationName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=SNMP Destination fileId=snmpDestinationId dvc=snmpDestinationIp dvchost=snmpDestinationHost cn1Label=SNMP Destination Port cn1=snmpDestinationPort cs1Label=ConnectorName cs1=connectorName cs2Label=Connector Location cs2=connectorLocation cs3Label=Logger Location cs3=loggerLocation
Syslog Destinations			
logger:647	/Logger/Resource/SyslogDestination/Configuration/Add	Syslog destination [name] has been added	fname=syslogDestinationName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Syslog Destination fileId=syslogDestinationId dvc=syslogDestinationIp dvchost=syslogDestinationHost cn1Label=Syslog Destination Port cn1=syslogDestinationPort

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:648	/Logger/Component/SyslogDestination/Configuration/Delete	Syslog destination [name] has been deleted	fname=syslogDestinationName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Syslog Destination fileId=syslogDestinationId dvc=syslogDestinationIp dvchost=syslogDestinationHost cn1Label=Syslog Destination Port cn1=syslogDestinationPort
logger:649	/Logger/Component/SyslogDestination/Configuration/Update	Syslog destination [name] has been updated	fname=syslogDestinationName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Syslog Destination fileId=syslogDestinationId dvc=syslogDestinationIp dvchost=syslogDestinationHost cn1Label=Syslog Destination Port cn1=syslogDestinationPort
Archives			
logger:520	/Logger/Resource/Archive/Configuration/Add	Archive [archiveName] has been added	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger:521	/Logger/Resource/Archive/Configuration/Delete	Archive [archiveName] has been deleted	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger:523	/Logger/Resource/Archive/Configuration/Load	Archive [archiveName] has been loaded	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:524	/Logger/Resource/Archive/Configuration/Unload	Archive [archiveName] has been unloaded	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger:525	/Logger/Resource/Archive/Configuration/Archive	Archive [archiveName] has been archived	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger:526	/Logger/Resource/Archive/Add	Event archive settings added	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger:527	/Logger/Resource/Archive/Update	Daily archive task settings updated	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger:528	/Logger/Resource/Archive/Failed	Event archive failed	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
Dashboards			
logger:580	/Logger/Resource/Dashboard/Configuration/Add	Dashboard [name] has been added	fname=dashboardName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Dashboard fileId=DashboardId rt=receiptTime

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:581	/Logger/Resource/Dashboard/Configuration/Add	Dashboard [name] has been deleted	fname=dashboardName duser=UserName duid=userId cs4=sessionIdfile fileType=Dashboard fileId=DashboardId rt=receiptTime
logger:582	/Logger/Resource/Dashboard/Configuration/Update	Dashboard [name] has been updated	fname=dashboardName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Dashboard fileId=DashboardId rt=receiptTime
Devices			
logger:510	/Logger/Resource/Device/Configuration/Add	Device [deviceName] has been added	fname=deviceName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Device fileId=deviceId
logger:511	/Logger/Resource/Device/Configuration/Delete	Device [deviceName] has been deleted	fname=deviceName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Device fileId=deviceId
logger:512	/Logger/Resource/Device/Configuration/Update	Device [deviceName] has been updated	fname=deviceName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Device fileId=deviceId
Filters			
logger:500	/Logger/Resource/Filter/Configuration/Add	Filter [filterName] has been added	fname=filterName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Filter fileId=filterId

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:501	/Logger/Resource/Filter/Configuration/Delete	Filter [filterName] has been deleted	fname=filterName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Filter fileId=filterId
logger:502	/Logger/Resource/Filter/Configuration/Update	Filter [filterName] has been updated	fname=filterName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Filter fileId=filterId
Groups			
logger:513	/Logger/Resource/Group/Configuration/Add	Group [groupName] has been added	fname=groupName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Group fileId=groupId
logger:514	/Logger/Resource/Group/Configuration/Delete	Group [groupName] has been deleted	fname=groupName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Group fileId=groupId
logger:515	/Logger/Resource/Group/Configuration/Update	Group [groupName] has been updated	fname=groupName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Group fileId=groupId
Peer Loggers			
logger:550	/Logger/Resource/PeerLogger/Configuration/Add	Peer Logger [name] has been added	fname=Name duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger fileId=LoggerId

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:551	/Logger/Resource/PeerLogger/Configuration/Delete	Peer Logger [name] has been deleted	fname=Name duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger fileId=LoggerId
logger:570	/Logger/Resource/PeerLogger/Authorizations/Configuration/Added	Peer Logger authorization [name] has been added	fname=Name duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger Authorization
logger:571	/Logger/Resource/PeerLogger/Authorizations/Configuration/Delete	Peer Logger authorization [name] has been deleted	fname=Name duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger Authorization fileId=LoggerId
Parsers			
logger:590	/Logger/Resource/ParserDescription/Configuration/Add	Parser Description [name] has been added	fileType=Parser Description duid=1 cs4=sessionIdfile cs4Label=Session ID duser=UserName rt=receiptTime fname=parserName
logger:591	/Logger/Resource/ParserDescription/Configuration/Delete	Parser Description [name] has been deleted	fileType=Parser Description cs4=sessionIdfile duser=UserName fileId=ParserID 710 duid=1 cs4Label=Session ID rt=receiptTime fname=parserName
logger:592	/Logger/Resource/ParserDescription/Configuration/Update	Parser Description [name] has been updated	fileType=Parser Description cs4=sessionIdfile duser=UserName fileId=ParserID duid=1 cs4Label=Session ID rt=receiptTime fname=parserName

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
Saved Searches			
logger:540	/Logger/Resource/SavedSearch/Configuration/Add	Saved search [name] has been added	fname=savedSearchName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Saved Search fileId=savedSearchId
logger:541	/Logger/Resource/SavedSearch/Configuration/Delete	Saved search [name] has been deleted	fname=savedSearchName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Saved Search fileId=savedSearchId
logger:542	/Logger/Resource/SavedSearch/Configuration/Update	Saved search [name] has been updated	fname=savedSearchName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Saved Search fileId=savedSearchId
Source Types			
logger:596	/Logger/Resource/SourceType/Configuration/Add	Source Type [name] has been added	cs4=sessionIdfile fileType=Source Type duid=1 cs4Label=Session ID duser=UserName rt=receiptTime fname=SourceTypeName
logger:597	/Logger/Resource/SourceType/Configuration/Delete	Source Type [name] has been deleted	fileType=Source Type cs4=sessionIdfile duser=UserName fileId=SourceTypeID duid=1 cs4Label=Session ID rt=receiptTime fname=SourceTypeName
logger:598	/Logger/Resource/SourceType/Configuration/Update	Source Type [name] has been updated	fileType=Source Type cs4=sessionIdfile duser=UserName fileId=1SourceTypeID duid=1 cs4Label=Session ID rt=receiptTime fname=SourceTypeName

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
Storage Groups			
logger:530	/Logger/Resource/StorageGroup/Configuration/Add	Storage group [storageGroupName] has been added	fname=storageGroupName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Group fileId=storageGroupId
logger:532	/Logger/Resource/StorageGroup/Configuration/Update	Storage group [storageGroupName] has been updated	fname=storageGroupName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Group fileId=storageGroupId
Storage Rules			
logger:533	/Logger/Resource/StorageRule/Configuration/Add	Storage rule [name] has been added	fname=storageRuleName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Rule
logger:535	/Logger/Resource/StorageRule/Configuration/Update	Storage rule [name] has been updated	fname=storageRuleName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Rule
Storage Volume			
logger:536	/Logger/Resource/StorageVolume/Configuration/Add	Storage volume [name] has been added	fname=storageVolumeName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Volume fileId=storageVolumeId
Search			
logger:680	/Logger/Search/Index/Update	Search indices have been added OR Search index has been added	cs4=sessionId fileType=Search Index Configuration duser=UserName msg=Search index has been added cn1=1 duid=1 cs4Label=Session ID rt=receiptTime cn1Label=No. of fields added

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:690	/Logger/Search/Options/Update	Search options have been updated	cs6=false cs7=true cs4=sessionId cs5=false cs2=false cs3=false cs1=true cs8=false cs1Label=Field Search Case Sensitivity duid=1 cs7Label=Field Summary cs8Label=Field Summary Field Discovery cs6Label=Display options rawEvent cs3Label=Regex Search Unicode Case Sensitivity fileType=Search Options duser=UserName cs5Label=Regex Search Canonical Equality Check cs4Label=Session ID rt=receiptTime cs2Label=Regex Search Case Sensitivity
logger:710	/Logger/Search/Cancelled	Search session [sessionId] has been cancelled by [user]	cs1Label=Session ID duid=1 cs1=sessionIdfile duser=UserName rt=receiptTime
Maintenance Mode			
logger:700	/Logger/Server/MaintenanceMode/Enter	Maintenance mode entered	fname=Maintenance Mode duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Maintenance Mode

Appendix C: Examples of System Health Events

The following table provides examples of system health events generated on OLI. These examples are intended to help you understand the format and various fields of the generated events.



Note

You can set up Alerts to be triggered to let you know when system health events are generated. [“Alerts” on page 172.](#)

The table includes information on the following system health event classes:

- [“cpu” on page 317](#)
- [“disk” on page 317](#)
- [“eps” on page 318](#)
- [“hardware” on page 318](#)
- [“memory” on page 321](#)
- [“network” on page 321](#)
- [“raid” on page 321](#)
- [“search” on page 323](#)
- [“storagegroup” on page 323](#)

Device Event Class: ID	Example
cpu	
cpu:100	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 cpu:100 CPU Usage 1 cat=/Monitor/CPU/Usage cn1=3 cn1Label=Percent Usage cs2=CurrentValue cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1302739080014 rt=1302739080014
disk	
disk:101	CEF:0 HP Operations Log Intelligence 5.1.0.5803.0 disk:101 Root Disk Space Remaining 1 cat=/Monitor/Disk/Space/Remaining/Root cn1=99 cn1Label=Percent Available cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=Ok cs4Label=Raw Status cs5=Root cs5Label=Location cs6=Disk/Space/Remaining/Root cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303927171790 rt=1303927171790

Device Event Class: ID	Example
disk:102	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 disk:102 Disk bytes read 1 cat=/Monitor/Disk/Read cn1=373524 cn1Label=Kb Read cs2=SinceStartup cs2Label=timeframe cs6=c0d0 cs6Label=Partition Name dst=192.168.35.115 dvc=192.168.35.115 end=1302743760036 rt=1302743760036
disk:103	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 disk:103 Disk bytes written 1 cat=/Monitor/Disk/Write cn1=24474998 cn1Label=Kb Written cs2=SinceStartup cs2Label=timeframe cs6=c0d0 cs6Label=Partition Name dst=192.168.35.115 dvc=192.168.35.115 end=1302743760038 rt=1302743760038
eps	
eps:100	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 eps:100 Overall Receiver EPS 1 cat=/Monitor/Receiver/All/EPS cn1=0 cn1Label=EPS cs2=SinceLastMonitorEvent cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1302733680034 rt=1302733680034
eps:101	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 eps:101 Overall Forwarder EPS 1 cat=/Monitor/Forwarder/All/EPS cn1=0 cn1Label=EPS cs2=SinceLastMonitorEvent cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115
eps:102	CEF:0 HP Operations Log Intelligence 5.1.0.5803.0 eps:102 Individual Receiver EPS 1 cat=/Monitor/Receiver/EPS/Individual cn1=0 cn1Label=EPS cs1=N/A cs1Label=Receiver Type cs2=SinceLastMonitorEvent cs2Label=timeframe cs6=udp514 cs6Label=name dst=192.168.35.115 dvc=192.168.35.115 end=1303927500046 rt=1303927500046
eps:103	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 eps:103 Individual Forwarder EPS 1 cat=/Monitor/Forwarder/One/EPS cn1=0 cn1Label=EPS cs1=N/A cs1Label=Forwarder Type cs2=SinceLastMonitorEvent cs2Label=timeframe cs6=esm cs6Label=name dst=192.168.35.115 dvc=192.168.35.115 end=1302733080052 rt=1302733080052
hardware	
hardware:101	CEF:0 HP Operations Log Intelligence 5.1.0.5784.0 hardware:101 Electrical (Current) OK 1 cat=/Monitor/Sensor/Current/Ok cs1=0.80 Amps cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=10.2 (Power Supply) cs5Label=Location cs6=Current 2 cs6Label=Sensor Name dst=192.168.36.5 dvc=192.168.36.5 end=1303937520837 rt=1303937520837

Device Event Class: ID	Example
hardware:102	CEF:0 HP Operations Log Intelligence 5.1.0.5776.0 hardware:102 Electrical (Current) Degraded 5 cat=/Monitor/Sensor/Current/Degraded cs1=126 Watts cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Inc cs4Label=Raw Status cs5Label=Location cs6=Power Meter cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302817019262 rt=1302817019262
hardware:103	CEF:0 HP Operations Log Intelligence 5.1.0.5776.0 hardware:103 Electrical (Current) Failed 8 cat=/Monitor/Sensor/Current/Failed cs1=126 Watts cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=lcr cs4Label=Raw Status cs5Label=Location cs6=Power Meter cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302817019262 rt=1302817019262
hardware:111	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 hardware:111 Electrical (Voltage) OK 1 cat=/Monitor/Sensor/Voltage/Ok cs1=State Deasserted cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=3.1 (Processor) cs5Label=Location cs6=VCORE cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302819047959 rt=1302819047959
hardware:112	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 hardware:112 Electrical (Voltage) Degraded 5 cat=/Monitor/Sensor/Voltage/Degraded cs1=State Deasserted cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Inc cs4Label=Raw Status cs5=3.1 (Processor) cs5Label=Location cs6=VCORE cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302819047959 rt=1302819047959
hardware:113	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 hardware:113 Electrical (Voltage) Failed 8 cat=/Monitor/Sensor/Voltage/Failed cs1=State Deasserted cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=lcr cs4Label=Raw Status cs5=3.1 (Processor) cs5Label=Location cs6=VCORE cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302819047959 rt=1302819047959
hardware:121	CEF:0 HP Operations Log Intelligence 5.1.0.5803.0 hardware:121 Battery OK 1 cat=/Monitor/Sensor/Battery/Ok cs1=cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=CMOS Battery cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303937972008 rt=1303937972008

Device Event Class: ID	Example
hardware:122	CEF:0 HP Operations Log Intelligence 5.1.0.5803.0 hardware:122 Battery Degraded 5 cat=/Monitor/Sensor/Battery/Degraded cs1=cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Inc cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=CMOS Battery cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303937972008 rt=1303937972008
hardware:123	CEF:0 HP Operations Log Intelligence 5.1.0.5803.0 hardware:123 Battery Failed 8 cat=/Monitor/Sensor/Battery/Degraded cs1=cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=lcr cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=CMOS Battery cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303937972008 rt=1303937972008
hardware:131	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 hardware:131 Fan OK 1 cat=/Monitor/Sensor/Fan/Ok cs1=29.01 unspecified cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=Inc cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=Fan Block 1 cs6Label=Sensor Name dst=192.168.37.21 dvc=192.168.37.21 end=1302823237825 rt=1302823237825
hardware:132	
hardware:133	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 hardware:133 Fan Failure 8 cat=/Monitor/Sensor/Fan/Failed cs1=29.01 unspecified cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=lcr cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=Fan Block 1 cs6Label=Sensor Name dst=192.168.37.21 dvc=192.168.37.21 end=1302823237825 rt=1302823237825
hardware:141	CEF:0 HP Operations Log Intelligence 5.1.0.5803.0 hardware:141 Power Supply OK 1 cat=/Monitor/Sensor/PowerSupply/Ok cs1=cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=10.1 (Power Supply) cs5Label=Location cs6=Status cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303938572149 rt=1303938572149
hardware:142	
hardware:143	CEF:0 HP Operations Log Intelligence 5.1.0.5776.0 hardware:143 Power Supply Failed 8 cat=/Monitor/Sensor/PowerSupply/Failed cs1=0 Watts cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=lcr cs4Label=Raw Status cs5=10.2 (Power Supply) cs5Label=Location cs6=Power Supply 2 cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302817019263 rt=1302817019263

Device Event Class: ID	Example
hardware:151	CEF:0 HP Operations Log Intelligence 5.1.0.5776.0 hardware:151 Temperature OK 1 cat=/Monitor/Sensor/Temperature/Ok cs1=17 degrees C cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=64.1 (Unknown (0x40)) cs5Label=Location cs6=Temp 1 cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302823560051 rt=1302823560051
hardware:152	
hardware:153	
memory	
memory:100	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 memory:100 Platform Memory Usage 1 cat=/Monitor/Memory/Usage/Platform cn1=2757 cn1Label=MB Used cs2=CurrentValue cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1302797940018 rt=1302797940018
network	
network:100	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 network:100 Network Usage - Inbound 1 cat=/Monitor/Network/Usage/In cn1=41837428 cn1Label=Bytes Received cs2=SinceStartup cs2Label=timeframe cs6=eth0 cs6Label=Interface Name dst=192.168.35.115 dvc=192.168.35.115 end=1302733620026 rt=1302733620026
network:101	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 network:101 Network Usage - Outbound 1 cat=/Monitor/Network/Usage/Out cn1=158442791 cn1Label=Bytes Sent cs2=SinceStartup cs2Label=timeframe cs6=eth0 cs6Label=Interface Name dst=192.168.35.115 dvc=192.168.35.115 end=1302733620028 rt=1302733620028
raid	
raid:101	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 raid:101 RAID Controller OK 1 cat=/Monitor/RAID/Controller/Ok cs1=Type: RAID-5 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=Optimal cs4Label=Raw Status cs5=RAIDController/Configuration cs5Label=Location cs6=RAIDController/Configuration cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302886250104 rt=1302886250104

Device Event Class: ID	Example
raid:102	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 raid:102 RAID Controller Degraded 5 cat=/Monitor/RAID/ControllerDegraded cs1=Type: RAID-5 Critical Disks: 0 Failed Disks: 0 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Degraded cs4Label=Raw Status cs5=RAIDController/Configuration cs5Label=Location cs6=RAIDController/Configuration cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302826128482 rt=1302826128482
raid:103	
raid:111	CEF:0 HP Operations Log Intelligence 5.1.0.5776.0 raid:111 RAID BBU OK 1 cat=/Monitor/RAID/BBU/Ok cs1=Battery/Capacitor Count: 1 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=OK cs4Label=Raw Status cs5=RAIDController/Battery/bbu cs5Label=Location cs6=RAIDController/Battery/bbu cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302890169285 rt=1302890169285
raid:112	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 raid:112 RAID BBU Degraded 5 cat=/Monitor/RAID/BBU/Degraded cs1=Fully Charged: false Remaining Time Alarm: false Remaining Capacity Alarm: false Over Charged: false isSOHGood: true cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Degraded cs4Label=Raw Status cs5=RAIDController/Battery/bbu cs5Label=Location cs6=RAIDController/Battery/bbu cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302820608015 rt=1302820608015
raid:113	
raid:121	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 raid:121 RAID Disk OK 1 cat=/Monitor/RAID/DISK/Ok cs1=Port: 1I Box: 1 Bay: 1 Size: 500 GB Serial Number: 9SP24JD5 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=OK cs4Label=Raw Status cs5=Port: 1I Box: 1 Bay: 1 Serial Number: 9SP24JD5 cs5Label=Location cs6=RAIDController/Port/p0 cs6Label=Sensor Name dst=192.168.37.21 dvc=192.168.37.21 end=1302849041777 rt=1302849041777
raid:122	CEF:0 HP Operations Log Intelligence 5.1.0.5776.0 raid:122 RAID Disk Rebuilding 5 cat=/Monitor/RAID/DISK/Rebuilding cs1=Port: 2I Box: 1 Bay: 1 Size: 1 TB Serial Number: WMATV6348517 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Rebuilding cs3Label=Status cs4=Rebuilding cs4Label=Raw Status cs5=Port: 2I Box: 1 Bay: 1 Serial Number: WMATV6348517 cs5Label=Location cs6=RAIDController/Port/p0 cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302826980530 rt=1302826980530

Device Event Class: ID	Example
raid:123	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 raid:123 RAID Disk Failed 8 cat=/Monitor/RAID/DISK/Failed cs1=Port: 1 Box: 1 Bay: 2 Size: 500 GB Serial Number: 9SP23M08 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=Failed cs4Label=Raw Status cs5=Port: 1 Box: 1 Bay: 2 Serial Number: 9SP23M08 cs5Label=Location cs6=RAIDController/Port/p1 cs6Label=Sensor Name dst=192.168.37.21 dvc=192.168.37.21 end=1302826358346 rt=1302826358346
search	CEF:0 HP Operations Log Intelligence 5.1.0.5780.0 search:100 Number of Searches Performed 1 cat=/Monitor/Search cn1=0 cn1Label=Number of Searches cs2=SinceStartup cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1302741300026 rt=1302741300026
storagegroup	CEF:0 HP Operations Log Intelligence 5.1.0.5803.0 storagegroup:100 Storage Group Space Used 1 cat=/Monitor/StorageGroup/Space/Used cn1=1 cn1Label=Percent Used cn2=7 cn2Label=retention period (days) cn3=1024 cn3Label=used (MB) cs2=CurrentValue cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1303928072008 fileType=storageGroup fname=Default Storage Group fsize=1534 rt=1303928072008

Appendix D: Using the Rex Operator

The `rex` operator is a powerful operator that enables you to extract information that matches a specified regular expression and assigns it to a field, whose field name you specify. You can also specify an optional start point and an end point in the rex expression between which the information matching the regular expression is searched.

This appendix describes the `rex` search operator in detail. It includes information on the following topics.

[“Syntax of the rex Operator” on page 325](#)

[“Ways to Create a rex Expression” on page 326](#)

[“Samples of rex Expressions” on page 327](#)

When a rex expression is included in a search query, it must be preceded by a basic search query that finds events from which the rex expression will extract information. For example:

```
failed | rex "(?<srcip>[^\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
```

Syntax of the rex Operator

```
| rex "text1 (?<field1>text2regex) "
```

text1—The text or point in the event AFTER which information extraction begins. The default is the beginning of the event.

text2—The text or point in the event at which information extraction ends.

field1—The name of the field to which the extracted information is assigned.

regex—The pattern (regular expression) used for matching information to be extracted between *text1* and *text2*.



If you are an experienced regular expression user, see the Note in the next section for a quick understanding of how rex enables you to capture named input and reference it for further processing.

Understanding the rex Operator Syntax

Extract all information AFTER *text1* and until *text2* that matches the specified *regex* (regular expression) and assign TO *field1*.

- **text1** and [**text2**] can be any points in an event—start and end of an event, specific string in an event (even if the string is in the middle of a word in the event), a specific number of characters from the start or end of an event, or a pattern.
- To specify the next space in the event as **text2**, enter [^].
This is interpreted as “not space.” Therefore, entering a “not” results in the capture to stop at the point where the specified character, in this case, a space, is found in the event.
- To specify [**text2**] to be the end of the line, enter [^\$].
This is interpreted as “not end of line.” Therefore, when an end-of-line in an event is encountered, the capture will stop at that point. The [^\$] usage only captures one character if it is not an end-of-line character. However, by specifying [^\$]* in a rex expression, the usage captures all characters until end-of-line.

You can also specify .* to capture all characters in an event instead of [^\$]. Examples in this document, however, use [^\$].
- Any extra spaces within the double quotes of the rex expression are treated literally.
- The characters that need to be escaped for rex expressions are the same as the ones for regular expressions. Refer to a regular expressions document of your choice to obtain a complete list of such characters.
- Information captured by a rex expression can be used for further processing in a subsequent rex expression as illustrated in the following example in which an IP address is captured by the first rex expression and the network ID (assuming the first three bytes of the IP address represent it) to which the IP address belongs is extracted from the captured IP address:

```
OLI | rex "(?<srcip>[^ ]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" |
rex field=srcip "(?<netid>\d{1,3}\.\d{1,3}\.\d{1,3})"
```



Note

If you are an experienced regular expression user, you can interpret the rex expression syntax as follows:

```
rex "(?<field1>regex)"
```

where the entire expression in the parentheses specifies a named capture. That is, the captured group is assigned a name, which can be referenced later for further processing. For example, in the following expression “srcip” is the name assigned to the capture.

```
failed | rex "(?<srcip>[^ ]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
```

Once named, use “srcip” for further processing as follows:

```
failed | rex "(?<srcip>[^ ]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" |
top srcip
```

Ways to Create a rex Expression

You can create a rex expression in two ways:

- Manually—Follow the syntax and guidelines described in this appendix to create a rex expression to suit your needs.

- **Regex Helper**—Use the Regex Helper tool, as described in [“Regex Helper Tool” on page 91](#). This tool not only simplifies the process, it also makes it less error prone and more efficient.

Creating a rex Expression Manually

Start with a simple search that finds the events that contains the information in which you are interested. Once the events are displayed, identify a common starting point in those events that precedes the information.

For example, you are interested in extracting the client IP address, which always appears after the word “[client” in the following event.

```
[Thu Jul 30 01:20:06 2009] [error] [client 69.63.180.245] PHP
Warning: memcache_pconnect() [a
href='function.memcache-pconnect'>function.memcache-pconnect</a>]:
Can't connect to 10.4.31.4:11211
```

Therefore, “[client” is the starting point. A good end point is the “]” after the last byte of the client IP address. Now, we need to define the regular expression that will extract the IP address. Because in this example, only the client IP address appears after the word “client”, we use “*” as the regular expression, which means “extract everything”. (We could be more specific and use `\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}` for the IP address.) We assign the extracted IP address to a field name “clientIP”. We are almost ready to create a rex expression, except that we need to escape the “[” and “]” characters in the expression. The escape character to use is “\”.

Now, we are ready to create the rex expression to extract the IP address that appears after the word “client” in the event shown above.

```
| rex "\[client(?<clientip>[^\]]*)" "
```

Samples of rex Expressions

This section contains several sample examples for extracting different types of information from an event. The specificity of the information extracted increases with each example. Use these examples as a starting point for creating rex expressions to suit your needs. Also, use the Regex Helper tool that simplifies rex expression creation.

This event is used as an example to illustrate the information the following rex expressions will extract:

```
2010/07/01 13:46:00 PDT      unknown      Local      HP      OLI      4.5.0.4836.0      eps:100
CEF:0|HP|4.5.0.4836.0|eps:100|OLI Internal Event|1| cat=/Monitor/Receiver/A11/EPS cs2=SinceLastMonitorEvent cnt=1 dvc=192.168.36.3 cs1=Events/sr
2010/07/01 13:46:00 PDT      unknown      Local      HP      OLI      4.5.0.4836.0      eps:100
CEF:0|HP|4.5.0.4836.0|eps:100|OLI Internal Event|1| cat=/Monitor/Forwarder/A11/EPS cs2=SinceLastMonitorEvent cnt=1 dvc=192.168.36.3 cs1=Events/sr
```

- **Capture matching events from the left of the pipeline and assign them to the field, message.** The entire event is assigned to the “message” field.

```
| rex "(?<message>[^\$]*) "
```

This expression extracts the entire event (as shown above), starting at the word “CEF:0”.

- Specifying the starting point as number of characters from the start of an event instead of a specific character or word

```
| rex "[a-zA-Z0-9:\.\s]{16}(<message>[^$]*)"
```

This expression starts extracting after 16 **consecutive** occurrences of the characters specified for *text1*—alphanumeric characters, colons, periods, or spaces. Although the first 16 characters of the first event are “CEF:0|HP|O”, the extraction does not begin at “LI|4.5.0...” because the pipeline character is not part of the characters we are matching, but this character is part of the beginning of the event. Therefore, the first 16 consecutive occurrences are “OLI Internal “. As a result, information starting at the word “Event”, is extracted from our example event.

- Extract a specified number of characters instead of specifying an end point such as the next space or the end of the line

```
| rex "[a-zA-Z0-9:\.\s]{16}(<message>[^$]{5})"
```

This expression only extracts the word “Event”. (See the previous sample rex expression for a detailed explanation of the reason extraction begins at the word “Event”.)

- Extract everything after “CEF:0|” into a field, *message*. Then, pipe events for which the *message* field is not null through another rex expression to extract the IP address contained in the matching events and assign the IP addresses to another field, *msgip*. Only display events where *msgip* is not null.

```
| rex "CEF:0|(<message>[^$]*)" | where message is not null |
rex "dvc=(<msgip>[^\s]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" |
where msgip is not null
```



The “: and =” characters do not need to be escaped; however, “[” must be escaped. The characters that need to be escaped for rex expressions are the same as the ones for regular expressions. Refer to a regular expressions document of your choice to obtain a complete list of such characters.

This expression extracts the device IP address from the event.

The following rex examples use this event for illustration:

```
Nov 10 03:04:24 192.168.20.111 192.168.20.112 [redacted] C007:4028:EvilPackets;Line 16:'New Group','My 80280150','11/10/2005 11:02:05.000','21561','11/10/2005 11:02:05.000','3106004','generator','1','192.168.20.111','http:80','192.168.20.112','32771','tcp','Alert','47302','47285','RPC Incomplete Segment','0','0','00:00:00:00:00:00','00:00:00:00:00:00'
```

- Extract the first two IP addresses from an event and assign them to two different fields, IP1 and IP2.

```
| rex "(?<IP1>[^$]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" | rex
"\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}(<IP2>[^$]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
```

This expression extracts the first and second IP addresses in the above event.

Because the two IP addresses are right after one another in this event, you can also specify the extraction of the two IP addresses in a single rex expression as follows:

```
| rex
"(?<IP1>[^\$]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}) (?<IP2>[^\$]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
```



Do not specify a space in the above expression.

- Building on the previous example, add a new field called Ignore. Assign the value "Y" to this field if the two IP addresses extracted in the previous example are the same and assign the value "N" if the two IP addresses are different. Then, list the top IP1 and IP2 combinations for events for which Ignore field is "N".

```
| rex (?<IP1>[^\$]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" | rex
"\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3} (?<IP2>[^\$]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" | eval Ignore=if(IP1==IP2,"Y","N") | where Ignore="N" | top IP1 IP2
```



The eval command uses double == to equate the two fields.

- Information captured by a rex expression can be used for further processing in a subsequent rex expression as illustrated in the following example. The first IP address is captured by the first rex expression and the network ID (assuming the first three bytes of the IP address represent it) to which the IP address belongs is extracted from the captured IP address:

```
logger | rex "(?<srcip>[^\$]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" | rex field=srcip
"(?<netid>\d{1,3}\.\d{1,3}\.\d{1,3})"
```

The following rex example uses this event for illustration:

```
127.0.0.1 - name [10/Oct/2010:13:55:36 -0700] "GET /apache_pb.gif HTTP/1.0" 200 2326
"http://www.example.com/start.html" "Mozilla/4.08 [en] (Win98; I ;Nav)" ¶
```

- Extract all URLs from events and generate a chart of the URL counts, excluding blank URLs

```
| rex "http://(?<customURL>[^\$]*)" | where customURL is not null
| chart count by customURL | sort - customURL
```



- The events contain the URL string in "http://" format.
- Meta character / needs to be enclosed in squarebrackets [] to be treated literally.

The following rex example uses this event for illustration:

D Using the Rex Operator

1	2010/04/06 22:11:46 CDT 10.0.10.222 [Raw_Syslog] Local	root
RAW	Feb 25 14:03:24 beach login(pam_unix)[123]: session closed for user root	
2	2010/04/06 22:11:46 CDT 10.0.10.222 [Raw_Syslog] Local	sysadmin
RAW	Feb 25 14:03:24 beach sshd(pam_unix)[123]: authentication failure; logname=uid=0 euid=0 tty=ssh ruser= rhost=123.123.123 USER =sysadmin	
3	2010/04/06 22:11:46 CDT 10.0.10.222 [Raw_Syslog] Local	padmin
RAW	Feb 25 14:03:24 beach sshd(pam_unix)[123]: session closed for USER padmin	
4	2010/04/06 22:11:46 CDT 10.0.10.222 [Raw_Syslog] Local	sysadmin
RAW	Feb 25 14:03:24 beach sshd(pam_unix)[123]: session opened for USER sysadmin by (uid=500)	

- Extract the first word after the word “user” (one space after the word) or “user=”. The word “user” is case-insensitive in this case and must be preceded by a space character. That is, words such as “ruser” and “suser” should not be matched.

```
| rex "\s[u|U] [s|S] [e|E] [r|R] [\s|=] (?<CustomUser> [^ ]*)" "
```

Index

Numerics

32-bit compatibility libraries not found 30

A

accounts, user. See users.

aggregated storage limit 28

alerts

about 172

adding 176

device event class id 301

disabling 178, 184

enabling 178, 184

real time 172

remove 178

saved search 172

Apache URL Access error log 26

archives

device event class id 308

scheduled 134

ArcSight Manager 42, 165

audit events 293

audit log 26

audit.log 26

authentication

CAC 255

client certificate 255

LDAP 255

RADIUS 255

B

backup 207

browser requirements 45

C

CAC support 239, 244

canonical equality check 47

case-sensitive search 47

CEF 119

event filters 119

TCP receiver 140

UDP receiver 140

certificate revocation list 246

Certificate Signing Request (CSR) 242

certificates

device event class id 303

character encoding 147, 148, 150, 152, 153

client certificate authentication 255

command-line installation 30

Common Access Card (CAC) 244, 255

common event format (CEF) 119

configuration backup 207

creating 208

device event class id 303

guidelines 209

restoring 209

settings 208

Configuration tab 127

connect

web browsers 45

connector forwarder 166

constraints, search 68

content export 229

guidelines 231

content import 229

guidelines 230

CSR

generating a certificate signing request

242

D

dashboards

device event class id 309

data limit 28

date/time format 154

default storage group 15, 40

device event class id 293

device event class ids

alerts 301

archives 308

certificates 303

configuration backup 303

dashboards 309

devices 310

filters 310

- forwarders 304
 - groups 311
 - maintenance mode 315
 - parsers 312
 - platform 294
 - receivers 306
 - saved searches 313
 - search 314
 - SNMP destinations 307
 - source types 313
 - storage groups 314
 - storage rules 314
 - storage volume 314
 - syslog destinations 307
 - device groups 129
 - creating 129
 - deleting 130
 - editing 130
 - maximum number 129
 - devices 127
 - delete 129
 - device event class id 310
 - edit 128
 - maximum number 128
 - pre-defining 128
 - dynamic search 74
- E**
- encoding 147, 148, 150, 152, 153
 - event archival, scheduled 134
 - event archives 131, 133
 - adding 133
 - deleting 134
 - loading 136
 - settings 135
 - unloading 136
 - event input 140, 159
 - folder follower receivers 141
 - receivers 140
 - source type 155
 - event output 165
 - forwarders 165
 - event storage, remote 140
 - events
 - searching 68
 - export
 - alerts 231, 232
 - dashboards 231, 232
 - filters 231, 232
 - OLI content 232
 - parsers 231, 232
 - saved searches 231, 232
 - search results 111
 - source types 231, 232
 - exporting OLI content 231
 - extract parser 160
- F**
- field query
 - indexing fields 112
 - field set, search 68
 - fields, indexing 112
 - file receivers 140, 141
 - multi-line 141
 - file transfer receiver 141
 - filters 189
 - copying 190
 - creating 189
 - deleting 190
 - device event class id 310
 - editing 190
 - search 68
 - system 119
 - finding events 68
 - FIPS 140-2
 - enabling on Connector Appliance 247
 - Firefox (web browser) 45
 - folder follower receivers 141, 142
 - /userdata/logs/apache/http_error_log 26
 - /var/log/audit/audit.log 26
 - /var/log/messages 26
 - OLI
 - folder follower receivers 26
 - forwarder types
 - connector 166
 - TCP 166
 - UDP 166
 - forwarders 165
 - creating 167
 - deleting 171
 - device event class id 304
 - editing 170
 - function tabs 46
- G**
- gauge range 47
 - gauges 46
 - glibc-2.12-1.25.el6.i686 30
 - groups, device event class ID 311
- H**
- health, system 266
 - help 46, 47
 - http_error_log 26

I

- i18n options 47
- import
 - alerts 230
 - dashboards 230
 - filters 230
 - parsers 230
 - saved searches 230
 - source types 230
- import OLI content 230
- importing OLI content 230
- indexing fields 112
- install OLI 31
- installation wizard 30
- internal storage group 15, 136
- internationalization options 47
- Internet Explorer (web browser) 45

L

- launch OLI 36
- LDAP Authentication 255
- localization options 47
- log in 38, 45, 46
- loggerd 36
 - quit 36, 37
 - restart 37
 - start 36, 37
 - status 37
 - stop 36, 37
- login banner 261
- login screen 38, 46
- logout 46, 49
 - automatic 49
- logs
 - internal 229
 - internal, retrieving 229

M

- maintenance mode 210
 - device event class id 315
- Manager 42, 165
- monitor a directory 142
- Monitor tab 62
- multi-line receivers 141

N

- navigation 46
- nss-softokn-freebl-3.12.9-3.el6.i686 30

O

- OLI content
 - export 231, 232

- import 230
- online help 47
- options 46, 47

P

- parse command 160
- parsers 159
 - device event class id 312
 - extract 160
 - REX 160
 - using with source types 159
- password
 - changing 266
- password changing 266
- Peer OLI
 - device event class id 311
- peer OLIs 203
 - adding 205
 - deleting 206
 - searching 97
- platform
 - device event class id 294
- pre-configured folder follower receivers 26
- predefined filters 119

Q

- query
 - controls 46
 - events 68

R

- RADIUS authentication 255
- RAID controller status 237
- range, gauge 47
- real time alerts 172
- receiver types
 - CEF TCP 140
 - CEF UDP 140
 - file receiver 140, 141
 - file transfer 141
 - file transfer receiver 141
 - folder follower 141
 - folder follower receiver 141
 - multi-line 141
 - SmartMessage 141
 - TCP 140
 - UDP 140
- receivers 140, 144
 - creating 144
 - deleting 146
 - device event class id 306
 - disabling 145
 - editing 145

- enabling 145
 - folder follower 142
 - regular expressions (regex)
 - predefined 119
 - Remote Authentication Dial-In User Service (RADIUS) 255
 - remote event storage 140
 - restart 36
 - restart OLI 36
 - retrieve logs 229
 - REX parser 160
- S**
- saved
 - filters 116
 - search 116
 - saved search alerts 172
 - saved search files 197
 - saved search Job 193
 - saved search job
 - adding 194
 - deleting 196
 - editing 196
 - saved searches 192
 - adding 192
 - deleting 193
 - device event class id 313
 - editing 193
 - scheduled event archive 134
 - scheduled tasks 187
 - currently running 188
 - finished 188
 - scheduling
 - export of search results 97
 - SCP file receivers 150
 - search
 - constraints 68
 - device event class id 314
 - events 68
 - exporting results 111
 - field set 68
 - filters 68, 116
 - peer OLIs 97
 - query, defining a 68
 - results, scheduling export of 97
 - saved 116
 - system filters 119
 - time range 68
 - search group filters 191
 - associating with user group 191
 - search operators
 - cef (deprecated) 269
 - chart 269
 - dedup 276
 - eval 277
 - extract 277
 - fields 279
 - head 280
 - keys 280
 - parse 281
 - rare 282
 - regex 283
 - rename 283
 - replace 284
 - rex 286
 - sort 288
 - tail 289
 - top 289
 - transaction 290
 - where 292
 - Search Results tab 99
 - SFTP file receivers 150
 - silent installation 30
 - SmartConnector 26, 42
 - SmartConnectors
 - configuring 42
 - SmartMessage receivers 26, 141
 - configuring 42
 - SNMP 266
 - SNMP destinations
 - device event class id 307
 - source types
 - device event class id 313
 - SSL 239, 242
 - Certificate Signing Request 242
 - start OLI 36
 - starting OLI 36
 - statistics 46
 - status
 - 3Ware RAID Controller 237
 - stop OLI 36
 - storage 136
 - storage groups 136
 - default 15, 40, 136, 137
 - device event class id 314
 - editing 137
 - internal 15
 - storage rules 40, 138
 - adding 138
 - deleting 139
 - device event class id 314
 - editing 139
 - storage settings 140
 - storage volume 140
 - device event class id 314
 - Streaming SmartConnector 166
 - structured data 16, 41

- exporting 19
- syslog destinations
 - device event class id 307
- system audit log 26
- system filters 119
- system health, monitoring 266
- system messages log 26

T

- TCP forwarders 166
- TCP receivers 140
 - default port 26
- time range
 - dynamic 74
 - search 68
- trial license 27

U

- UDP forwarders 166
- UDP receivers 140
 - default port 25
- Unicode options 47
- uninstalling OLI 37
- unstructured data 16
 - exporting 19

- searching 73
- US-ASCII encoding 147, 148, 150, 152, 153
- user groups
 - associating with search group filters 191
 - creating 265
 - deleting 266
 - editing 265
- user interface 46
 - Search Results tab 99
- user rights for content export
 - content export
 - user rights 231
- user rights for content import
 - content import
 - user rights 230
- users
 - changing password 266
 - creating 262
 - deleting 263
 - editing 263
- UTF-8 encoding 147, 148, 150, 152, 153

W

- web browser requirements 45

