
Micro Focus Security ArcSight ESM

Software Version: 7.2 Service Pack 1

Administrator's Guide

Document Release Date: April 2020

Software Release Date: April 2020



Legal Notices

Copyright Notice

© Copyright 2001-2020 Micro Focus or one of its affiliates

Confidential computer software. Valid license from Micro Focus required for possession, use or copying. The information contained herein is subject to change without notice.

The only warranties for Micro Focus 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. Micro Focus shall not be liable for technical or editorial errors or omissions contained herein.

No portion of this product's documentation may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's internal use, without the express written permission of Micro Focus.

Notwithstanding anything to the contrary in your license agreement for Micro Focus ArcSight software, you may reverse engineer and modify certain open source components of the software in accordance with the license terms for those particular components. See below for the applicable terms.

U.S. Governmental Rights. For purposes of your license to Micro Focus ArcSight software, "commercial computer software" is defined at FAR 2.101. If acquired by or on behalf of a civilian agency, the U.S. Government acquires this commercial computer software and/or commercial computer software documentation and other technical data subject to the terms of the Agreement as specified in 48 C.F.R. 12.212 (Computer Software) and 12.211 (Technical Data) of the Federal Acquisition Regulation ("FAR") and its successors. If acquired by or on behalf of any agency within the Department of Defense ("DOD"), the U.S. Government acquires this commercial computer software and/or commercial computer software documentation subject to the terms of the Agreement as specified in 48 C.F.R. 227.7202-3 of the DOD FAR Supplement ("DFARS") and its successors. This U.S. Government Rights Section 18.11 is in lieu of, and supersedes, any other FAR, DFARS, or other clause or provision that addresses government rights in computer software or technical data.

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Support

Contact Information

Phone	A list of phone numbers is available on the Technical Support Page: https://softwaresupport.softwaregrp.com/support-contact-information
Support Web Site	https://softwaresupport.softwaregrp.com/
ArcSight Product Documentation	https://community.softwaregrp.com/t5/ArcSight-Product-Documentation/ct-p/productdocs

Contents

- Chapter 1: Starting and Stopping the Manager and Components12
 - Restarting the Manager - Stop the Manager and Start All Services12
 - Starting the ArcSight Command Center12
 - Starting ArcSight SmartConnectors 13
 - Stopping and Starting ArcSight Services 13
 - Starting the ArcSight Console13
 - Reconnecting ArcSight Console to the Manager14

- Chapter 2: Basic Configuration Tasks15
 - References to ARCSIGHT_HOME 15
 - Managing and Changing Properties File Settings15
 - Property File Format15
 - Defaults and User Properties 15
 - Editing Properties Files16
 - Dynamic Properties17
 - Example18
 - Changing Manager Properties Dynamically20
 - Changing the Service Layer Container Port20
 - Securing the Manager Properties File21
 - Adjusting Console Memory21
 - Restricting File Types for Attachments22
 - Adjusting Pattern Discovery22
 - Improving Annotation Query Performance23
 - Installing New License Files23
 - Configuring Manager Logging24
 - Sending Logs and Diagnostics to ArcSight Support26
 - Guidelines for Using the sendlogs Command27
 - Gathering Logs and Diagnostic Information27
 - Reconfiguring the ArcSight Console After Installation30
 - Reconfiguring ArcSight Manager30
 - Changing ArcSight Command Center Session Timeout30
 - Configuring Email for Transport Layer Security31

Managing Password Configuration	31
Enforcing Good Password Selection	31
Password Length	31
Restricting Passwords Containing User Name	32
Password Character Sets	32
Requiring Mix of Characters in Passwords	32
Checking Passwords with Regular Expressions	33
Password Uniqueness	34
Setting Password Expiration	34
Restricting the Number of Failed Log Ins	35
Disabling Inactive User Accounts	35
Re-Enabling User Accounts	36
Advanced Configuration for Asset Auto-Creation	36
Asset Auto-Creation from Scanners in Dynamic Zones	36
Create Asset with Either IP Address or Host Name	37
Preserve Previous Assets	38
Changing the Default Naming Scheme	39
Compressing SmartConnector Events	40
Reducing Event Fields with Turbo Modes	40
Monitoring ESM Appliance with SNMP	41
Sending Events as SNMP Traps	42
Configuration of the SNMP Trap Sender	42
Configuring Asset Aging	44
Excluding Assets from Aging	44
Disabling Assets of a Certain Age	45
Deleting an Asset	45
Amortize Model Confidence with Scanned Asset Age	46
Tuning for Supporting Large Actor Models	47
About Exporting Actors	47
Viewing License Tracking and Auditing Reports	48
Setting Up ESM for MSSP Environments	48
Setting up a Custom Login Message for ArcSight Console and Command Center	48
Setting Checkpoint Parameters	49
Preventing Rules Recovery Timeout	50
Enable Iframe of ArcSight Command Center Pages	50
Enabling Scaling for Bytes In and Bytes Out Event Fields	51

Converting an ESM Appliance to IPv6	52
Importing an Archive of 300MB Maximum Size	53
Customizing Product Image on Login Screen and Navigation Bar in the ArcSight Command Center	53
Changing the Hostname of Your Machine	54
Rule Actions Queue Full - Set rules.action.capacity Property	56
Chapter 3: Configuring and Managing a Distributed Correlation Cluster	57
Understanding the Cluster Services	57
Implementing a Distributed Correlation Cluster	58
Setting Up Key-Based Passwordless SSH	59
Set Up Key-Based Passwordless SSH	59
Verify Key-Based Passwordless SSH	60
Configuring Services in a Distributed Correlation Cluster	60
Configuring Message Bus Control and Message Bus Data Services	60
Adding Correlators and Aggregators	63
Configuring Correlators and Aggregators	64
Configuring Distributed Cache	65
Restoring the State of Distributed Cache Nodes	66
Configuring a Repository	67
Starting Distributed Correlation Services	69
Forwarding Events in Distributed Correlation Mode	69
Obtaining and Importing the Transformation Hub CA Certificate	70
Configuring the Filter and Destination	70
Modifying the Filter and Configuration	72
Troubleshooting Event Forwarding Throughput	73
Managing Distributed Correlation Services - Basic Commands	74
Start and Stop Order of Distributed Correlation Processes	76
Monitoring the Cluster Using the Cluster View Dashboard	78
Certificate-Based Admission of Services to a Cluster	79
Dynamic Ports in the Distributed Correlation Environment	79
Viewing Port Numbers for Dynamically Allocated Ports	80
Changing Authentication in a Distributed Correlation Environment	80
Changing Hostnames or IP Addresses in a Cluster	80
Changing the Internet Protocol Version in a Distributed Correlation Environment	85
Removing a Node from a Cluster	87

Troubleshooting and Frequently Asked Questions for Distributed Correlation Mode	88
What is compact mode?	88
How does distributed correlation mode differ from compact mode in terms of what I see in the ArcSight Console?	88
How do I verify that my system is running in distributed correlation mode?	88
How do I verify that the correlators, aggregators, and other services in my cluster are working?	88
Why are services stopped or in a prolonged initializing state on non-persistor nodes?	89
If I make a change to a system in the cluster, do I have to synchronize the cluster to make that change known throughout the cluster?	89
Why can't I start a service from a non-persistor node?	89
Are there special memory requirements for the cluster?	90
How are audit events generated for distributed correlation services?	90
Will content behave differently in distributed correlation mode?	90
Where do queries run in distributed correlation mode?	90
How does upgrade work for distributed correlation?	90
I changed the friendly name for one of my correlators. Why is the change not reflected in audit events?	90
 Chapter 4: SSL Authentication	 91
SSL Authentication Terminology	91
Understanding Cipher Suites	93
How SSL Works	94
Certificate Types	95
SSL Certificate Tasks	96
Export a Key Pair	98
Exporting a Key Pair Using bin/arc sight keytool	98
Exporting a Key Pair Using keytoolgui	98
Import a Key Pair	99
Importing a Key Pair Using bin/arc sight keytool	99
Importing a Key Pair Using keytoolgui	99
Export a Certificate	100
Exporting a Certificate Using bin/arc sight keytool	100
Exporting a Certificate Using keytoolgui	100
Import a Certificate	101
Importing a Certificate Using bin/arc sight keytool	101
Importing a Certificate Using keytoolgui	101

Creating a Keystore	102
Creating a Keystore Using jre/bin/keytool	102
Creating a Keystore Using keytoolgui	103
Generating a Key Pair	104
Generating a Key Pair Using bin/arcsight keytool	104
Generating a Key Pair Using keytoolgui	105
View Certificate Details From the Store	106
Viewing a Certificate Details from the Store Using bin/arcsight keytool	106
Viewing a Certificate Details from the Store Using keytoolgui	106
Delete a Certificate	107
Deleting a Certificate Using bin/arcsight keytool	107
Deleting a Certificate Using keytoolgui	107
Changing Keystore/Truststore Passwords	107
Using a Self-Signed Certificate	108
When Clients Communicate With One Manager	108
When Clients Communicate With Multiple Managers	110
Using a CA-Signed SSL Certificate	112
Create a Key Pair for a CA-Signed Certificate	113
Send for the CA-Signed Certificate	114
Sending a CA-Signed Certificate Using keytool	114
Sending a CA-Signed Certificate Using keytoolgui	114
Import the CA Root Certificate	115
Import the CA-Signed Certificate	115
Start the Manager Again (Restart the Manager)	118
Using CA-Signed Certificates with Additional Components	119
Removing a Demo Certificate	119
Replacing an Expired Certificate	120
Establishing SSL Client Authentication	120
Setting up SSL Client-Side Authentication on ArcSight Console- Self-Signed Certificate	122
Setting up SSL Client-Side Authentication on ArcSight Console- CA-Signed Certificate	122
Setting Up Client-Side Authentication for ArcSight Command Center	123
Setting Up Client-Side Authentication on SmartConnectors	124
Setting Up Client-Side Authentication for Utilities on the ESM Server	125
SSL Authentication - Migrating Certificate Types	126
Migrating from Demo to Self-Signed	127
Migrating from Demo to CA-Signed	127
Migrating from Self-Signed to CA-Signed	127

Verifying SSL Certificate Use	127
Sample Output for Verifying SSL Certificate Use	128
Using Certificates to Authenticate Users to the Manager	128
Using the Certificate Revocation List (CRL)	129
Chapter 5: Running the Manager Configuration Wizard	131
Running the Wizard	131
Authentication Details	135
How External Authentication Works	135
Guidelines for Setting Up External Authentication	135
Password-Based Authentication	136
Built-In Authentication	136
Setting up RADIUS Authentication	136
Setting up Active Directory User Authentication	137
Configuring AD SSL	138
Setting up LDAP Authentication	138
Configuring LDAP SSL	139
Password Based and SSL Client Based Authentication	139
Password Based or SSL Client Based Authentication	139
SSL Client Only Authentication	139
One SSO Provider (OSP) Authentication	140
OSP Client Only Authentication	140
External SAML2 Client Only Authentication	140
Appendix A: Administrative Commands	143
ArcSight_Services Command - Compact Mode	143
ArcSight_Services Command - Distributed Correlation Mode	145
ArcSight Commands	147
ACLReportGen	148
agent logfu	148
agent tempca	149
agentcommand	149
agents	149
agentsvc	150
agentup	150
aggregatorthreaddump	150
arcdt	151
archive	153

archivefilter	161
bleep	162
bleepsetup	163
changepassword	163
checklist	164
certadmin	164
configure-event-forwarding	166
console	166
consolesetup	167
correlationsetup	168
correlatorthreaddump	168
dcachesetup	168
deploylicense	168
downloadcertificate	169
exceptions	170
export_system_tables	171
flexagentwizard	172
groupconflictingassets	172
import_system_tables	173
keytool	174
keytoolgui	175
kickbleep	176
listsubjectdns	176
logfu	176
managerinventory	179
manager-reload-config	180
managersetup	181
managerthreaddump	182
managerup	183
mbussetup	183
monitor	183
netio	184
package	184
portinfo	186
reenableuser	187
refcheck	187
regex	188
replayfilegen	188
reposeup	188
resetpwd	189
resvalidate	189

searchindex	190
sendlogs	191
syncpreferip	192
tee	193
tempca	193
threaddumps	194
tproc	194
updaterepohostconfig	195
whois	195
zoneUpdate	196
CORR-Engine ArcSight Commands	199
Appendix B: Troubleshooting	202
General Troubleshooting	202
Pattern Discovery Performance Troubleshooting	207
Query and Trend Performance Tuning Troubleshooting	207
SmartConnectors Troubleshooting	209
ArcSight Console Troubleshooting	209
Manager Troubleshooting	212
CORR Engine Troubleshooting	214
SSL Troubleshooting	215
Appendix C: Event Data Transfer Tool	218
ESM and Hadoop - Benefits	218
Setting Up the Event Data Transfer Tool	219
Using the Event Data Transfer Tool Command	219
Event Data Transfer Tool Usage Notes	221
File Names	221
Threads	221
Data Compression	222
Transfer Failures	222
Transfer Performance	222
Size of Transferred Files	223
Column Names	223
Appendix D: Creating Custom E-mails Using Velocity Templates	224

Notification Velocity Templates - Example	224
Velocity Template #if statement	224
Using Email.vm and Informative.vm	225
Understanding the Customization Process	226
Customizing the Template Files	227
Velocity Template Sample Output	227
Appendix E: Configuration Changes Related to FIPS	230
FIPS Encryption Cipher Suites	230
Key Pair Types Used in FIPS Mode	231
Enabling Use of CA-Signed Certificates Over Port 9000 in FIPS Mode	231
Configuring ESM to Use TLS 1.1 or 1.0	232
Generating a New Key Pair When Changing a Manager Hostname for FIPS Mode	233
Changing a Default Mode Installation to FIPS 140-2	234
Manager	235
ArcSight Console	236
Connectors	237
Changing Keystore/Truststore Passwords in FIPS Mode	237
Configure Your Browser for FIPS	238
Send Documentation Feedback	239

About this PDF Version of Online Help

This document is a PDF version of the online help. This PDF file is provided so you can easily print multiple topics from the help information or read the online help in PDF format. Because this content was originally created to be viewed as online help in a web browser, some topics may not be formatted properly. Some interactive topics may not be present in this PDF version. Those topics can be successfully printed from within the online help.

Chapter 1: Starting and Stopping the Manager and Components

Start the Manager from a command or console window. The remainder of this section provides more information about command line options to start, shut down, configure, or reconfigure ESM components.

Restarting the Manager - Stop the Manager and Start All Services

To restart the Manager, as user *arcsight*, you must stop the Manager, and then start all services.

To stop the Manager and start all services:

1. Stop the Manager:

```
/etc/init.d/arcsight_services stop manager
```

2. Start all services:

```
/etc/init.d/arcsight_services start all
```

Note: If your system is in distributed correlation mode, refer to the *ESM Release Notes* topic "Stop and Start All Services if a Major Service is Stopped" for special instructions.

Starting the ArcSight Command Center

To start the Command Center from a supported browser enter the following URL:

```
https://<hostname>:8443/
```

Where **<hostname>** is the host name or IP address of the Manager that you specified when you first configured ESM.

Starting ArcSight SmartConnectors

This procedure is only for SmartConnectors that are *not* running as a service. Before you start ArcSight SmartConnectors, make sure the Manager is running. It's also a good idea for the ArcSight Console to also be running, so that you can see the status of the configured SmartConnectors and view messages as they appear on the Console.

To start up an ArcSight SmartConnector:

1. Open a command window and navigate to the connector's `/current/bin` directory.
2. Type in the following line and press **Enter**:

```
./arcsight agents (on Linux)  
arcsight agents (on windows)
```

The connector in that folder starts.

Stopping and Starting ArcSight Services

Before performing tasks such as rebooting the server or installing a patch, you must stop ArcSight services. Performing a clean shutdown of services in this way will ensure the integrity of your ESM databases.

To stop ArcSight services, run the following command as the user *arcsight*:

```
/etc/init.d/arcsight_services stop all
```

To start ArcSight services, run the following command as the user *arcsight*:

```
/etc/init.d/arcsight_services start all
```

Important: In distributed correlation mode, you must run:

```
/etc/init.d/arcsight_services start all
```

after a system reboot to start **all services on all cluster nodes**. If you do not do so, only the persistor node-related services will start; services on other cluster nodes will not start automatically in this case.

Starting the ArcSight Console

To start up the ArcSight Console:

1. Open a command window on `<ARCSIGHT_HOME>/bin`.
2. Type in the following line and press **Enter**.

```
./arcsight console (on Linux)  
arcsight console (on Windows)
```

Reconnecting ArcSight Console to the Manager

If the ArcSight Console loses its connection to the Manager (because the Manager was started again, for example) a dialog box appears in the ArcSight Console stating that your connection to the Manager has been lost. Wait for the Manager to finish starting, if applicable. Click **Retry** to re-establish a connection to the Manager or click **Relogin**.

Note: The connection to the Manager cannot be re-established while the Manager is starting. In some cases, a connection cannot be established without resetting one or both machines.

Clicking **Retry** may display connection exceptions while the Manager is starting again, or as the connection is re-established.

Chapter 2: Basic Configuration Tasks

References to ARCSIGHT_HOME

<ARCSIGHT_HOME> in the paths represents:

- /opt/arcsight/manager for the ArcSight Manager
- Whatever path you specified when you installed the ArcSight Console
- Whatever path you specified when you installed an ArcSight SmartConnector.

Managing and Changing Properties File Settings

Various components use properties files for configuration. Many sections of this documentation require you to change properties in those files. Some of the properties files are also modified when you use one of the configuration wizards.

Property File Format

Properties files are text files containing pairs of keys and values. The keys specify the setting to configure. For example, the following property configures the port on which the Manager listens:

```
servletcontainer.jetty311.encrypted.port=8443
```

Blank lines and lines that start with a pound sign (#) are ignored. Use the pound sign for comments.

Defaults and User Properties

Most properties files come in pairs. The first is the defaults properties file, such as `server.defaults.properties`. It contains the default settings. Do not modify these files; use them as a reference. They are overwritten upon upgrade.

The second file is the user properties file, such as `server.properties`. It can contain any properties from the defaults properties file, but the property values in this file override those in the defaults file. Thus, it contains settings that are specific to a particular

installation. Typically, the user properties file for a component is created and modified automatically when you configure the component using its configuration wizard.

Because the user properties file contains settings you specify to suit your environment, it is never replaced by an upgrade. If an upgrade, such as a service pack or a version update, changes any properties, it does so in the `defaults` file.

The following table lists the most important properties files. The paths are relative to `<ARCSIGHT_HOME>`.

Default Properties	User Properties	Purpose
<code>config/server.defaults.properties</code>	<code>config/server.properties</code>	Manager Configuration
<code>config/esm.defaults.properties</code>	<code>config/esm.properties</code>	Cluster configuration properties and SSL properties common to persistor, correlator, and aggregator services on the node. This properties file is present on each node in a distributed correlation cluster.
<code>config/console.defaults.properties</code>	<code>config/console.properties</code>	ArcSight Console Configuration
<code>config/client.defaults.properties</code>	<code>config/client.properties</code>	ArcSight Common Client Configuration
<code>/opt/arcsight/logger/current/arcsight/logger/config/logger.defaults.properties</code>	<code>/opt/arcsight/logger/userdata/logger/user/logger.properties</code>	Features exposed on the ArcSight Command Center

Editing Properties Files

When you edit a properties file, copy the property to edit from the `*.defaults.properties` to `*.properties` and change the setting to your new value in `*.properties`. When you install an upgrade, and the `*.defaults.properties` file is updated, the properties you customized in `*.properties` remain unchanged.

You can edit the properties using any text editor. Make sure you use one that does not add any characters such as formatting codes.

If you configured the Console and SmartConnectors using default settings in the configuration wizard, a user properties file is not created automatically for that component. If you need to override a setting on such a component, use a text editor to create this file in the directory specified in the above table.

When you edit a property on a component, you must restart the component for the new values to take effect except for the dynamic Manager properties listed in the next section.

If you change a communication port, be sure to change both sides of the connection. For example, if you configure a Manager to listen to a different port than 8443, be sure to configure all the Manager's clients (Consoles, SmartConnectors, and so on) to use the new port as well.

Protocol	Port	Configuration
ICMP	none	ArcSight Console to Target communication (ping tool)
UDP	1645 or 1812	Manager to RADIUS server (if enabled)
	9090	ESM Service Layer Container Port
	9000	Used by the Manager for peering.
TCP	8443	SmartConnectors, ArcSight Command Center, and ArcSight Console to Manager communication
TCP	636	Manager to LDAP server (with SSL if enabled)
TCP	389	Manager to LDAP server (without SSL if enabled)
TCP	143	Manager to IMAP server (for Notifications)
TCP	110	Manager to POP3 server (for Notifications)
UDP/TCP	53	ArcSight Console to DNS Server communication (<code>nslookup</code> tool)
UDP/TCP	43	ArcSight Console to Whois Server communication (<code>whois</code> tool)
TCP	25	Manager to SMTP server (for Notifications)

Dynamic Properties

When you change the following properties in the `server.properties` file on the Manager, you do not need to start the Manager again for the changes to take effect:

- `auth.auto.reenable.time`
- `auth.enforce.single.sessions.console`
- `auth.enforce.single.sessions.web`
- `auth.failed.max`
- `auth.password.age`
- `auth.password.age.exclude`
- `auth.password.different.min`
- `auth.password.length.max`
- `auth.password.length.min`
- `auth.password.letters.max`
- `auth.password.letters.min`

- `auth.password.maxconsecutive`
- `auth.password.maxoldsubstring`
- `auth.password.numbers.max`
- `auth.password.numbers.min`
- `auth.password.others.max`
- `auth.password.others.min`
- `auth.password.regex.match`
- `auth.password.regex.reject`
- `auth.password.unique`
- `auth.password.userid.allowed`
- `auth.password.whitespace.max`
- `auth.password.whitespace.min`
- `external.export.interval`
- `notification.aggregation.max_notifications`
- `process.execute.direct`
- `servletcontainer.jetty311.log`
- `servletcontainer.jetty311.socket.https.expirationwarn.days`
- `ssl.debug`
- `whine.notify.emails`
- `xmlrpc.accept.ips`

After you make the change, you use the `manager-reload-config` command to load those changes to the Manager. Every time the `manager-reload-config` command is successful, a copy of the `server.properties` file it loaded is placed in `<ARCSIGHT_HOME>/config/history` for backup purposes. The `server.properties` file in `<ARCSIGHT_HOME>/config/history` is suffixed with a timestamp and does not overwrite the existing versions, as described in the following example.

Example

Manager M1 starts successfully for the first time on September 26, 2013, at 2:45 p.m. A backup copy of its `server.properties` file is written to `<ARCSIGHT_HOME>/config/history` with this timestamp:

```
server.properties.2013_09_26_14_45_27_718
```

On September 27, 2013, the M1 administrator adds the following property to the `server.properties` file:

```
notification.aggregation.max_notifications=150
```

When the administrator runs the `manager-reload-config` command at 1:05 p.m. the same day, it runs successfully because this property can be loaded dynamically.

As soon as the updated `server.properties` file is loaded in M1's memory, a backup copy of the updated `server.properties` file is written to `<ARCSIGHT_HOME>/config/history` with appropriate timestamp.

Now, `<ARCSIGHT_HOME>/config/history` contains these two backup files:

```
server.properties.2013_09_26_14_45_27_718
```

```
server.properties.2014_09_27_01_05_40_615
```

On September 28, 2014, the M1 administrator adds this property to the `server.properties` file:

```
notification.aggregation.time_window=2d
```

As this property can be also loaded dynamically, similar to the previous change, after the updated `server.properties` is loaded in M1's memory, a backup copy of the `server.properties` file is written to `<ARCSIGHT_HOME>/config/history` with appropriate timestamp.

Now, `<ARCSIGHT_HOME>/config/history` contains these three backup files:

```
server.properties.2014_09_26_14_45_27_718
```

```
server.properties.2014_09_27_01_05_40_615
```

```
server.properties.2014_09_28_03_25_45_312
```

On September 30, 2014, the M1 administrator updates the `log.channel.file.property.maxsize` property in the `server.properties` file. When the administrator runs the `manager-reload-config` command, the command fails because this property cannot be loaded dynamically. As a result, these things happen:

- The updated `server.properties` file is not loaded into M1's memory, however, changes made to it are not reverted.
- M1 continues to use the properties that were loaded on September 29th.
- No backup copy is made. The `<ARCSIGHT_HOME>/config/history` directory continues to contain the same three backup files:

```
server.properties.2014_09_26_14_45_27_718
```

```
server.properties.2014_09_27_01_05_40_615
```

```
server.properties.2014_09_28_03_25_45_312
```

The changes made on September 30th are not effective until M1 is restarted.

Changing Manager Properties Dynamically

To change any of the properties listed previously, do these steps:

1. Change the property in the `server.properties` file and save the file.
2. **(Optional)** Use the `-diff` option of the `manager-reload-config` command to view the difference between the server properties the Manager is currently using and the properties loaded after you run this command:

```
arcsight manager-reload-config -diff
```

Note: The `-diff` option compares all server properties—default and user properties. For all options available with the `manager-reload-config` command, see ["manager-reload-config" on page 180](#).

3. Run this command in `<ARCSIGHT_HOME>/bin` to load the new property values:

```
arcsight manager-reload-config
```

If this command fails with a warning, it means you are changing properties that require a Manager restart. In that case, none of the property changes are applied, including ones that do not require a restart. You can:

- Revert changes to properties that require you to stop the Manager and then start all services and rerun the `manager-reload-config` command.
- Force an update of all properties using the `-as` option, as follows:

```
arcsight manager-reload-config -as
```

When you use the `-as` option, the properties that can be changed without starting the Manager take effect immediately. The properties that require the Manager started again are updated in the `server.properties` but are not effective until the Manager is started.

For example, if you change `auth.password.length.min` to 7 and `search.enabled` to false, you get the above warning because only `auth.password.length.min` can be updated without starting the Manager. If you force an update of the `server.properties` file, `auth.password.length.min` is set to 7, but `search.enabled` continues to be set to true until the Manager is started.

Note: Be careful in using the `-as` option to force reload properties. If an invalid static change is made, it may prevent the Manager from starting up after it reboots.

Changing the Service Layer Container Port

By default the service layer container port is 9090. You can change this port:

1. Modifying the following files located in the Manager's <ARCSIGHT_HOME>:
 - <ARCSIGHT_HOME>/arcsight-dm/plugins/com.arcsight.dm.plugins.tomcatServer/conf/server.xml
 - <ARCSIGHT_HOME>/config/proxy.rule.xml
 - <ARCSIGHT_HOME>/config/rewriteProxy.rule.xmlMake sure to replace the references to port 9090 with an unused port number.
2. Stop the Manager by running the following command as user *arcsight*:

```
/etc/init.d/arcsight_services stop manager
```
3. Start all services:

```
/etc/init.d/arcsight_services start all
```

Securing the Manager Properties File

The Manager's `server.properties` file contains sensitive information such as database passwords, keystore passwords, and so on. Someone accessing the information in this file can do a number of things, such as tampering with the database and acting as a Manager. Protect the `server.properties` file so that only the user account under which the Manager is running is able to read it. For example, in Unix you can use the `chmod` command:

```
chmod 600 server.properties
```

This operation is performed during the Manager installation. As a result, only the owner of the file, which must be the user that runs the Manager, may read or write to the file. For all other users, access to the file is denied.

Adjusting Console Memory

Because the ArcSight Console can open up to ten independent event-viewing channels, out-of-memory errors may occur. If such errors occur, or if you simply anticipate using numerous channels for operations or analysis, please make the following change to each affected Console installation.

In the `bin/scripts` directory, edit the `console.bat` (Windows) or `console.sh` (Linux) and modify the `ARCSIGHT_JVM_OPTIONS -Xms` and/or the `-Xmx` parameters for the memory usage range of the Java Virtual Machine.

Restricting File Types for Attachments

In the ArcSight Console and the Command Center, users can attach files to cases. To prevent users from accidentally attaching malicious files, Administrators can restrict the types of files users are allowed to add by adding or modifying the following key in the `server.properties` file:

```
file.attachments.unacceptable.extensions
```

Separate multiple file types with a comma.

For example to prevent users from uploading files with `.exe`, `.bin`, or `.sh` extensions, configure the key as follows:

```
file.attachments.unacceptable.extensions=exe,bin,sh
```

In addition to attaching files to cases, this setting also affects:

- ArcSight files in the Command Center
- File resources

After updating the `server.properties` file, you must restart the Manager service.

Note: Changing the key does not affect any currently attached files.

Adjusting Pattern Discovery

Note: Pattern Discovery is not supported on ESM on an appliance.

By default, Pattern Discovery limits its memory usage to about 4 GB of memory. However, if the search for patterns involves too many transactions and events, the task can run out of memory and abort. To control the memory limit indirectly, change the maximum number of transactions and events the Pattern Discovery task can hold in memory. The settings for these values are in the `server.defaults.properties` file in the `config` folder. Place the changed versions in the `server.properties` file to supersede the default.

- `patterns.transactionbase.max`: The maximum transactions allowed in memory. If you exceed this, these transactions are stored as a page file. The default is 10000.
- `patterns.maxSupporterCost`: The maximum supporters allowed in memory. If you exceed this number, the Pattern Discovery task aborts. The default is 80000.
- `patterns.maxUniqueEvents`: The maximum unique events allowed in memory. If you exceed this number, the Pattern Discovery task aborts. The default is 20000.
- `patterns.timeSpreadCalculation`: Set to false avoid calculating timespread statistics,

which can take a lot of resources. If you experience performance issues while "Extracting Pattern for Snapshot," try scheduling Pattern Discovery for off-peak times. If you run Pattern Discovery against millions of matched events, try reducing the time frame to half to see how long it takes to complete. Use that information to plan when to run it. You can also make the filter condition more granular so there are fewer matches.

If the Pattern Discovery task aborts, a message to that effect appears in the console. Run the Pattern Discovery task again after increasing the Pattern Discovery memory usage limits. To increase the memory usage limit increase the three values proportionally. For example, to add 25 percent more memory capacity, you would change the values to:

- `patterns.transactionbase.max=12500`
- `patterns.maxSupporterCost=100000`
- `patterns.maxUniqueEvents=25000`

After changing these values, stop the manager and start all services for the new values to take effect:

1. Stop the Manager by running the following command as user *arcsight*:

```
/etc/init.d/arcsight_services stop manager
```

2. Start all services:

```
/etc/init.d/arcsight_services start all
```

Improving Annotation Query Performance

If you have annotation queries, their performance can be improved by adding the following property to the Manager's `server.properties` file:

```
event.annotation.optimization.enabled=true
```

You can edit the properties file using a regular text editor. After adding this property, stop and start the manager for it to take effect.

Installing New License Files

You receive new license files as `.xml` files that are sent through email. To deploy the new license file and replace the old license file, use the `deploylicense` command. For more information, see [deploylicense](#).

Configuring Manager Logging

The Manager writes logging information to log files, which by default are located in:

```
/opt/arcsight/var/logs/manager/default
```

Various Manager utilities write logging information to different sets of log files. Each of which can consist of multiple files.

The number and size of log files are configurable, a typical setting is 10 files with 10 megabytes each. When a log file reaches a maximum size, it is copied over to a different location. Depending on your system load, you may have to change the default settings. To make changes to the logging configuration, change the log channel parameters. The default log channel is called *file*.

For the main Manager log file, called `server.log`, the following `server.properties` settings are used:

```
# Maximum size of a log file.  
log.channel.file.property.maxsize=10MB  
  
# Maximum number of roll over files.  
log.channel.file.property.maxbackupindex=10
```

The first setting affects the size of each individual log file; the second affects the number of log files created. The log file currently in use is always the one with no number appended to the name. The log file with the largest number is the oldest. All log files are written to the

```
/opt/arcsight/var/logs/manager/default directory.
```

The Manager and its related tools write the following log files:

Log File	Description
<code>server.log*</code>	The main Manager log.
<code>server.status.log*</code>	System status information, such as memory usage.
<code>server.channel.log*</code>	Active Channel logs.
<code>server.std.log*</code>	All output that the Manager prints on the console (if run in command line mode)
<code>server.pulse.log*</code>	The Manager writes a line to this set of logs every ten seconds. Used to detect service interruptions.
<code>server.sql.log*</code>	If database tracing is enabled, the SQL statements are written to this set of log files.

Log File	Description
execproc.log*	Log information about externally executed processes (only on some platforms)
serverwizard.log*	Logging information from the <code>managersetup</code> command.
archive.log*	Logging information from the <code>arcsight archive</code> utility, which works with the Oracle database.
Correlator	
correlator.log*	The main correlator log.
correlatorpulse.log*	The correlator writes a line to this set of logs every ten seconds. Used to detect service interruptions.
correlator.status.log*	Correlator status information, such as memory usage.
velocity.log	Information about the velocity template engine.
Aggregator	
aggregator.log*	The main aggregator log.
aggregator.pulse.log*	The aggregator writes a line to this set of logs every ten seconds. Used to detect service interruptions.
aggregator.status.log*	Aggregator status information, such as memory usage.
velocity.log	Information about the velocity template engine.
zookeeper.log	The main ZooKeeper log.
zookeeper.std.log	ZooKeeper status information, such as memory usage.
Distributed Cache	
dcache.log	The main dcache log.
dcache.std.log	Dcache status information, such as memory usage.
Repository	
repo.log	Miscellaneous information output by ZooKeeper.
repo.std.log	ZooKeeper status information.
Message Bus Data	
controller.log	Information about the cluster's management activity.
kafka-authorizer.log	Information about secure connections and access control. Because ESM does not use these features, the log might be empty.
kafka.log	Information about Kafka broker activity.
kafka-request.log	Data about each request received by the Kafka broker.
kafka.std.log	Kafka status information, such as memory usage.

Log File	Description
log-cleaner.log	Information about Kafka's log compaction.
state-change.log	Information about cluster resource state changes.
Message Bus Control	
zookeeper.log	Information about ZooKeeper server activity.
zookeeper.std.log	ZooKeeper status information, such as memory usage.

Logs for distributed correlation services are located in:

Distributed Correlation Service	Log Location
persistor	/opt/arcsight/var/logs
correlators	/opt/arcsight/var/logs/<correlator_serviceId>
aggregators	/opt/arcsight/var/logs/<aggregator_serviceId>
message bus control	/opt/arcsight/var/logs/<mbus_control_serviceId>
message bus data	/opt/arcsight/var/logs/<mbus_data_serviceId>
distributed cache	/opt/arcsight/var/logs/<dcache_serviceId>
repository	/opt/arcsight/var/logs/<repo_serviceId>

Sending Logs and Diagnostics to ArcSight Support

Customer Support may request log files and other diagnostic information to troubleshoot problems. You can use the Log Retrieval feature in ArcSight Command Center. Check the online help for that feature for more information.

In the Console, the `sendlogs` command automatically locates the log files and compresses them. You can send the compressed files to Customer Support. For details on the `sendlogs` command, see ["sendlogs" on page 191](#).

- You can run this command as a wizard directly from the Console interface (GUI) in addition to the command-line interface of each component.
- Optionally, gather diagnostic information such as session wait times, thread dumps, and database alert logs about your ESM system, which helps Customer Support analyze performance issues on your ESM components.
- When you run this command from the Console or Manager, you can gather logs and diagnostic information for all components of the system.

Guidelines for Using the `sendlogs` Command

When using the `sendlogs` command:

- You can be connected as any valid user on an ESM component to collect its local logs; however, you must have administrator access to collect logs from other components. For example, if you are connected as user 'joe' to the Console, you can collect its logs. But if you need to collect logs for the Manager and the database, you must connect to the Console as the administrator.
- You can only collect local logs on SmartConnectors or the CORR-Engine. The Send Logs utility only collects logs for the component on which you run it. In order to collect the CORR-Engine logs, the Manager needs to be running.
- All log files for a component are gathered and compressed. That is, you cannot select a subset of log files that the utility should process.
- The `sendlogs` command generates a compressed file on your local system that you can send to Customer Support by e-mail, if they request it.
- You can review the compressed file to ensure that only a desired and appropriate amount of information is sent to support.
- You can remove or sanitize information such as IP addresses, host names, and e-mail addresses from the log files before compressing them. The options are:
 - Send log as generated
This default option does not remove any information from the log files.
 - Only remove IP address
This option removes IP addresses, but not host names or e-mail addresses, from the log files.
 - Remove IP address, host names, e-mail addresses
This option removes all IP addresses and enables you to specify a list of host-name suffixes for which all host names and e-mail addresses are removed from the log files.

For example, if you specify 'company.com' as a host-name suffix to remove, the Send Logs utility removes all references to domains such as 'www.company.com' and e-mail addresses such as 'john@company.com' from the logs.

Gathering Logs and Diagnostic Information

When you run the `sendlogs` command on SmartConnectors, it gathers logs and diagnostic information (if applicable) for only those components. However, when you run this utility on ArcSight Console or Manager, you can gather logs and diagnostic information for all or a selected set of ESM components.

To run this command on SmartConnectors, enter this in `<ARCSIGHT_HOME>/bin`:

```
./arcsight agent sendlogs
```

To gather logs and diagnostic information for all or a selected set of components, do one of the following:

- On the ArcSight Console, click **Tools > SendLogs**.
- Enter this command in <ARCSIGHT_HOME>/bin on the Console or Manager machine:

```
./arcsight sendlogs
```

The above action starts the Send Logs wizard. In the wizard screens, perform these steps:

Note: The Send Logs wizard remembers most of the choices you make when you run it for the first time. Therefore, for subsequent runs, if you choose to use the previous settings, you do not need to re-enter them.

1. Decide whether you want the wizard to gather logs only from the component on which you are running it or from all components.

Choose either **Use current settings to gather logs** or **Change/Review settings before gathering logs**.

If you select **Use current settings to gather logs** Logs for all components are gathered thus: If this is the first sendlogs is run after installation, then all the logs are gathered. If this is not the first time you have sendlogs has run, it uses the same setting as the previous run.

- a. Enter the Manager's login information.
- b. Go to the step [Sanitize logs](#).

If you select **Change/Review settings before gathering logs**, you can to select the components for which you want logs gathered.

Choose either **Local Logs Only** or **Logs from other components (Requires Manager credentials)**. These choices allow you to select whether you want only the local (the component from where you ran the sendlogs command) logs selected or to select logs from other components to be collected as well.

Local logs only:

If you select **Local logs only**, you can choose either **Include all time ranges** or **Choose a specific time range**.

If you select **Include all time ranges**, go to the step [Sanitize logs](#).

If you select **Choose a specific time range**, you are prompted to enter a **Start Time** and **End Time**, which is a time range for which the wizard gathers the logs.

Go to the step [Sanitize logs](#).

Logs from other components (Requires Manager credentials):

If you select **Logs from other components (Requires Manager credentials)**, you are prompted to choose the components.

- a. Select the components (for example, Manager, or Connectors) and the time range for which you want to gather logs. In addition, select whether you want to run the diagnostic utilities to gather additional information for those components.

If you choose to specify the diagnostic utilities to run, you are prompted to select the utilities from a list in a later screen. The diagnostic utilities you can select are described in ["arcdt" on page 151](#).

- b. If you chose to gather logs from the SmartConnectors, select those SmartConnectors in the next screen.
- c. If you chose to select the diagnostic utilities you want to run earlier in this wizard, select them in the next screen.

2. Sanitize logs

Select whether you want to sanitize the logs before collecting them. For more information about sanitizing options, see ["Guidelines for Using the sendlogs Command" on page 27](#).

If you choose **Do not sanitization logs (fastest)**, go to the step [Incident Number](#)

If you choose **Change/Review Logs sanitization settings**, you are prompted to select what you want to sanitize.

If you chose one of the first two options, go to the step [Incident Number](#).

If you selected **Remove IP addresses, host names, and e-mail addresses (Slowest)**, you are prompted to enter what you want removed. Click **Add** to add a suffix to remove. Highlight an entry and click **Remove** to remove it from the list.

3. Incident Number

Enter the Customer Support incident number.

The `sendlogs` command uses this number to name the compressed file it creates. Use the incident number that Customer Support gave you when you reported the issue for which you are sending the logs. Doing so helps Customer Support relate the compressed file to your incident.

In case you do not have an incident number at this time, you can continue by entering a meaningful name for the compressed file to be created. After you obtain the incident number from Customer Support, you can rename the file with the incident number you received.

4. Click **Next** to start the compression.

Note: Most of the values you entered during the first run of the Send Logs wizard are retained. The next time you run this wizard, you need to enter only a few settings.

5. Click **Done** on the final screen.

Reconfiguring the ArcSight Console After Installation

You can reconfigure ArcSight Console at anytime by typing `arcsight consolesetup` within a command window.

Run the ArcSight Console Configuration Wizard by entering the following command in a command window in the `<ARCSIGHT_HOME>/bin` directory:

```
./arcsight consolesetup
```

To run the ArcSight Console Setup program without the graphical user interface, type:

```
./arcsight consolesetup -i console
```

The ArcSight Console Configuration Wizard launches.

Reconfiguring ArcSight Manager

To reconfigure Manager settings made during installation, run the Manager Configuration Wizard. The Manager Configuration Wizard is covered in ["Running the Manager Configuration Wizard" on page 131](#).

To change advanced configuration settings (port numbers, database settings, log location, and so on) after the initial installation, change the `server.properties` file. ArcSight's default settings are listed in the `server.defaults.properties` file. You can override these default settings by adding the applicable lines from `server.defaults.properties` to the `server.properties` file. If a property exists in both the `server.defaults.properties` file and the `server.properties` file, the value in the `server.properties` file is used. These files are located in `<ARCSIGHT_HOME>/config`. Values in the `server.properties` file supersede those in `server.defaults.properties`.

Changing ArcSight Command Center Session Timeout

ArcSight Command Center will automatically log out if it has been inactive for a certain amount of time. This duration is defined by the configurable `service.session.timeout` property. The default timeout is 900 seconds (15 minutes). If the session duration is too short, increase the value set for the `service.session.timeout` property in the `<ARCSIGHT_HOME>/config/server.properties` file.

Configuring Email for Transport Layer Security

Note: ESM supports TLS only.

The `email.tls.desired` server property allows you to configure email for SMTP servers configured to use Transport Layer Security (TLS).

If your SMTP server is configured to use TLS, you do not need to do anything because, by default, this property is set to `true`.

If your SMTP server is not set to use TLS, then add the property `email.tls.desired=false` to the `server.properties` file. See ["Managing and Changing Properties File Settings" on page 15](#), for information on editing the `server.properties` file.

If the TLS configurations do not match:

- SMTP server uses TLS and `email.tls.desired=false`, emails are sent without TLS.
- SMTP server does not use TLS and `email.tls.desired=true`, emails are not sent.

If emails fail for any reason, they are not re-sent.

Managing Password Configuration

The Manager supports a rich set of functionality for managing users passwords. This section describes various password configuration options. Generally, all the settings are made by editing the `server.properties` file. See ["Managing and Changing Properties File Settings" on page 15](#). Some of these control character restrictions in passwords.

Enforcing Good Password Selection

There are a number of checks that the Manager performs when a user picks a new password in order to enforce good password selection practices.

Password Length

The simplest one is a minimum and, optionally, a maximum length of the password. The following keys in `server.properties` affect this:

```
auth.password.length.min=6
```

```
auth.password.length.max=20
```

By default, the minimum length for passwords is six characters and the maximum length is 20 characters and can contain numbers and/or letters.

Configuring the above properties to a value of -1 sets the password length to unlimited characters.

Restricting Passwords Containing User Name

Another mechanism that enforces good password practices is controlled through the following `server.properties` key:

```
auth.password.userid.allowed=false
```

When this key is set to false (the default), a user cannot include their user name as part of the password.

Password Character Sets

For appliance users, the Manager comes installed using the UTF-8 character set. If you install the Manager, it allows you to set the character set encoding that the Manager uses. When you install the ArcSight Console, the operating system on that machine controls the character set the Console uses. Be sure the operating system uses the same character set as the Manager if:

- A user password contains "non-English" characters (in the upper range of the character set: values above 127)
- That user wants to log in with that ArcSight Console.

This is not an issue if you log in from the web-based ArcSight Command Center.

For passwords that are in the ASCII range (values up to 127), the character set for the ArcSight Console does not matter.

Requiring Mix of Characters in Passwords

Strong passwords consist not only of letters, but contain numbers and special characters as well. This makes them more difficult to guess and can prevent dictionary attacks.

By default, the minimum length for passwords is six characters and the maximum length is 20 characters and can contain numbers and/or letters.

The following properties control the distribution of characters allowed in new passwords:

```
auth.password.letters.min=-1
```

```
auth.password.letters.max=-1
```

```
auth.password.numbers.min=-1
```

```
auth.password.numbers.max=-1
```



```
auth.password.whitespace.min=0
```

```
auth.password.whitespace.max=0
```

```
auth.password.others.min=-1
```

```
auth.password.others.max=-1
```

The `*.min` settings can be used to enforce that each new password contains a minimum number of characters of the specified type. The `*.max` settings can be used to limit the number of characters of the given type that new passwords can contain. Letters are all letters from A-Z, upper and lowercase, numbers are 0-9; "whitespace" includes spaces, etc.; "others" are all other characters, including special characters such as `#$%@!`.

Additionally, the following `server.properties` key lets you restrict the number of consecutive same characters allowed.

```
auth.password.maxconsecutive=3
```

For example, the default setting of 3 would allow "adam999", but not "adam9999" as a password.

Furthermore, the following `server.properties` key enables you to specify the length of a substring that is allowed from the old password in the new password.

```
auth.password.maxoldsubstring=-1
```

For example, if the value is set to 3 and the old password is "secret", neither "secretive" nor "cretin" is allowed as a new password.

Checking Passwords with Regular Expressions

To accommodate more complex password format requirements, the Manager can also be set up to check all new passwords against a regular expression. The following `server.properties` keys can be used for this purpose:

```
auth.password.regex.match=
```

```
auth.password.regex.reject=
```

The `auth.password.regex.match` property describes a regular expression that all passwords have to match. If a new password does not match this expression, the Manager rejects it. The `auth.password.regex.reject` property describes a regular expression that no password may match. If a new password matches this regular expression, it is rejected.

Note: Backslash (`\`) characters in regular expressions must be duplicated (escaped)—instead of specifying `\`, type `\\`.

For more information on creating an expression for this property, see <http://www.regular-expressions.info/>. The following are a few examples of regular expressions and a description of what they mean.

- `auth.password.regex.match= /^\d.*\d$/`

Only passwords that do not start or end with a digit are accepted.

- `auth.password.regex.match= ^(?=[A-Z].*[A-Z])(?=[a-z].*[a-z])(?=[0-9].*[0-9])(?=[^a-zA-Z0-9].*[^a-zA-Z0-9]).{10,}$`

Only passwords that contain at least 10 characters with the following breakdown are accepted:

- At least two upper case letters
 - At least two lower case letters
 - At least two digits
 - At least two special characters (no digits or letters)
- `auth.password.regex.reject= ^(?=[A-Z].*[A-Z])(?=[a-z].*[a-z])(?=[0-9].*[0-9])(?=[^a-zA-Z0-9].*[^a-zA-Z0-9]).{12,}$`

The passwords that contain 12 characters with the following breakdown are rejected:

- At least two upper case letters
- At least two lower case letters
- At least two digits
- At least two special characters (no digits or letters)

Password Uniqueness

In some environments, it is also desirable that no two users use the same password. To enable a check that ensures this, the following `server.properties` key can be used:

```
auth.password.unique=false
```

If set to true, the Manager checks all other passwords to make sure nobody is already using the same password.

Note: This feature may not be appropriate for some environments as it allows valid users of the system to guess other user's passwords.

Setting Password Expiration

The Manager can be set up to expire passwords after a certain number of days, forcing users to choose new passwords regularly. This option is controlled by the following key in `server.properties`:

```
auth.password.age=60
```

By default, a password expires 60 days from the day it is set.

When this setting is used, however, some problems arise for user accounts that are used for automated log in, such as the user accounts used for Manager Forwarding Connectors. These user accounts can be excluded from password expiration using the following key in `server.properties`:

```
auth.password.age.exclude=username1,username2
```

This value is a comma-separated list of user names. The passwords of these users never expire.

The Manager can also keep a history of a user's passwords to make sure that passwords are not reused. The number of last passwords to keep is specified using the following key in `server.properties`:

```
auth.password.different.min=1
```

By default, this key is set to check only the last password (value = 1). You can change this key to keep up to last 20 passwords.

Restricting the Number of Failed Log Ins

The Manager tracks the number of failed log in attempts to prevent brute force password guessing attacks. By default, a user's account is disabled after three failed log in attempts. This feature is controlled through the following key in `server.properties`:

```
auth.failed.max=3
```

Change this to the desired number or to `-1` if you do not wish user accounts to be disabled, regardless of the number of failed log in attempts.

After a user account has been disabled, the Manager can be configured to automatically re-enable it after a certain period of time. This reduces administrative overhead, while effectively preventing brute force attacks. This mechanism is controlled by the following key in `server.properties`:

```
auth.auto.reenable.time=10
```

This value specifies the time, in minutes, after which user accounts are automatically re-enabled after they were disabled due to an excessive number of incorrect log ins. Set the property key to `-1` to specify that user accounts can only be re-enabled manually.

Disabling Inactive User Accounts

By default, if a user does not log in for 90 days, the account is automatically disabled. To change the number of days of inactivity before the account is disabled, add the following

property to the `server.properties` file:

```
auth.user.account.age=<days>
```

Change `<days>` to the number of days of inactivity allowed before the account is disabled.

Re-Enabling User Accounts

Under normal circumstances, user accounts that have been disabled—for example, as a result of too many consecutive failed log ins—can be re-enabled by any user with sufficient permission. Check the **Login Enabled** check box for a particular user in the User Inspect/Editor panel in the ArcSight Console.

If the only remaining administrator user account is disabled, a command line tool can be run on the system where the Manager is installed to re-enable user accounts. First, ensure that the Manager is running. Then, from the command line, run the following commands as user *arcsight*:

```
cd <ARCSIGHT_HOME>/bin  
./arcsight reenabler <username>
```

where `username` is the name of the user you want to re-enable. After this procedure, the user can log in again, using the unchanged password.

Advanced Configuration for Asset Auto-Creation

Assets are automatically created for all components and, if applicable, for assets arriving from scan reports sent by vulnerability scanners via scanner SmartConnectors. This is done by the asset auto-creation feature.

If the profile of events in your network causes asset auto creation feature to create assets in your network model inefficiently, you can modify the asset auto creation default settings in the user configuration file, `server.properties`.

The `server.properties` file is located at `<ARCSIGHT_HOME>/config/server.properties`.

Asset Auto-Creation from Scanners in Dynamic Zones

The following properties relate to how assets are created from a vulnerability scan report for dynamic zones.

Create Asset with Either IP Address or Host Name

By default, an asset is not created in a dynamic zone if there is no host name present. The property set by default is:

```
scanner-event.dynamiczone.asset.nonidentifiable.create=false
```

You can configure ESM to create the asset as long as it has either an IP address or a host name. In `server.properties`, change `scanner-event.dynamiczone.asset.nonidentifiable.create` from `false` to `true`. ESM discards conflicts between an IP address and host name (similar IP address, but different host name and/or MAC address).

Caution: Creating an asset if no host name is present can result in an inaccurate asset model.

Setting `scanner-event.dynamiczone.asset.nonidentifiable.create` to `true` means that assets are created if the asset has an IP address or hostname.

This could lead to disabled assets or duplicated assets being created. Change this configuration only if you are using a dynamic zone to host ostensibly static assets, such as long-lived DHCP addresses.

When this property is set to `true`, the following takes place:

Example	Action taken if no conflicts	Action taken if previous asset with similar information
IP=1.1.1.1 hostname=myhost mac=0123456789AB	Asset created.	Asset created, previous asset is deleted.
ip=1.1.1.1 hostname=myhost mac=null	Asset created.	Asset created, previous asset is deleted.
ip=1.1.1.1 hostname=null mac=0123456789AB	Asset created.	Asset created, previous asset is deleted.
ip=1.1.1.1 hostname=null mac=null	Asset created.	Asset created, previous asset is deleted.

Example	Action taken if no conflicts	Action taken if previous asset with similar information
ip=null hostname=myhost mac=null	Asset created.	Previous asset deleted.
ip=null hostname=null mac=0123456789AB	Asset not created.	Asset not created.
ip=null hostname=myhost mac=0123456789AB	Asset created.	Previous asset deleted.

Preserve Previous Assets

This setting applies when ESM creates assets from a vulnerability scan report for dynamic zones. By default, if a previous asset with similar information already exists in the asset model, ESM creates a new asset and deletes the old one.

To preserve the previous asset rather than delete it when a scan finds a new asset with similar information, you can configure ESM to rename the previous asset. In `server.properties`, change `scanner-event.dynamiczone.asset.ipconflict.preserve` from `false` to `true`.

Caution: Preserving previous assets results in a larger asset model.

Setting `event.dynamiczone.asset.ipconflict.preserve` to `true` means that assets are continually added to the asset model and not removed. Use this option only if you know you must preserve all assets added to the asset model.

When the system is configured with `scanner-event.dynamiczone.asset.nonidentifiable.create=false` and `scanner-event.dynamiczone.asset.ipconflict.preserve=true`, it takes the following actions:

Example	Action taken if previous asset with similar information and preserve = true
IP=1.1.1.1 hostname=myhost mac=0123456789AB	Asset created, previous asset is renamed.
ip=1.1.1.1 hostname=myhost mac=null	Asset created, previous asset is renamed.
ip=1.1.1.1 hostname=null mac=0123456789AB	Asset created, previous asset is renamed.
ip=1.1.1.1 hostname=null mac=null	No asset created. Either host name or MAC address is required.
ip=null hostname=myhost mac=null	Asset created.
ip=null hostname=null mac=0123456789AB	Asset not created.
ip=null hostname=myhost mac=0123456789AB	Asset created.

Changing the Default Naming Scheme

By default, the system names assets that come from scanners using the naming scheme outlined in the topic "Asset Names" in the *ArcSight Console User's Guide*.

	Static Zone	Dynamic Zone
Property	scanner-event.auto-create.asset.name.template	scanner-event.auto-create.dynamiczone.asset.name.template
Value	\$destinationAddress - \$!destinationHostName	\$destinationHostName
Example	1.1.1.1 - myhost	myhost

You can reconfigure this naming scheme. For example, if you want the asset name for an asset in a static zone to appear this way in the ArcSight Console:

```
myhost_1.1.1.1
```

In this case, change the default

```
$destinationAddress - $!destinationHostName
```

to

```
$!destinationHostName_$destinationAddress
```

Compressing SmartConnector Events

ArcSight SmartConnectors can send event information to the Manager in a compressed format using HTTP compression. The compression technique used is standard GZip, providing compression ratio of 1:10 or higher, depending on the input data (in this case, the events the ArcSight SmartConnector is sending). Using compression lowers the overall network bandwidth used by ArcSight SmartConnectors dramatically, without impacting their overall performance.

By default, all ArcSight SmartConnectors have compression enabled. To turn it off, add the following line to the <ARCSIGHT_HOME>/user/agent/agent.properties file:

```
compression.enabled = false
```

ArcSight SmartConnectors determine whether the Manager they are sending events to supports compression.

Reducing Event Fields with Turbo Modes

If your configuration, reporting, and analytic usage permits, you can accelerate the transfer of sensor information through SmartConnectors by choosing one of the "turbo" modes, which send fewer event fields from the connector. The default transfer mode is called Complete, which passes all the data arriving from the device, including any additional data (custom, or vendor-specific).

ArcSight SmartConnectors can be configured to send more or less event data, on a per-SmartConnector basis, and the Manager can be set to read and maintain more or less event data, independent of the SmartConnector setting. Some events require more data than others. For example, operating system syslogs often capture a considerable amount of environmental data that may or may not be relevant to a particular security event. Firewalls, on the other hand, typically report only basic information.

ESM defines the following Turbo Modes:

	Turbo Modes	
1	Fastest	Recommended for firewalls
2	Faster	Manager default

When Turbo Mode is not specified (mode 3, Complete), all event data arriving at the SmartConnector, including additional data, is maintained. Turbo Mode 2, Faster, eliminates the additional custom or vendor-specific data, which is not required in many situations. Turbo Mode 1, Fastest, eliminates all but a core set of event attributes, in order to achieve the best throughput. Because the event data is smaller, it requires less storage space and provides the best performance. It is ideal for simpler devices such as firewalls.

The Manager processes event data using its own Turbo Mode setting. If SmartConnectors report more event data than the Manager needs, the Manager ignores the extra fields. On the other hand, if the Manager is set to a higher Turbo Mode than a SmartConnector, the Manager maintains fields that are not filled by event data. Both situations are normal in real-world scenarios, because the Manager configuration reflects the requirements of a diverse set of SmartConnectors.

Event data transfer modes are numbered (1 for Fastest, 2 for Faster, 3 for Complete), and possible Manager-SmartConnector configurations are therefore:

- 1-1 Manager and SmartConnector in Fastest mode
- 1-2 SmartConnector sending more sensor data than Manager needs
- 1-3 SmartConnector sending more sensor data than Manager needs
- 2-1 SmartConnector not sending all data that Manager is storing*
- 2-2 Manager and SmartConnector in Faster mode
- 2-3 Default: Manager does not process additional data sent by SmartConnector
- 3-1 Manager maintains Complete data, SmartConnector sends minimum*
- 3-2 Manager maintains additional data, but SmartConnector does not send it
- 3-3 Manager and SmartConnector in Complete mode

*When the SmartConnector sends minimal data (Turbo Mode 1), the Manager can infer some additional data, creating a 2-1.5 or a 3-1.5 situation.

Monitoring ESM Appliance with SNMP

We now provide the necessary snmp packages on the appliance so that you can set up SNMP monitoring.

By default net-snmp comes set up using the community string *public*, and will work right out of the box using that community string.

If you would like to change the configuration to make it more secure, edit the `/etc/snmp/snmpd.conf` file. All the configuration about net-snmp goes in that file.

Sending Events as SNMP Traps

ESM can send a sub-stream of all incoming events (that includes rule-generated events) via SNMP to a specified target. A filter is used to configure which events are sent. ESM's correlation capabilities can be used to synthesize network management events that can then be routed to your enterprise network management console.

Configuration of the SNMP Trap Sender

The SNMP trap sender is configured using the Manager configuration file. The `<ARCSIGHT_HOME>/config/server.defaults.properties` file includes a template for the required configuration values. Copy those lines into your `<ARCSIGHT_HOME>/config/server.properties` file and make the changes there. After making changes to this file, you must stop the Manager and then start all services:

```
/etc/init.d/arcsight_services stop manager  
  
/etc/init.d/arcsight_services start all
```

Caution: Setting the Manager to send SNMP v3 traps is not FIPS compliant. This is because SNMP v3 uses the MD5 algorithm. However, SNMPv1 and v2 are FIPS compliant.

The following provides a description of specific SNMP configuration properties:

```
snmp.trapsender.enabled=true
```

Set this property to true in order to enable the SNMP trap sender.

```
snmp.trapsender.uri=  
/All Filters/Arcsight System/SNMP Forwarding/SNMP Trap Sender
```

The system uses the filter specified by the URI (it should all be on one line) to decide whether or not an event is forwarded. There is no need to change the URI to another filter. These contents are locked and are overwritten when the contents are upgraded to the next version. By default, the "SNMP Trap Sender" filter logic is Matches Filter (/All Filters/ArcSight System/Event Types/ArcSight Correlation Events)—that is, only rules-generated events are forwarded.

```
snmp.destination.host=
```

```
snmp.destination.port=162
```

The host name and the port of the SNMP listener that wants to receive the traps.

```
snmp.read.community=public
```

```
snmp.write.community=public
```

The SNMP community strings needed for the traps to make it through to the receiver. The read community is reserved for future use, however, the write community must match the community of the receiving host. This depends on your deployment environment and your receiving device. Please consult your receiving device's documentation to find out which community string to use.

```
snmp.version=1
```

```
snmp.fields=\
```

```
event.eventId,\
```

```
event.name,\
```

```
event.eventCategory,\
```

```
event.eventType,\
```

```
event.baseEventCount,\
```

```
event.arcsightCategory,\
```

```
event.arcsightSeverity,\
```

```
event.protocol,\
```

```
event.sourceAddress,\
```

```
event.targetAddress
```

These event attributes should be included in the trap. The syntax follows the SmartConnector SDK as described in the FlexConnector Developer's Guide. All the ArcSight fields can be sent. The identifiers are case sensitive, do not contain spaces and must be capitalized except for the first character. For example:

ArcSight Field	SDK/SNMP trap sender identifier
Event Name	eventName
Device Severity	deviceSeverity
Service	service

The SNMP field types are converted as:

ArcSight	SNMP
STRING	OCTET STRING
INTEGER	INTEGER32
Address	IP ADDRESS
LONG	OCTET STRING
BYTE	INTEGER

Additional data values are accessible by name, for example:

```
snmp.fields=event.eventName,additionaldata.myvalue
```

This sends the Event Name field and the value of `myvalue` in the additional data list part of the SNMP trap. Only the String data type is supported for additional data, therefore all additional data values are sent as `OCTET STRING`.

Configuring Asset Aging

The age of an asset is defined as the number of days since it was last scanned or modified. So, for example, if an asset was last modified 29 hours ago, the age of the asset is taken as 1 day and the remaining time (5 hours, in our example) is ignored in the calculation of the asset's age. You can use asset aging to reduce asset confidence level as the time since the last scan increases.

Note: Only the assets belonging to the following categories are considered for aging:

- /Site Asset Categories/Scanned/Open Ports
- /Site Asset Categories/Scanned Vulnerabilities

Excluding Assets from Aging

To exclude certain assets from aging, you can add those assets to a group and then set the property `asset.aging.excluded.groups.uris` in the `server.properties` file to the URI(s) of those groups.

For example, to add the groups `MyAssets` and `DontTouchThis` (both under `All Assets`) add the following to the `server.properties` file:

```
#Exclude MyAssets and DontTouchThis from aging
```

```
asset.aging.excluded.groups.uris=/All Assets/MyAssets,/All Assets/DontTouchThis
```

Note: When setting the `asset.aging.excluded.groups.uris` property keep in mind that the assets in this group are not disabled, deleted or amortized.

Disabling Assets of a Certain Age

By default, asset aging is disabled. There is a scheduled task that disables any scanned asset that has reached the specified age. By default, after the assets aging feature is turned on, this task runs every day half an hour after midnight (00:30:00). Add the following in the `server.properties` file to enable asset aging:

```
#-----  
# Asset aging  
#-----  
# Defines how many days can pass before a scanned asset is defined as old  
# after this time the asset will be disabled  
# Default value: disabled  
asset.aging.daysbeforedisable = -1
```

Note that the default value -1 means that asset aging is turned off, not that assets will be disabled.

The value is expressed in days that define how long an asset is allowed to age before it is disabled. For example:

```
asset.aging.daysbeforedisable = <number of days>
```

So, this setting:

```
asset.aging.daysbeforedisable = 4
```

means that after 4 days, assets will be considered old and disabled. Set this property to a reasonable value that makes sense for your assets.

Deleting an Asset

To delete the asset instead of disabling it, set the property `asset.aging.task.operation` to delete in `server.properties` file:

```
# Delete assets when they age  
asset.aging.task.operation = delete
```

Verify that this property is set to delete for deletion of aging assets to occur.

Amortize Model Confidence with Scanned Asset Age

The `IsScannedForOpenPorts` and `IsScannedForVulnerabilities` sub-elements in the `ModelConfidence` element are factored by the age of an asset. They are extended to include an optional attribute, `AmortizeScan`. If `AmortizeScan` is not defined (or defined with value `-1`), the assets are not amortized. A "new" asset gets the full value while and "old" asset gets no points. You can edit the `AmortizeScan` value (number of days) in the Manager's `/config/server/ThreatLevelFormula.xml` file:

```
<ModelConfidence>
  <Sum MaxValue="10" Weight="10">
    <!-- If target Asset is unknown, clamp modelConfidence to 0 -->
    <HasValue FIELD="targetAssetId" Value="-10" Negated="Yes" />
    <HasValue FIELD="targetAssetId" Value="4" Negated="NO" />
    <!-- Give 4 points each for whether the target asset has been scanned
for open port and vulnerabilities -->
    <!-- This values can be amortized by the age of the asset -->
    <!-- that means that the value will reduce constantly over time as the
asset age -->
    <!-- ie if you set the value to be 120 on the day the assets are
created they receive the four points, by day 60 they'll receive 2 points and
by day 120 they'll receive 0 points -->
    <IsScannedForOpenPorts Value="4" Negated="NO"
      AmortizeScan="-1" />
    <IsScannedForVulnerabilities Value="4" Negated="NO" AmortizeScan="-1" />
  </Sum>
</ModelConfidence>
```

For this example, the value is modified as follows:

Asset Age (in days)	AmortizeScan Value
0	4
60	2

Asset Age (in days)	AmortizeScan Value
120	0
240	0

Tuning for Supporting Large Actor Models

If your actor model contains tens of thousands of members, follow the guidelines in this section to allow adequate processing capacity for best results.

1. Shut down the Manager.

Note: In-memory capacity changes made to `arc_session_list` must match `sessionlist.max_capacity` in `server.properties`

If you update the in-memory capacity for the `arc_session_list` table to number other than the default 500,000, the value you enter must match the value set for `sessionlist.max_capacity` in `server.properties`.

2. **Adjust Java Heap Memory Size** using the Manager Configuration Wizard. Supporting 50,000 actors requires an additional 2 GB of Java heap memory in the Manager. An additional 300 MB is needed for each category model you construct that uses 50,000 actors. This additional memory is not in use all the time, but is needed for certain operations. The Manager Configuration Wizard is covered in ["Running the Manager Configuration Wizard" on page 131](#).
3. Re-start the Manager. See ["Restarting the Manager - Stop the Manager and Start All Services" on page 12](#) for details.
4. Proceed with importing the actor model.

About Exporting Actors

If you need to export your entire actor model to image another Manager, you can do it using the `export_system_tables` command with the `-s` parameter, which specifies the export of session list data. Additionally, the `-s` parameter captures the special session list infrastructure that is part of the Actor Resource Framework in addition to the actor resources themselves.

Viewing License Tracking and Auditing Reports

The system automatically maintains a license audit history that allows you to see how many licenses are in use. When users log into the Console they receive a warning notifying them if they have exceeded their current license. ESM creates an internal audit event for each licensable component to help users track which areas have been exceeded. There are licensing reports on individual features. These reports are located in `/All Reports/ArcSight Administration/ESM/Licensing/`. The reports provide a summary for the number of Actors, Assets, Users, Devices, and EPS identified over the last week.

Setting Up ESM for MSSP Environments

To set up ESM in a managed security service provider (MSSP) environment, do the following:

- Disable the search auto-complete feature. To do this, in the `logger.properties` file change the value of `auto-complete.fulltext.enabled` to **false**.
- Also, consult [ESM Best Practices: Multitenancy and Managed Security Services Providers](#).

Setting up a Custom Login Message for ArcSight Console and Command Center

You can configure the Manager to display a custom login message when users log in to the ArcSight Console and Command Center. The login message will be displayed once, at initial user login. Use this feature to display a legal disclaimer message, or other information you wish the user to see upon login.

To set up a custom login message:

1. Create a text file named `loginbanner.txt` in the `<ARCSIGHT_HOME>/config` directory. In the file, type the login message the user will see. If the message is very long, add physical line breaks so the message formats properly when displayed.
2. Set the following properties in `server.properties`:

Required Property Setting	Purpose
<code>server.staticbanner.text=true</code>	Enables the display of the custom login message
<code>auth.login.banner=config/loginbanner.txt</code>	Configures the Manager to send the text from the <code>loginbanner.txt</code> file

3. To implement the changes, stop the Manager by running the following command as user *arcsight*:

```
/etc/init.d/arcsight_services stop manager
```

4. Then, start all services:

```
/etc/init.d/arcsight_services start all
```

Setting Checkpoint Parameters

When you stop ESM, the system takes a checkpoint (snapshot) of the rules engine to record the actions that occurred until the system stopped. When ESM starts again, it uses the checkpoint to return the system to the state it was in just before it stopped.

Note: Duplicate rule actions after a crash recovery

If you stop ESM, it takes a checkpoint of the rules engine so that it knows what actions have been performed and where it stopped. If ESM crashes in such a way that it cannot take a checkpoint (during a power failure, for example), it returns to the last checkpoint when ESM starts again, and replays events from there. Any actions that occurred between that checkpoint and the ESM crash are therefore repeated. Repeated actions that generate audit events generate duplicate audit events. You should investigate repeated actions that do not duplicate well. For example, if an action adds an item to an Active List, that item's counter will be incremented. If the action runs a command, it will run the command again, and so on.

The following properties related to system checkpoint are configurable in the `server.properties` file on the Manager:

- `rules.checkpoint.enabled=true`
Use this property to set whether or not a rules engine checkpoint file is created. The default is `true`. If this property is set to `true`, then the property `rules.recovery.enabled` should also be set to `true`.
- `rules.checkpoint.interval=300`
This property sets the interval between checkpoints in seconds. The default is 300 seconds (5 minutes).
- `rules.recovery.enabled=true`

Use this property to specify if a checkpoint should be used when ESM starts. The default is `true`.

- `rules.recovery.time-limit=120`

This property sets the time limit on loading events from the database. The default is 120 seconds (2 minutes). For example, if it takes longer than 2 minutes to load the events, the system will stop event recovery at 2 minutes.

- `rules.recovery.event-query-time-range=1800`

This property sets the limit on how far in the past the system will go for events during recovery. The default is 1800 seconds (30 minutes). You can increase this parameter value to accommodate an extended system downtime, but prolonged recovery time can affect system performance. Also, if you change the value of this parameter, you might find you need to change the value of `rules.recovery.time-limit=120` as well.

Preventing Rules Recovery Timeout

Rules recovery can timeout if there is a high EPS on the system, which causes the server to stop loading events from the database for checkpoint. You can modify the `rules.recovery.time-limit` property in `server.properties` to set a higher recovery time limit to attempt to prevent this timeout. The default value for the `rules.recovery.time-limit` property is 120 seconds (two minutes).

Note: The timeout can still occur even after you increase the time limit, due to overall system load, high EPS, or a large number of rules to recover.

Enable Iframe of ArcSight Command Center Pages

To allow iframing of ArcSight Command Center pages, you can add the following optional setting in `server.properties`:

```
allow.from.domains=entries
```

Where `entries` are a comma separated list of the elements that could be of one of the following two forms:

- `origin` (for example, `https://mycompany.com`)
- `key::origin`

In this example, the `key` is any string uniquely identifying the origin within the comma-separated list. For the definition of origins, see <http://tools.ietf.org/html/rfc6454>.

Below is an example of `allow.from.domains` containing several entries. The first entry is origin, while the second is key-value pair:

```
allow.from.domains=https://mycompany.com,microsoft:::https://microsoft.com
```

Third party applications that need to `iframe` Command Center pages should add the parameter "origin" to URLs pointing to Command Center page and use that parameter to specify their origin. For example:

```
https://host:8443/www/ui-  
phoenix/com.arcsight.phoenix.PhoenixLauncher/?origin=microsoft#login
```

In that parameter the origin could be specified directly (`https://microsoft.com`) or with help of the key (`microsoft`) from the above ESM configuration setting.

ESM uses "origin" parameter from HTTP request to lookup an entry in "allow.from.domains" setting. If there is matching entry, then `iframe`ing is allowed for configured origin. If origin is specified in the HTTP request, but is not presented in "allow.from.domains", the request will fail with the exception "Not allowed request".

HTTP requests without "origin" parameter are handled by ESM the same way as before, so there are no changes for regular Command Center sessions. Here `iframe`ing is not allowed to prevent clickjacking vulnerability:

```
https://www.owasp.org/index.php/Clickjacking_Defense_Cheat_Sheet
```

The implementation requires enabling cookies in the browser. It might also be needed to login to Command Center without `iframe`ing from the browser once. Opening Command Center directly creates browser's cookie for the target host. By default, the cookies for `iframe`d pages are not created.

Enabling Scaling for Bytes In and Bytes Out Event Fields

When the values for `Bytes In` and `Bytes Out` event fields are larger than the maximum value of integer numbers, overflow can occur for these fields. You can control scaling for these fields using the server property `bytesInBytesOut.scaling.divider`. This property has a default value 1. If this value is set to be greater than 1, the values for `Bytes In` and `Bytes Out` event fields are scaled, and are saved in ESM in the scaled units.

For example, if the value of the property `bytesInBytesOut.scaling.divider` is 10, all bytes in and bytes out values are divided by 10. By default, the data are not changed. If the values after scaling are still larger than the maximum integer number ($2^{31}-1$), they are then rounded to the maximum integer number, which is 2,147,483,647.

For forwarded events (`Locality=forwarded`), bytes in and bytes out values will not be changed and scaling will not apply.

Notes for bytes in and bytes out behavior in relation to the different versions of SmartConnectors:

- For SmartConnectors version 7.4.0.XXXX.0 and newer: These versions of SmartConnectors send the Bytes In and Bytes Out field values as long numbers to ESM. If the property value is more than 1, scaling will be applied. If a value of Bytes In or Bytes Out event fields after scaling is more than the maximum integer value, it will be truncated to the maximum integer value.
- For SmartConnectors version 7.3.0.XXXX.0 and older: These versions of SmartConnectors do not support values for Bytes In and Bytes Out event fields that are larger than the maximum integer value; in these cases, the entire event is dropped. For Bytes In and Bytes Out fields that have values less than the maximum integer value, data will be scaled if the server property value is more than 1.

Converting an ESM Appliance to IPv6

You can convert your ESM appliance to use IPv6. Note that any connectors registered on the appliance will need to be re-registered after you perform this conversion because the IPv4 IP address will change to a hostname, and the Manager certificate will be regenerated.

To perform the conversion:

1. Stop all the services. As user *root* or *arcsight*, run:
`/etc/init.d/arcsight_services stop all`
2. Confirm if all services are stopped. As user *root* or *arcsight*, run:
`/etc/init.d/arcsight_services status all`
3. As user *root*, run the network configuration script:
`/opt/arcsight/services/bin/scripts/nw_reconfig.py`
4. Reboot the system.
5. As user *root*, edit the `/etc/hosts` file and comment out the line that contains "IPv4 address to hostname mapping if present".
6. Stop the Manager. As user *root*, run:
`/etc/init.d/arcsight_services stop manager`
7. Re-run `managersetup`. As user *arcsight*, run:
`opt/arcsight/manager/bin/arcsight managersetup`
Change the IP protocol to IPv6 and change the hostname to the appropriate IPv6 hostname.
8. Regenerate the Manager certificate.

9. Start all services. As user *root*, run:

```
/etc/init.d/arc_sight_services start all
```

ESM now uses IPv6.

Importing an Archive of 300MB Maximum Size

Use this procedure to import large files into the ArcSight Console. Or, you could export resources from the ArcSight Console to import them into another Manager.

1. In `server.properties`, change the value of the parameters `persist.file.size.total.max` and `persist.file.size.max` to 600:

```
persist.file.size.total.max=600  
persist.file.size.max=600
```

See ["Editing Properties Files" on page 16](#) for details on editing properties.
2. As user *arc_sight*, stop ArcSight services:

```
service arc_sight_servers stop all
```
3. Increase the heap size on the server machine:

```
export ARCSIGHT_JVM_OPTIONS="-Xms12288m -Xms12288m"
```
4. As user *arc_sight*, start ArcSight services:

```
service arc_sight_services start all
```
5. Import the archive using a command like this:

```
./arc_sight archive -f /home/arc_sight/test.xml -u admin -p password -m n15-214-137-h65.arst.swlab.net -action import
```

Customizing Product Image on Login Screen and Navigation Bar in the ArcSight Command Center

You can add your own product images to the ArcSight Command Center login page (up to four images) and to the top left of the navigation bar.

To customize the product images on the ArcSight Command Center login page:

1. Create your own product image or images with your logo and product label, or have one available. There are four images that you can customize on the login page:

Description of Image	Image Name	Image Dimensions
Micro Focus logo	login_logo.png	336W x 84H
ArcSight logo	login_graphic.png	312W x 356H
Background image for the login page	login_background.png	2075W x 1247H
ArcSight Command Center title graphic	login_header.png	331W x 124H

2. Rename the product image file that you want to customize to the image name listed in the table above. The image size must be as specified in the table.
3. Copy the customized images files to:
<ARCSIGHT_HOME>/arcsight-dm/dmapps/ui-phoenix-
1.<X.X>/dist/www/com.arcsight.phoenix.PhoenixLauncher/bannerImages.
4. Refresh the ArcSight Command Center to see the new product images.

To customize the image on the top left of the navigation bar:

1. Create your own product image with your logo and product label, or have one available.
2. Rename the product image file to ACC_NavMenu_Banner.png. The image size must be 256W x 50H.
3. Copy the file ACC_NavMenu_Banner.png to:
<ARCSIGHT_HOME>/arcsight-dm/dmapps/ui-phoenix-
1.<X.X>/dist/www/com.arcsight.phoenix.PhoenixLauncher/bannerImages.
4. Refresh the ArcSight Command Center to see the new product image.

Changing the Hostname of Your Machine

Wherever you see "hostname," you may assume it means "hostname or IP address." If you have configured peering, make sure to re-establish the peer relationship.

Note: You cannot use an IPv6 address as a hostname. IPv4 addresses can be used as hostnames.

If you are using the High Availability module, the procedure is different. Refer to the *ArcSight High Availability Module User's Guide* for the proper procedure.

In case you want to change the IP address of your machine after running the First Boot Wizard successfully, follow these steps:

Note: Run the Manager setup command when logged in as user *arcsight*.

1. Stop all ArcSight services by running (as user *arcsight*):

```
/etc/init.d/arcsight_services stop all
```

2. Change the hostname of your machine.
3. Reboot the machine.
4. As the user *arcsight*, stop the Manager by running:

```
/etc/init.d/arcsight_services stop manager
```

5. As the user *arcsight*, run the Manager's setup program from the <ARCSIGHT_HOME>/bin directory:

```
./arcsight managersetup
```

- a. Enter the new host name (that you set for your machine in the steps above), in the Manager Host Name field when prompted by the wizard - and in every other field where the old hostname is displayed.
- b. Make sure to select the self-signed keypair option when prompted by the wizard and enter the required information to generate the self-signed certificate containing the new host name.

If you are in FIPS mode, you do not get an option to regenerate the keypair. In this case, manually delete the keypair, regenerate it, and then go the next step to start all services.

6. As the user *arcsight*, start all services by running:

```
/etc/init.d/arcsight_services start all
```

7. As the user *arcsight*, see if the manager is running yet by running the command.

```
/etc/init.d/arcsight_services status manager
```

Run this command about once a minute, until you see the line, *manager service available*. Then continue to the next step.

8. Import the Manager's newly-generated certificate on all clients (Console and connectors) that access the Manager. See the "Import a Certificate" topic in the "SSL Authentication" chapter in the ESM Administrator's Guide.
9. Test to make sure that
 - The clients can connect to the Manager
 - Peer configuration works as expected. If not, redo the peer configuration.

Rule Actions Queue Full - Set `rules.action.capacity` Property

If cases are not generated by rules, it might be because the limit on the unprocessed rules action queue has been reached. In this case, you will observe log error messages indicating numerous pending actions. To alleviate this, you can set the `rules.action.capacity` property in the `server.properties` file to increase the queue size limit. See ["Editing Properties Files" on page 16](#) for details on editing properties.

Chapter 3: Configuring and Managing a Distributed Correlation Cluster

Distributed correlation mode allows for multiple correlators and aggregators running as services on nodes in a cluster. A distributed correlation deployment includes the persistor, information repository, correlators, aggregators, message bus data, message bus control, and distributed cache. Ideally, the correlators and aggregators in the cluster will keep up with event flow on your system.

You must balance system resources as you add these components (CPU and memory). Be somewhat generous in your cluster planning and add more correlators and aggregators than you think you need. For information about planning for a distributed correlation cluster, see the [ESM Installation Guide](#). To achieve maximum fault tolerance, configure the cluster over multiple physical systems. For more information about distributed correlation and fault tolerance, see [ESM 101](#).

When you add cluster services (such as instances of correlators, aggregators, distributed cache, or message bus) or change system content (for example, import a package or write a rule that updates an active list), that change is available to every part of the cluster. You do not have to manually synchronize the changes across the cluster nodes.

This chapter provides information about configuring and managing the cluster services, as well as security considerations and troubleshooting information.

Understanding the Cluster Services

You configure and manage cluster services on one or multiple systems. The cluster consists of the following services:

- **Persistor**

The persistor is the root of the cluster and executes processes across the cluster services. The persistor consists of several internal services. Typically, you do not need to manage the internal services:

- manager
- logger_httpd
- logger_servers
- logger_web
- mysqld

- `aps`
- `postgresql`
- `execprocsvc`

You can manage all cluster nodes from the persistor node. Tasks that you perform at the node level affect that node only. When you import content packages into the persistor system, the packages are propagated across the cluster. All database queries run on the persistor node.

- Correlator

The correlator uses filters and rule conditions to discover the relationships between events. You can have multiple instances of correlators.

- Aggregator

The aggregator summarizes or consolidates events with matching (or not matching) values over a timeframe. You can have multiple instances of aggregators.

- Message bus data and message bus control (`mbus_data` and `mbus_control`)

`mbus_data` and `mbus_control` manage messaging in the cluster.

- Distributed cache (`dcache`)

The distributed cache is a shared data cache for the cluster.

- Repository (`repo`)

The repository contains cluster service status information and an index of cluster resources. The repository is the shared information base for the cluster and central to cluster operation.

For conceptual information about distributed correlation clusters, see [ESM 101](#). For information about cluster planning and installation, see the [ESM Installation Guide](#).

Implementing a Distributed Correlation Cluster

To implement a distributed correlation cluster, you must complete the following tasks in the order in which they are listed:

1. Plan the cluster. For more information, see the [ESM Installation Guide](#).

When you are planning the cluster, be aware that all of the cluster nodes must run the same version (major, minor, and patch) of ESM. A program runs periodically to detect nodes that are not running the same ESM version as the persistor node and stops all services on the nodes that do not match the persistor.

2. Install ESM in distributed correlation mode on each server that will be a cluster node. The cluster installation includes the persistor and information repository on

the first cluster node (the persistor node). You can also add other cluster services (aggregators, correlators, distributed cache, and repository) during installation. For more information, see the [ESM Installation Guide](#).

Note: The message bus data and message bus control services are not added during installation. You must configure these services later.

3. The distributed correlation cluster uses key-based passwordless SSH to enable communication among the cluster services. Set up key-based passwordless SSH on the persistor node. For more information, see [Setting Up Key-Based Passwordless SSH](#).
4. Configure cluster services as needed. You must configure the message bus data and message bus control services. For more information, see [Configuring Services in a Distributed Correlation Cluster](#).
5. Start the cluster services on the persistor node. For more information, see [Starting Distributed Correlation Services](#).
6. Perform cluster management tasks. For more information, see [Managing Distributed Correlation Services - Basic Commands](#).

Setting Up Key-Based Passwordless SSH

The distributed correlation services cluster depends on key-based passwordless SSH to enable communication among the cluster services. In the distributed correlation environment, passwordless SSH must be implemented on the node in the cluster that contains the persistor.

The command `arcsight_services` uses passwordless SSH to allow starting and stopping of services on remote nodes through commands originating on the persistor node. In this instance, passwordless SSH works by generating a keypair on the persistor, and configuring the remote node to accept the login based on a public key for the persistor node. In the distributed correlation environment, ESM is configured to allow the user `arcsight` on the persistor node to connect to a remote node as the user `arcsight`. Only `arcsight` user to `arcsight` user passwordless SSH is supported, and only from the persistor node to other cluster nodes.

Set Up Key-Based Passwordless SSH

After installing ESM on all nodes in the cluster, use this command on the persistor node to setup passwordless SSH with cluster nodes. Run as user `arcsight`.

```
/etc/init.d/arcsight_services sshSetup <hostname_for_the_node>
```

If a node needs configuration, the command prompts you for the user *arcsight* password, so it can log in and complete the setup.

For ease of setup, you can elect not to specify the hostname, and in that case the setup is performed on all cluster nodes.

Verify Key-Based Passwordless SSH

On the persistor node, run the command `/etc/init.d/arcsight_services checkSshSetup`. This command verifies whether the nodes in the cluster are configured with passwordless SSH.

Configuring Services in a Distributed Correlation Cluster

To configure a service in a distributed correlation cluster means to add another instance of a service (such as a correlator or distributed cache) to the cluster or to edit an existing instance of a service. Adding message bus control and message bus data services after distributed cluster installation is mandatory.

Note: After you configure a cluster service, you must start all services from the persistor node. For more information, see [Starting Distributed Correlation Services](#).

As you add cluster services, ESM increments the instance IDs. For example, as you add correlators, ESM numbers them as `correlator1`, `correlator2`, and so on. If you delete an instance of a service (for example, `correlator2`), ESM increments the next instance of that service by one. In this example, if you delete `correlator2` and then add a new correlator instance, ESM numbers that instance as `correlator3`. Because ESM identifies all cluster services internally by the assigned instance ID, ESM does not re-use numbers.

You can add names to services to identify them for organization, as a reminder of the point of adding the services, or other purposes that help you work with the cluster. These service names do not replace the instance IDs that ESM provides, but exist in addition to those instance IDs.

Configuring Message Bus Control and Message Bus Data Services

Message bus control (`mbus_control1`) and message bus data (`mbus_data`) facilitate communication among cluster services. All instances of message bus control or

message bus data must be available and functioning for event flow to continue in the cluster.

Important: You must set up passwordless SSH on the persistor node in order to configure the message bus data and message bus control services. If passwordless SSH is not implemented, the `mbusetup` wizard will not perform any configuration. For information about setting up passwordless SSH, see [Setting Up Key-Based Passwordless SSH](#).

When you change message bus data or message bus control instances, ESM replaces all instances of these services and you lose the data in the replaced instances. Also, ESM increments the number in the name of the new instances. For example, assume that you have a four-node cluster with the following message bus data and message bus control instances:

Node 1 - no message bus instances

Node 2 - `mbus_data1`, `mbus_control1`

Node 3 - `mbus_data2`, `mbus_control2`

Node 4 - `mbus_data3`, `mbus_control3`

If you change this cluster to add a message bus data instance to a new node 5, ESM increments the existing instance IDs by one, based on the highest instance ID from the previous configuration:

Node 1 - no message bus instances

Node 2 - `mbus_data4`, `mbus_control4`

Node 3 - `mbus_data5`, `mbus_control5`

Node 4 - `mbus_data6`, `mbus_control6`

Node 5 - `mbus_data7`

Unlike other services, you can only stop message bus control services from the persistor node. When you run `/etc/init.d/arcsight_services stop mbus_control<#>` from the persistor node, it will stop all instances of the message bus data service.

Caution: If you plan to install the High Availability module on the persistor node, do not configure a message bus data instance on the persistor node. Otherwise, the cluster is likely to fail when ESM swaps the primary and secondary systems.

To configure multiple message bus control and message bus data instances in a cluster:

Note: Complete this task on the persistor node.

1. Verify that key-based passwordless SSH is set up on the persistor node.
2. Stop all services:

```
/etc/init.d/arcsight_services stop all
```
3. Start the information repository:

```
/etc/init.d/arcsight_services start repo
```
4. Run `<ARCSIGHT_HOME>/bin/arcsight mbussetup`. The command starts the configuration wizard on the persistor node, but does not add a message bus instance on the persistor node. Add message bus data and message bus control instances to non-persistor nodes.
5. Choose the appropriate action and provide the requested information:

Action	Description
I want to add, delete, or change an instance	<p>a. Enter the number of <code>mbus_data</code> instances you want on each node.</p> <p>You need a minimum of three message bus data instances per cluster.</p> <p>b. Enter the number of <code>mbus_control</code> instances you want on each node.</p> <p>Add one or three instances of the message bus control service per cluster.</p> <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;"> <p>Note: While it is possible to configure one message bus control instance in a three-node cluster, Micro Focus does not recommend this configuration.</p> </div> <p>You can specify different node locations to change the location of instances. You might want to do this if the server on which the persistor and the initial instance of the repository is running becomes unavailable.</p>
I want to update an <code>mbus_control</code> instance	<p>Enter a name for the service.</p> <p>This name does not replace the instance ID that ESM provides (for example, <code>mbus_control1</code>), but allows you to provide a friendly name to identify the service in logs and other service reporting.</p>
I want to update an <code>mbus_data</code> instance	<p>Enter a name for the service.</p> <p>This name does not replace the instance ID that ESM provides (for example, <code>mbus_data1</code>), but allows you to provide a friendly name to identify the service in logs and other service reporting.</p>

6. If you added services or changed services (such as adding a friendly name for a service), you must start the services in order for the configuration changes to take effect. For more information, see [Starting Distributed Correlation Services](#).

Adding Correlators and Aggregators

This section describes how to add correlators and aggregators in the following scenarios:

- The node to which you want to add correlators or aggregators does not currently include correlator or aggregator instances.
- The node to which you want to add correlators or aggregators already includes at least one correlator or aggregator instance.

To add correlator and aggregator instances:

1. Run `<ARCSIGHT_HOME>/bin/arcsight correlationsetup` on the cluster node to which you wish to add an instance.

Note: `correlationsetup` operates on a per instance basis. You must configure instances one at a time.

2. Select the appropriate action and provide the requested information:

Action	Description
Select the type of service you would like to configure	Select Correlator or Aggregator .
Add/Configure an Instance	Provide the following information: <ul style="list-style-type: none">• Host name, port, and physical location of the server to which you want to add the correlator or aggregator• Java heap size• Key pair• Service name This name does not replace the instance ID that ESM provides (for example, <code>correlator1</code>), but allows you to provide a friendly name to identify the service in logs and other service reporting.

3. If the node did not previously include correlator or aggregator instances, after you exit the wizard, complete the following steps:
 - a. As user `arcsight`, run the following command on the persistor node:

```
<ARCSIGHT_HOME>/bin/arcsight certadmin -list submitted
```

Review the output to verify that the certificates represent the cluster nodes. To view the certificate details, use the `-v` option.
 - b. After you confirm that the certificate list is correct, run the following command:

```
<ARCSIGHT_HOME>/bin/arcsight certadmin -approveall
```

Specify the cluster administration password that was provided when you installed ESM on the persistor node.

- c. Stop and start the ArcSight Manager so that it can read the certificates:

```
/etc/init.d/arcsight_services stop manager
```

```
/etc/init.d/arcsight_services start all
```

You must start all of the services for the configuration changes to take effect.

4. If the node already included at least one correlator or aggregator instance, run the following command to start the new instance:

```
/etc/init.d/arcsight_services start <correlator or aggregator ID>
```

Configuring Correlators and Aggregators

This section describes how to modify the configuration of correlator and aggregator instances.

To modify the configuration of correlator and aggregator instances:

1. Run `<ARCSIGHT_HOME>/bin/arcsight correlationsetup` on the cluster node on which you want to configure the instance.

Note: `correlationsetup` operates on a per instance basis. You must configure instances one at a time.

2. Choose the appropriate action and provide the requested information:

Action	Description
Select the type of service you would like to configure	Select Correlator or Aggregator .
Add/Configure an Instance	Provide the following information: <ul style="list-style-type: none"> • Host name, port, and physical location of the server where you want to modify the correlator or aggregator • Java heap size • Key pair • Service name This name does not replace the instance ID that ESM provides (for example, <code>correlator1</code>), but allows you to provide a friendly name to identify the service in logs and other service reporting.

Action	Description
Delete an Instance	Specify the instance to delete. You must stop the instance before you delete it: <pre>/etc/init.d/arcsight_services stop aggregator# or /etc/init.d/arcsight_services stop correlator#</pre>

3. Run the following commands to stop and then start the modified instance:

```
/etc/init.d/arcsight_services stop <correlator or aggregator ID>
/etc/init.d/arcsight_services start <correlator or aggregator ID>
```

Configuring Distributed Cache

The distributed cache contains the shared resources for the cluster.

In some cases, distributed cache nodes might lose contact with each other. For information about resolving the issue, see [Restoring the State of Distributed Cache Nodes](#).

To configure multiple distributed cache instances in a cluster:

1. Run `<ARCSIGHT_HOME>/bin/arcsight dcachesetup` on the cluster node to which you wish to add an instance.

Note: `dcachesetup` operates on a per instance basis. You must configure instances one at a time.

2. Select the appropriate action and provide the requested information:

Action	Description
Configure Common Settings	Specify the interval at which to report that distributed cache instances are running. The default is 180 seconds.
Add/Configure an Instance	Specify the service name. This name does not replace the instance ID that ESM provides (for example, <code>dcache1</code>), but allows you to provide a friendly name to identify the service in logs and other service reporting.
Delete an Instance	Specify the instance to delete. You must stop the instance before you delete it: <pre>/etc/init.d/arcsight_services stop dcache#</pre>

3. If you added a distributed cache instance, run the following command to start the instance:

```
/etc/init.d/arcsight_services start <distributed cache ID>
```

4. If you changed an existing distributed cache (such as adding a friendly name), run

the following commands to stop and then start the modified instance:

```
/etc/init.d/arcsight_services stop <distributed cache ID>
```

```
/etc/init.d/arcsight_services start <distributed cache ID>
```

Restoring the State of Distributed Cache Nodes

In some cases, distributed cache nodes might lose contact with each other because of network interruptions or a heavily-loaded system. If this occurs, not all data is shared between correlators, aggregators, and the persistor. As a result, some data monitors and dashboards will show no data, and there might be a drop in EPS.

If there is a connection issue, ESM displays a `Connection to DC` message in the ArcSight Command Center Cluster View dashboard. To correct the issue, you must identify the distributed cache (`dcache`) instances that are causing the problem and need to be restarted.

To restore the state of distributed cache nodes:

1. From the ArcSight Command Center Cluster View dashboard, check the audit events and look for the event **DCache connection is down** and the associated service message **Hazelcast cluster inconsistency**
2. Hover your mouse over the **Hazelcast cluster inconsistency . . .** service message to identify the service that is causing the issue. For example:

```
Hazelcast cluster inconsistency. Some DCache instances are not accessible.  
Restart them if they are running (split-brain), otherwise unregister them  
using command "dcache-repo-records". Troubled instances: dcache1@host2,  
dcache2@host3
```

In this example, the name of the distributed cache instance that is causing the issue is `dcache2`. The host name is the name of the server in the cluster on which that particular distributed cache instance resides.

3. Stop the services:

```
/etc/init.d/arcsight_services stop dcache2
```

4. Run the following command to remove information repository records from non-responsive distributed cache instances:

```
bin/arcsight dcache-repo-records -r dcache#
```

where `#` is the `dcache` ID

The command cleans internal runtime records for the `dcache` instance in the information repository. The instance automatically resets the records after the network connection is restored.

5. Restart the services:

```
/etc/init.d/arcsight_services start dcache2
```

Configuring a Repository

By default, ESM creates a repository instance on the persistor node (the first node on which you installed ESM in distributed correlation mode). After installation, you can configure three repository instances. A cluster can have either one repository instance or three instances, with one repository instance per node. Other numbers of repository instances are not supported. The benefit of having multiple repository instances is that you have more than one place to store cluster service status information (including certificate status and port assignments for services); all repository instances work together to provide the repository service. If a node that contains a repository becomes unavailable, the other instances of the repository will maintain the cluster status information. The repository is central to cluster operation, and cluster membership is based on the information in the cluster repository. If the repository is unavailable, the cluster stops working.

When you change repository instances, ESM replaces all instances of the service and you lose the data in the replaced instances. Also, ESM increments the number in the name of the new instances.

To configure multiple repository instances in a cluster:

1. Verify that key-based passwordless SSH is set up on all cluster nodes.

2. Stop all services:

```
/etc/init.d/arcsight_services stop all
```

3. Start the information repository:

```
/etc/init.d/arcsight_services start repo
```

4. Run `<ARCSIGHT_HOME>/bin/arcsight repositup` on the persistor node.

5. Choose the appropriate action and provide the requested information:

Action	Description
Update an Information Repository Instance	<p>Specify the service name.</p> <p>This name does not replace the instance ID that ESM provides (for example, <code>repo1</code>), but allows you to provide a friendly name to identify the service in logs and other service reporting.</p>
Change the list of Information Repository Instances	<p>From the list of existing nodes, select two nodes to add a total of three repositories. The initial repository instance was created during installation.</p> <p>You can specify different node locations to change the location of instances. You might want to do this if the server on which the initial instance of the repository is running is due to be taken offline or will otherwise be unavailable.</p> <p>Note: After you change the location of repository instances, ESM restarts existing instances, starts new instances, and stops deleted instances.</p>
Change the TCP Port Range for ESM Processes	<p>Specify a range of ports to use for dynamic assignment to correlator and aggregator services.</p> <p>Enter the lowest port value and the highest port value. The lowest value must be lower than the original lowest value, and the highest value must be higher than the original highest value. For example, if the original lowest value was 10000, and the original highest value was 12000, then a valid set of new port designations would be the lowest ESM server port at 9000, and the highest ESM server port at 14000. You can change the lowest value, the highest value, or both.</p> <p>The lowest supported port value is 1024 and the highest supported port value is 32767.</p> <p>You specified the original range during installation and configuration of the first cluster node (the persistor).</p>
Remove a node from the Information Repository	<p>Select the nodes to remove.</p> <p>For more information about removing a node, see Removing a Node from a Cluster.</p> <p>Note: Preserve at least one repository instance. New repository instances receive data from existing instances.</p>

- If you added or removed an information repository instance, ESM stops all of the instances and you must run the following command to start them:

```
/etc/init.d/arcsight_services start repo
```

Starting Distributed Correlation Services

After configuring new services or changing the cluster configuration (such as adding a friendly name for a service), you must start all services. From the persistor node, run `/etc/init.d/arcsight_services start all`.

If a non-persistor node is not running the same version (major, minor, and patch) of ESM as the persistor node, the services will not start on the non-persistor node. For information about resolving this issue, see [Why are services stopped or in a prolonged initializing state on non-persistor nodes?](#)

Forwarding Events in Distributed Correlation Mode

Distributed event forwarding is available when ESM is installed in distributed correlation mode. The feature allows you to forward events from ESM to Transformation Hub at a much higher rate than the Forwarding Connector supports. Distributed event forwarding leverages the distributed infrastructure of ESM to allow ESM to spread the work of event forwarding across the cluster, similar to how ESM distributes event correlation. This allows event forwarding to scale horizontally.

Events that ESM forwards to Transformation Hub can subsequently be read by another ESM instance or multiple ESM instances. Those ESM instances do not have to be installed in distributed correlation mode in order to read events from Transformation Hub.

If you need to forward events from ESM to Transformation Hub at a high rate (generally higher than 10K events per second) Micro Focus recommends that you use ESM in distributed correlation mode and use distributed event forwarding instead of the Forwarding Connector.

Distributed event forwarding requires the following:

- CA certificate of the Transformation Hub cluster
- ESM filter that determines which events to forward
You can create a filter or use one that is installed with ESM.
- List of broker socket addresses (in the form `host:port`) of the Transformation Hub cluster to which you want to forward events
The Transformation Hub documentation refers to these as "worker nodes."
- Transformation Hub topic name to which you want to publish events

Note: Perform the configuration steps that are described below on the persistor node of the ESM cluster.

Obtaining and Importing the Transformation Hub CA Certificate

Transformation Hub maintains its own certificate authority (CA) to issue certificates for individual nodes in the Transformation Hub cluster. ESM needs that CA certificate in its truststore so that it will trust connections to Transformation Hub. For information about obtaining the certificate, see the information about viewing and changing the certificate authority in the [Transformation Hub Deployment Guide](#). You might need to contact the Transformation Hub administrator to obtain the CA certificate if you do not have sufficient privileges to access the Transformation Hub cluster.

If you have root access to the Transformation Hub cluster (version 3.x or later), you can obtain the CA certificate as follows:

```
master=<master node host name or IP address>
ssh root@$master env K8S_HOME=/opt/arcsight/kubernetes
/opt/arcsight/kubernetes/scripts/cdf-updateRE.sh > /opt/home/arcsight/th-
ca.crt
```

The command copies the CA certificate to the file `/opt/home/arcsight/th-ca.crt`. You must then import that certificate into the ESM cluster using the [certadmin](#) tool.

To import the Transformation Hub CA certificate to ESM:

1. Ensure that the ESM cluster is running. If it is not, run the following command:

```
/etc/init.d/arcsight_services start
```

Wait for the cluster to completely start.
2. Import the Transformation Hub CA certificate:

```
bin/arcsight certadmin -importcert /opt/home/arcsight/th-ca.crt
```

Note the alias that is reported in the output.
3. Run the following command and verify that the alias that was reported in the output from the previous command is listed as an approved certificate:

```
bin/arcsight certadmin -list approved
```

Configuring the Filter and Destination

After you import the Transformation Hub CA certificate to ESM, configure a filter and the destination properties. You specify the destination properties in a file that is used as

input to the [configure-event-forwarding](#) command.

To configure the filter and destination:

1. Create a filter to select the events that you want to forward or use a filter that is installed with ESM.

For more information about creating event filters, see the information about filtering events in the [ArcSight Console User's Guide](#).

2. Edit `forwarding.properties` in `/opt/arcsight/manager/config/` with the following information:
 - Name of the filter that you want to use to select the events to be forwarded.
 - Comma separated list of socket addresses of the worker nodes of the Transformation Hub cluster. Each socket address should be in the form `host:port`.

Note: Optionally, you can copy `forwarding.properties` to a different file and edit that file. If so, you must specify the file name when running the `configure-event-forwarding` utility. Otherwise, `configure-event-forwarding` will read the default `forwarding.properties` file.

The properties file is in the following format:

```
# Provide the filter that determines which events to forward.
# The filter may be specified by ID or by Name.
# If name is used then it can be the simple name or the fully qualified
URI.
#
# e.g. filterID=2jq50g-sAABCAFyopCdLChw==
# or
# filterName=Non-ArcSight Internal Events
# or
# filterName=/All Filters/ArcSight System/Event Types/Non-ArcSight
Internal Events
filterName=Non-ArcSight Internal Events
# List the socket addresses of the destination worker nodes
hostPortCSV=host1.example.com:9093,host2.example.com:9093,host3.example.co
m:9093
# Specify the destination topic name
topicName=esm-forwarded-events
```

3. Run the following command to validate the settings in the properties file:

```
bin/arcsight configure-event-forwarding -validate <file>
```

Note: If you do not specify a file name, <ARCSIGHT_HOME>/config/forwarding.properties is the default file.

The command output indicates whether the filter and connections to the forwarding destination are valid.

4. If the settings in `forwarding.properties` are valid, run the following command to set the configuration and save it in the information repository:

```
bin/arcsight configure-event-forwarding -commit <file>
```

Note: If you do not specify a file name, <ARCSIGHT_HOME>/config/forwarding.properties is the default file.

5. When you are ready to actively filter and forward events, run the following command to enable event forwarding:

```
bin/arcsight configure-event-forwarding -enable
```

To disable event forwarding, run the following command:

```
bin/arcsight configure-event-forwarding -disable
```

Note: ESM does not cache events when you disable distributed event forwarding. Events that ESM ingests while forwarding is disabled will not be forwarded, even when you subsequently enable forwarding. Events that ESM ingests after you enable forwarding will be forwarded.

6. To view the current settings for distributed event forwarding, including whether it is enabled or disabled, run the following command:

```
bin/arcsight configure-event-forwarding -print
```

Modifying the Filter and Configuration

The properties file is only used to validate and commit the configuration. The information repository stores the configuration itself. If you need to modify the configuration, it is not sufficient to simply edit the file. You must edit the file and then run `bin/arcsight configure-event-forwarding -validate <file>` and `bin/arcsight configure-event-forwarding -commit <file>` again to overwrite the old configuration. It is not necessary to stop and start the ESM cluster to modify the configuration. Distributed event forwarding will automatically detect the updated configuration. You can run `bin/arcsight configure-event-forwarding -print` to view the configuration that is currently in use.

If you modify the filter that distributed event forwarding uses, event forwarding automatically detects the updated filter and begins using it to select events for forwarding as soon as you save the filter update. Therefore, be cautious when modifying the filter.

Instead of modifying the filter, you can create a new filter and test it before you commit to using it. When you are sure that the filter is correct, enter the new filter in the properties file, and then run `bin/arcsight configure-event-forwarding -validate <file>` and `bin/arcsight configure-event-forwarding -commit <file>` to apply the change to the current configuration.

Troubleshooting Event Forwarding Throughput

The maximum rate at which events can be forwarded depends on many factors, including the following:

- Amount of other work that distributed correlation must do to support all of the rules, data monitors, and other content that you have defined
- File input/output contention
- CPU contention
- Memory contention
- Network contention, especially the network between distributed correlation nodes and the Transformation Hub worker nodes
- Maximum rate at which Transformation Hub can accept messages

It might be that the maximum throughput of event forwarding is not enough to support the number of events that need to be forwarded in a given period of time. When that happens, events yet to be forwarded will build up inside message bus. This buildup of unforwarded events is called **lag**. If the lag is too high, the ESM cluster will stop ingesting events to reduce the lag. When this happens it is called **backpressure**.

The Acceptable Lag value in the Cluster View dashboard of Command Center defines the amount of lag that can occur before backpressure is applied. For more information, see the information about using the Cluster View dashboard in the [Command Center User's Guide](#).

Distributed event forwarding leverages the ESM distributed infrastructure to gain horizontal scalability. The correlator does the work of event forwarding. You might be able to add event forwarding throughput by adding correlator(s), either on an existing node or on a new distributed correlation node. However, this will only increase throughput if it is the correlators that are causing the bottleneck (for example, because of CPU or memory limitations). If the network or Transformation Hub is causing the bottleneck, adding correlators might not have any effect. For information about adding correlators, see [Adding Correlators and Aggregators](#).

Optionally, you can disable backpressure that might occur as a result of unforwarded events, but you should only use this option if you accept that some events might never be forwarded. To disable backpressure, run the following command:

```
bin/arcsight configure-event-forwarding -backpressure disable
```

Note: When you disable backpressure, ESM will not slow the ingestion of events if event forwarding cannot keep up with the rate at which events are being filtered for forwarding. The lag might build up inside message bus to the point that the oldest unforwarded events are dropped. Those events will not be forwarded.

To enable backpressure, run the following command:

```
bin/arcsight configure-event-forwarding -backpressure enable
```

Managing Distributed Correlation Services - Basic Commands

Use the `arcsight_services` command to manage ArcSight distributed correlation services on the cluster. You can start and stop and otherwise manage services from the machine that hosts the persistor or from the specific nodes where services are installed in the distributed correlation environment.

While you can view the status of all services in a cluster from one node, you can only start and stop services from the local node or from the persistor node. This means that if you make a request from a non-persistor node (Node A), and the process is on another node in the cluster (Node B), `arcsight_services` cannot start or stop it. Only commands issued from the persistor node or from the local node will be successful. In this case, the command to start a service on Node B must come from the persistor node or Node B. Node A cannot start services on Node B.

In the distributed correlation environment, there can be multiple instances of services (such as correlators and aggregators). In this case, you must identify which of several processes you want to start. To support that, each process has an instance ID, consisting of its type, followed by a number (for example, *correlator1*). You can start or stop all processes by type (for example, *correlator*) as well. In the case of persistor services, there will only be one process for each service, and no number is used (for example, *manager*, or *mysqld*).

See ["ArcSight_Services Command - Distributed Correlation Mode" on page 145](#) for further information on using the `arcsight_services` command in a distributed correlation environment.

This section shows examples of the commands you will use most of the time for cluster management.

Stop Services

To stop service by service type

```
/etc/init.d/arcsight_services stop <service_type>
```

For example:

```
/etc/init.d/arcsight_services stop logger_servers
```

To stop a specific instance of a service type

```
/etc/init.d/arcsight_services stop <instance_ID>
```

For example:

```
/etc/init.d/arcsight_services stop correlator3
```

Note: You can specify a friendly name for a service, but when you stop services, you must use the instance ID provided by the cluster.

To stop all services (run from the persistor node only)

```
/etc/init.d/arcsight_services stop all
```

If a service fails to stop, you will receive a message like this (for example, for the service manager):

```
FAIL: timed out awaiting stop of manager service after PT00:050:00.000
```

Start Services

To start service by service type

```
/etc/init.d/arcsight_services start <service type>
```

For example:

```
/etc/init.d/arcsight_services start dcache
```

will start all instances of the distributed cache service in the cluster.

To start a specific instance of a service type

```
/etc/init.d/arcsight_services start <instance_ID>
```

For example:

```
/etc/init.d/arcsight_services start correlator3
```

Note: You can specify a friendly name for a service, but when you start services, you must use the instance ID provided by the cluster.

To start all services (run from the perisistor node only)

```
/etc/init.d/arcsight_services start all
```

Check the Status of Services

To check a specific service status

```
/etc/init.d/arcsight_services status <service_type>
```

For example:

```
/etc/init.d/arcsight_services status repo
```

will display the status of all repository instances in the cluster.

To check the status of all services (run from the persistor node only)

You can display service status from the persistor system or from individual nodes. Status for persistor services can be viewed only on the persistor machine.

```
/etc/init.d/arcsight_services status all
```

To check service status by node (run from the persistor node only)

```
/etc/init.d/arcsight_services statusByNode <hostname>
```

Displays service status information grouped by node, and the status of the node. If a hostname is specified, only services are listed; if not, all services on all hosts are listed. The process name displays as well.


See "[ArcSight_Services Command - Distributed Correlation Mode](#)" on page 145 for details on running `arcsight_services`.

Start and Stop Order of Distributed Correlation Processes

The services in a distributed correlation cluster are started and stopped in a predetermined order. The `arcsight_services` command starts servers (for example, databases) first, followed by clients (for example, `logger_web`). Instances of services of the same type are started at the same time, meaning, for example, that all repository instances are started at the same time and all correlator instances are started at the same time.

First, `arcsight_services` starts all the repository instances. This is important because the repository contains information on the status of all processes.


The diagram below illustrates the order of starting process types, as well as when `arcsight_services` waits for servers to start before starting clients.



In this diagram, the order of process start is indicated by the numbers below each process type. As shown, the repository instances start first and the aggregator instances start last. The arrows in the graphic point from clients to servers. Note that the processes are started in the reverse direction of the arrows. The arrows indicate which processes to start. For instance, in order to start an aggregator instance, besides that aggregator, all correlators, distributed cache (dcache), manager, execprocsvc, logger_servers, postgresql, mysqld, mbus_data (message bus data), and mbus_control (message bus control) are started.

The arrows also indicate when arcsight_services waits for servers to start before starting clients. For instance, logger_servers do not start until mysqld and postgresql are running. By contrast, requests to start aggregators usually immediately follow the request to start correlators.

The following diagram illustrates processes stopping.



The stop order is roughly the reverse of the start order. Client-server relationships remain the same. However, more servers must wait for clients to stop before the server itself can stop to ensure there is no data loss.

Monitoring the Cluster Using the Cluster View Dashboard

Monitoring the processing of the existing cluster with the Cluster View Dashboard and data monitors will allow you to make decisions about cluster management and sizing. The Cluster View dashboard is provided in the ArcSight Command Center to allow you to monitor cluster activity and determine if the correlators and aggregators you are managing are keeping up with the event load.

This dashboard includes a visualization of the entire cluster and status information for each service in the cluster (available or not available). The status of a service running on a node in a cluster is based on the heartbeat from that service that is collected in the repository. The dashboard shows this status as derived from the cluster's repository. See the *ArcSight Command Center User's Guide* for details.

Certificate-Based Admission of Services to a Cluster

Admission to a cluster is enabled using a master certificate with the digital signature for the cluster. All services in the cluster share the same keypair and the same keystore. Each cluster component has its own SSL certificate. This certificate is imported into the cluster trust store for each cluster component (service) during ESM distributed correlation installation. Each cluster node must have a client-side SSL certificate (one certificate per node). This certificate can be generated by ArcSight or be one that you provide.

When you add a service to the cluster, it submits an admission request to the cluster repository. The request includes the client-side and server-side SSL certificate and the name of the service to be admitted to the cluster (for example, *correlator2*). Services (clients) authenticate themselves to the persistor (which acts as a server).

Use the `certadmin` command to view all certificates based on status: submitted, approved, or revoked. Use this command to approve or revoke SSL certificates used by the aggregator, correlator, and persistor services. See ["certadmin" on page 164](#) for details.

Dynamic Ports in the Distributed Correlation Environment

Each instance of a correlator or aggregator service that is part of a cluster has a port associated with that service. That port is permanently identified with the individual instance of the service, and cannot be reused. For example, if port 20201 is associated with the cluster service *correlator3*, if you happen to remove *correlator3* from the cluster, the port 20201 will not be made available for use by another service.

You must specify a range of available ports when you install ESM and specify the first node (the persistor node) of your cluster. This range of ports is made available for dynamic assignment to services as they are added to a cluster.

The default range offered during install is:

- Lowest ESM server port: 10000
- Highest ESM server port: 10100

You can set other values during installation; the lowest value can be 1024 and the highest value 32767. The difference between the lowest value and the highest value specified must be at least 100.

If, after working with the cluster, you find you need to extend the range of available ports, you can do so using `reposestap`. See ["Configuring Services in a Distributed Correlation Cluster" on page 60](#) for details.

Viewing Port Numbers for Dynamically Allocated Ports

If you want to find the port number of a dynamically allocated port for a specific service, you can:

- Run the `correlationsetup` wizard on the node for that service, and choose an existing correlator or aggregator; the port number associated with that service is one of the properties displayed. See ["Configuring Services in a Distributed Correlation Cluster" on page 60](#) for details on running `correlationsetup`.
- In the ArcSight Console, view the Distributed Correlation Audit Events Active Channel. The field `deviceCustomString1` contains the instance ID (for example, `correlator4`). The field `source port` displays the port number for that service.

Changing Authentication in a Distributed Correlation Environment

To change authentication for distributed correlation:

1. Stop all services on the cluster persistor node:

```
/etc/init.d/arcsight_services stop all
```
2. Using `managersetup`, change the authentication on the persistor node. See ["Running the Wizard" on page 131](#) for details on running `managersetup`.
3. Start all services on the cluster persistor node:

```
/etc/init.d/arcsight_services start all
```

Changing Hostnames or IP Addresses in a Cluster

This information applies to ESM in distributed mode only.

About:

Each node in a cluster is identified either by hostname or IP address. A single host may have multiple hostnames and IP addresses. The procedures refer to the specific

hostname or IP address where ArcSight services are stored.

To determine the hostnames for ArcSight services:

Run the command from any node as arcsight:

```
/etc/init.d/arcsight_services version
```

Following is an example of the returned information (your results may not match exactly, but the versions should all be identical in all returned nodes):

```
[arcsight@n15-nnn-nnn-nnnn esm]$ /etc/init.d/arcsight_services version
```

Build versions:

myhost1:

```
esm: 7.0.0.XXXXX.0(BEXXXXX)
storage: 7.0.0.AAAAA.0(BLAAAAA)
process_management: 7.0.0-BBBBB
installer: 7.0.0-CCCCC
highavail: (HA2345)
```

myhost2:

```
esm: 7.0.0.XXXXX.0(BEXXXXX)
storage: 7.0.0.AAAAA.0(BLAAAAA)
process_management: 7.0.0-BBBBB
installer: 7.0.0-CCCCC
highavail: (HA2345)
```

myhost3:

```
esm: 7.0.0.XXXXX.0(BEXXXXX)
storage: 7.0.0.AAAAA.0(BLAAAAA)
process_management: 7.0.0-BBBBB
installer: 7.0.0-CCCCC
highavail: (HA2345)
```

In the above example, the nodes are identified by the hostnames myhost1, myhost2, and myhost3. If the name(s) myhost*n* are changing to something else, then use the procedures in this topic to apply the new hostname in ESM.

Typically a single host may have multiple hostnames and IP addresses. If the hostname or IP Address shown above are to be changed, then you are required to perform the documented procedure.

Initial steps:

On the persistor, log in the user `arcsight` and stop all services:

```
/etc/init.d/arcsight_services stop
```

Run on each node where the hostnames will be changed:

On all nodes where the hostnames will be changed, run the following command as root:

```
<ARCSIGHT_HOME>/bin/remove_services.sh
```

This prevents services from restarting, in case a reboot is necessary.

Change the hostnames on each node:

Use the OS-specified process for changing the IP address or hostname on the node.

Note: If the persistor is an HA setup refer to the *ESM High Availability Module User's Guide*. Follow the instructions in the topics "Changing Hostname, IP Address, or Service IP" for both the primary and secondary ESM systems.

If the persistor is an HA setup, as the user `arcsight`, confirm the change to the HA subsystem by running the command

```
/usr/lib/arcsight/highavail/bin/arcsight firstBootWizard
```

The wizard should display the new information. Exit the wizard.

After changing hostnames:

After the hostnames have been changed in the system, perform these steps on the persistor as user `arcsight`.

1. Run

```
<ARCSIGHT_HOME>/bin/arcsight updaterepohostconfig
```

Enter the updated hostnames when prompted. This is only necessary if the hostnames where one or more of the repository instances runs, has changed. It updates the repository configuration so that the repository can be accessed using the new hostnames.

In most cases, this step will fail the first time it is run, because `sshSetup` needs to be run again. The message will display a similar information:

6 errors were detected.

See `/opt/arcsight/var/logs/misc/updaterepohostconfig.log` for details.

2. Run

```
/opt/arcsight/services/init.d/arcsight_services sshSetup
```

This fixes the errors that caused `updaterepohostconfig` to fail. Note that the `arcsight_services` in `/etc/init.d` is not used. It was removed when you ran `remove_services.sh` on the persistor. Expect to see errors like:

```
ERROR: Information Repository is down.
```

Ignore such messages.

3. Run

```
<ARCSIGHT_HOME>/bin/arcsight updaterepohostconfig
```

but do not change hostnames at this point, since these have already been updated.

At the prompt to continue, enter `yes`.

4. If the persistor hostname was also changed, run the command as user `root`:

```
<ARCSIGHT_HOME>/bin/setup_services.sh
```

This enables `monit` and also starts local services.

5. As user `arcsight`, run:

```
/etc/init.d/arcsight_services stop all
```

This is necessary because in the command in the previous step started the local services.

6. Run

```
/etc/init.d/arcsight_services start repo
```

```
/etc/init.d/arcsight_services start mysql
```

Update the hostname(s) in the repository:

As user `arcsight`, run

```
<ARCSIGHT_HOME>/bin/arcsight updatehostnameinrepo
```

When prompted, enter the old and new hostnames for this node. Repeat the process on each node. This updates the hostnames in the repository.

Set up the persister:

Updating the manager keypair is only necessary if the persister hostname or IP address is changed.

Run these commands as user `arcsight`.

1. If this is a FIPS installation, and/or CA signed certificates are being used, replace the existing certificate with the alias `mykey` in the `managerkeys` keystore with a new certificate with the new hostname using the instructions for creating FIPS and/or CA certificates.
2. Run

```
bin/arcsight managersetup
```

Even if you have already generated a new keypair outside of `managersetup`, you still need to run the `managersetup` wizard to completion so that the newly generated certificate can be submitted to the repository.

Set up other nodes:

These steps are only need on nodes where the hostname/IP was changed. Run the commands as user `arcsight`.

1. If there are correlators or aggregators on this node, do the following:
 - a. If this is a FIPS installation, and/or CA signed certificates are being used, replace the existing certificate with the alias `mykey` in the `managerkeys` keystore with a new certificate with the new hostname using the instructions for creating FIPS or CA certificates.
 - b. Run:

```
<ARCSIGHT_HOME>/bin/arcsight correlationsetup
```

Pick an already configured correlator or aggregator, and select the current options.

Even if you have already generated a new keypair outside of `correlationsetup`, you still need to run the `correlationsetup` wizard to completion (choosing any existing instance) so that the newly generated certificate can be submitted to repository.
2. As the user `root`, run the command

```
<ARCSIGHT_HOME>/bin/setup_services.sh
```

This re-enables `monit`. Local services are also restarted.
3. Run

```
/etc/init.d/arcsight_services stop all
```

This is necessary because local services were started by the previous step.

4. Run

```
/etc/init.d/arcsight_services start repo
```

Final steps:

Run these commands on the persistor as the user arcsight.

1. Run

```
<ARCSIGHT_HOME>/bin/arcsight certadmin -list submitted
```

Review the output to verify that the certificates represent the cluster nodes that were reconfigured. To view the certificate details, use the `-v` option.

2. After you ensure that the certificates are correct, run

```
<ARCSIGHT_HOME>/bin/arcsight certadmin -approveall
```

Specify the cluster administration password that was provided when you installed ESM on the persistor node.

This approves newly-created certificates.

3. Run

```
/opt/arcsight/services/init.d/arcsight_services sshSetup
```

This is needed for the following step.

4. If one or more of the nodes with `mbus_data` or `mbus_control` instance had its hostname changed, then run:

```
bin/arcsight mbussetup
```

Select **I want to add, delete, or change Message Bus Instances**. When prompted for the number of instances, make no changes.

The message confirms you did not make changes and you are prompted to continue.

Enter **yes** to continue.

5. Run

```
/etc/init.d/arcsight_services start
```

The cluster should be running normally.

Changing the Internet Protocol Version in a Distributed Correlation Environment

This information applies to ESM in distributed mode only.

About:

You can change the Internet Protocol version (IP) of your cluster. You can change the cluster preference from IPv4 to IPv6, or from IPv6 to IPv4. You first must run the Manager Configuration wizard on the persistor node to select the IP preference, and then use the `syncpreferip` utility to synchronize the cluster nodes with the choice of IP preference you made on the persistor node.

Note: Passwordless SSH must be set up in order to be able to perform this task. See ["Setting Up Key-Based Passwordless SSH" on page 59](#) for details.

All nodes in a distributed correlation cluster must have the same IP preference. When you change the IP preference in a cluster, this change affects all nodes in the cluster. Note that **all** nodes in the cluster will have the same Internet Protocol version when you have completed this task.

To change the IP preference in a distributed correlation cluster:

1. On the persistor node in a cluster, and as user *arcsight*, run:

```
bin/arcsight managersetup
```

See ["Running the Manager Configuration Wizard" on page 131](#) for details on starting the wizard.

2. Step through the wizard until it displays Preferred IP Version. Choose **0** for IPv4 or **1** for IPv6.
3. Continue through the wizard without making any other configuration changes. After the wizard completes, you receive a message that your IP preference has changed. For example:

NOTE :

The IP preference has changed from IPv4 to IPv6.

ESM is set up in distributed mode. Please run the 'arcsight syncpreferip' tool on the persistor node to sync the IP preference across all nodes in the cluster.

Exit the wizard.

4. On the persistor node, and as user *arcsight*, run:

```
bin/arcsight syncpreferip
```

This command synchronizes the IP preference for all of the nodes in the cluster to the same IP preference you specified in the persistor node when you ran `managersetup`.

5. Stop and then start all services by running the following commands as user *arcsight*,

enabling the changes to take effect:

```
/etc/init.d/arcsight_services stop all
```

```
/etc/init.d/arcsight_services start all
```

Removing a Node from a Cluster

You can remove any cluster node that is not the persistor node. You might, for example, want to redeploy a system that is a node in your cluster.

To remove a cluster node:

1. Stop all services on the cluster, except for the information repository. Run the following commands as user *arcsight* from the persistor node:

```
/etc/init.d/arcsight_services stop all  
/etc/init.d/arcsight_services start repo
```
2. On the persistor node, run the configuration wizards and delete all services on the node you want to remove. This may require reducing repository or message bus control from three to one if there are no longer three nodes in the cluster. See ["Configuring Services in a Distributed Correlation Cluster" on page 60](#) for details on the cluster configuration wizards.
3. If `repositup` stopped the repository, run this command on the persistor:

```
/etc/init.d/arcsight_services start repo
```
4. On the node to be removed, remove services by running, as user *root*:

```
<ARCSIGHT_HOME>/bin/remove_services.sh
```

Do not attempt to add the node back into the cluster using `<ARCSIGHT_HOME>/bin/setup_services.sh`.
5. Uninstall the ESM on the node. See "Uninstalling ESM - Distributed Correlation Mode" in the *ESM Installation Guide* for details.
6. As user *arcsight*, run the information repository wizard to ensure that the node has been removed. Select Remove a Node, and if you see that the node is still in the cluster, remove it:

```
<ARCSIGHT_HOME>/bin/arcsight repositup
```

See ["Configuring Services in a Distributed Correlation Cluster" on page 60](#) for details on the information repository configuration wizard.
7. On the persistor node, start all services:

```
/etc/init.d/arcsight_services start
```

Troubleshooting and Frequently Asked Questions for Distributed Correlation Mode

What is compact mode?

Before distributed correlation mode was introduced, ESM always operated in compact mode. In compact mode, event processing occurs in the core of ESM and is not extensible. In contrast, distributed correlation mode allows you to extend processing by adding instances of services like aggregators, correlators, and repositories, among others. You can spread processing over several servers that you configure as nodes in a distributed correlation cluster. Each node can run groups of cluster services.

How does distributed correlation mode differ from compact mode in terms of what I see in the ArcSight Console?

Other than some distributed correlation audit events, there is no difference.

How do I verify that my system is running in distributed correlation mode?

Look at the `esm.distributed` property in the `<ARCSIGHT_HOME>/config/esm.properties` file. If distributed correlaton mode is enabled, the property is set to `true`. If the component is running in default mode, the property is set to `false`. Do not modify the value of the `esm.distributed` property.

How do I verify that the correlators, aggregators, and other services in my cluster are working?

On the persistor node, run `/etc/init.d/arcsight_services statusByNode` after you complete the cluster installation and configure the services.

Why are services stopped or in a prolonged initializing state on non-persistor nodes?

On the persistor node, ESM periodically runs a program that detects nodes that are not running the same version (major, minor, and patch) of ESM as the persistor node and attempts to stop all services on the nodes that do not match the persistor. To determine whether this is the cause of the problem, run `/etc/init.d/arcsight_services statusByNode` on the persistor node. If a non-persistor node is not running the same version of ESM as the persistor node, the output will contain information similar to the following:

```
Node <node identifier>.arcsight.com: available
aggregator4: unavailable
correlator7: initializing
dcache5: unavailable
```

Verify the version of ESM that is running on the node where services are not available. If the node is not running the same version of ESM as the persistor node, upgrade the mismatched node and then start the services.

If I make a change to a system in the cluster, do I have to synchronize the cluster to make that change known throughout the cluster?

When you add cluster services (such as instances of correlators, aggregators, distributed cache, or message bus) or change system content (for example, import a package or write a rule that updates an active list), that change is available to every part of the cluster. Manual synchronization is not required.

Why can't I start a service from a non-persistor node?

Access to the cluster nodes is granted through passwordless SSH. Only commands that you issue from the persistor node (through passwordless SSH) or from the local node will be successful.

Are there special memory requirements for the cluster?

Yes, the persistor needs more memory because it interacts with the distributed cache.

How are audit events generated for distributed correlation services?

Each service in a cluster node has a heartbeat. These heartbeats generate audit events.

Will content behave differently in distributed correlation mode?

Content will behave the same as in compact mode.

Where do queries run in distributed correlation mode?

Queries run on the persistor node.

How does upgrade work for distributed correlation?

You upgrade your system to compact mode, and then use configuration wizards to create a cluster by adding services such as distributed cache, message bus, correlators, and aggregators. For more information, see the [ESM Upgrade Guide](#).

I changed the friendly name for one of my correlators. Why is the change not reflected in audit events?

After making any configuration change to the cluster (including adding friendly names for services), you must start all services in the cluster. For more information, see [Starting Distributed Correlation Services](#).

Chapter 4: SSL Authentication

ESM uses Secure Socket Layer (SSL) technology for communication between the Manager and its clients, ArcSight Console, Transformation Hub, and SmartConnectors. It is not used between the Manager and the database.

Note: TLS is based on SSL 3.0, so you can read this chapter to get a better understanding of how TLS works as well.

SSL enables the Manager to authenticate its clients and communicate information over an encrypted channel, thus providing the following benefits:

- **Authentication:** Ensuring that clients send information to an authentic server and not to a machine pretending to be that server.
- **Encryption:** Encrypting information sent between the clients and the server to prevent intentional or accidental modification.

By default, clients submit a valid user name and password to authenticate with the server; however, these clients can be configured to use SSL client authentication.

Note that in the following topics in this chapter, for paths to commands on the Manager are:

- Interpret `bin/arcsight` as `<ARCSIGHT_HOME>/bin/arcsight`
- Interpret `jre/bin/keytool` as `<ARCSIGHT_HOME>/jre/bin/keytool`

SSL Authentication Terminology

- **Certificate**

A certificate is an entry in the keystore file that contains the public key and identifying information about the machine such as machine name and the authority that signs the certificate. SSL certificates are defined in the ISO X.509 standard.

- **Key pair**

A key pair is a combination of a private key and a public key that encrypts and decrypts information. A machine shares only its public key with other machines; the private key is never shared. The public and private keys are used to set up an SSL session. For details, see ["How SSL Works" on page 94](#).

- **SSL server-SSL client**

An SSL session is set up between two machines: a server and a client. In client-side SSL authentication, the server and its clients authenticate each other before communicating.

The Manager is an SSL server, while SmartConnectors, Console, and browsers are SSL clients.

- **Keystore**

A keystore file is an encrypted repository on the SSL server that holds the SSL certificate and the server's private key. The following table lists the ESM component, the name of the keystore on that component, and its location. Do not change the keystore file name.

Keystore password

Use a keystore password to encrypt the keystore file. Without this password, you cannot open the keystore file. The default is *password* for the Manager and *changeit* for the ArcSight Console's client keystore. The default password for the key pair for any component is the same as for the component's keystore.

You specify a keystore password when creating a key pair, which is discussed in later sections of this chapter. The password is obfuscated and stored in the ESM component's *.properties file. The following table lists the property name where the obfuscated keystore passwords are stored.

Keystore	Property File	Property Name
Client*	config/client.properties**	ssl.keystore.password
Manager	config/esm.properties	server.privatekey.password.encrypted Default password is password.
Connector	user/agent/agent.properties**	ssl.keystore.password.encrypted Default password is changeit.

*For client-side authentication

** If config/client.properties or user/agent/agent.properties does not exist, create it using an editor of your choice.

Whenever you change a password for the keystore, you must make the same change in the password entry in the corresponding properties file.

- **Truststore**

Truststore is an encrypted repository on SSL clients that contains a list of certificates from the issuers that a client trusts. Use either the `keytool` or `keytoolgui` command to view a truststore. See ["View Certificate Details From the Store"](#) on [page 106](#) for details on viewing a truststore.

A certificate is signed by the issuer with its private key. When the server presents this certificate to the client, the client uses the issuer's public key from the certificate in its truststore to verify the signature. If the signature matches, the client accepts the certificate. For more details, see how SSL handshake occurs in ["How SSL Works" on the next page](#).

- **Alias**

Certificates and key pairs in a keystore or a truststore are identified by an alias.

- **Truststore password**

The `*.defaults.properties` file contains the default truststore password for each ESM component (By default this password is *changeit*). Use a truststore password to encrypt a truststore file. Without this password, you cannot open the truststore file. The password is in clear text. To change or obfuscate it, use the [changepassword](#) command.

The following table lists the property name where the obfuscated truststore passwords are stored.

Truststore	Property File	Property Name
Client	<code>client.properties**</code>	<code>ssl.truststore.password.encrypted</code>
Manager*	<code>server.properties</code>	<code>servletcontainer.jetty311.truststore.password.encrypted</code>
Connector	<code>agent.properties**</code>	<code>ssl.truststore.password</code>

*For client-side authentication

** If `config/client.properties` or `user/agent/agent.properties` does not exist, create it using an editor of your choice.

Whenever you change a password for the truststore, you must make the same change in the password entry in the corresponding properties file.

Understanding Cipher Suites

In general, cipher suites are a set of authentication, encryption, and data integrity algorithms used for securely exchanging data between an SSL server and a client.

The cipher suites that are enabled are configured by ArcSight Wizards in property files. Although in most cases you do not need to change the cipher suites, you can configure them in the corresponding properties file for an ArcSight component:

Component	Property File	Property
Manager	config/esm.properties	servletcontainer.jetty311.socket.https.ciphersuites
Clients	config/client.properties	ssl.cipher.suites
Connectors	user/agent/agent.properties	ssl.cipher.suites

Cipher suites are set as a comma-delimited list. During the SSL handshake, the endpoints provide these lists as the cipher suites that they can accept, in descending order of preference. One of the cipher suites is chosen by SSL negotiation process and that cipher suite is used for the entire communication session between these two components. This means that in order to limit cipher suites, it is sufficient to restrict the list of enabled cipher suites on one side only, for example, on the Manager side.

For information on specific cipher suites for FIPS encryption, see ["FIPS Encryption Cipher Suites" on page 230](#)


How SSL Works

When a client initiates communication with the SSL server, the server sends its certificate to authenticate itself to the client. The client validates the certificate by verifying:

- The hostname is identical to the one with which the client initiated communication.
- The certificate issuer is in the list of trusted certificate authorities in the client's truststore (<ARCSIGHT_HOME>/jre/lib/security/cacerts) and the client is able to verify the signature on the certificate by using the CA's public key from the certificate in its truststore.
- The current time on the client machine is within the validity range specified in the certificate to ensure that the certificate is valid.

If the certificate is validated, the client generates a random session key, encrypts it using the server's public key, and sends it to the server. The server decrypts the session key using its private key. This session key is used to encrypt and decrypt data exchanged between the server and the client from this point forward.

The following figure illustrates the handshake that occurs between the client and Manager.



With client-side authentication, the server requests the client's certificate when it sends its certificate to the client. The client sends its certificate along with the encrypted session key.

Certificate Types

There are three types of SSL certificates:

- CA-signed
- Self-signed

- Demo (applicable to default mode only)

CA-signed certificates are issued by a third party you trust. The third party may be a commercial Certificate Authority (CA) such as VeriSign and Thawte or you might have designated your own CA. Because you trust this third party, your client's truststores might already be configured to accept its certificate. Therefore, you may not have to do any configuration on the client side. See ["Using a CA-Signed SSL Certificate" on page 112](#).

You can create your own self-signed certificates. A self-signed certificate is signed using the private key from the certificate itself. Each server is an issuer. Configure clients to trust each self-signed certificate you create.

Self-signed certificates are as secure as CA-signed, however, CA-signed certificates scale better as illustrated in this example:

If you have three SSL servers that use self-signed certificates, you configure your clients to accept certificates from all of them (the three servers are three unique issuers). If you add a new server, you configure all the clients, again, to accept the additional certificate. However, if these servers use a CA-signed certificate, all servers use copies of the same one. You only configure the clients once to accept that certificate. If the number of Managers grows in the future, you do not need to do any additional configuration on the clients.

Demo certificates are useful in isolated test environments. Using one in a production environment is not recommended.

SSL Certificate Tasks

The command `<ARCSIGHT_HOME>/bin/arcsight keytool` (runs from the command line in a terminal window) and the command `keytoolgui` (provides a graphical user interface) enable you to perform SSL certificate configuration tasks.

Note: We recommend that you use `bin/arcsight keytool`, which does not require the X Window system. Using the `keytoolgui` interface requires that the X Window system be installed on your system, and only works in a non-FIPS implementation. Using the X Window system is not preferred on the Manager machine. Also, note that the X Window system is not present on an appliance. The command `keytoolgui` is not supported on the Mac, so for managing the keystore and certificates and so on, on a Mac, use `bin/arcsight keytool`. The command `keytoolgui` is not supported in FIPS mode, so for managing the keystore and certificates, use `bin/arcsight keytool`.

The command `bin/arcsight keytool` simplifies usage by pre-populating several command line arguments as defaults of the command based on component's configured

values. The following subsections discuss the use of `bin/arcsight keytool`, and show the simplified command lines. The command `bin/arcsight keytool` provides default values for the following parameters based on the value of the `-store` parameter:

- `-keystore`
- `-storetype`
- `-storepass`
- `-keypass`
- `-srckeystore`
- `-srcstoretype`
- `-srcstorepass`
- `-destkeystore`
- `-deststoretype`
- `-deststorepass`

The table below shows what values each of the different `-store` values uses by default.

-store	Description	Configuration Files	Password Property	Keystore Property	Keystore non-FIPS	Keystore FIPS
managercerts	Truststore for Manager	config/esm.properties	servletcontainer.jetty311.truststore.password	servletcontainer.jetty311.truststore.file	config/jetty/truststore	config/jetty/keystore.bcdfs
managerkeys	Keystore for Manager	config/esm.properties	server.privatekey.password	servletcontainer.jetty311.keystore.file	config/jetty/keystore	config/jetty/keystore.bcdfs
clientcerts	Truststore for Console	config/client.properties	ssl.truststore.password	ssl.truststore.path	jre/lib/security/cacerts	config/keystore.client.bcdfs
clientkeys	Keystore for Console	config/client.properties	ssl.keystore.password	ssl.keystore.path	config/keystore.client	config/keystore.client.bcdfs

Parameters ending in `storetype` default to JKS in the non-FIPS case, and BCFKS for FIPS. Default values for parameters ending in `pass` can be found by looking up the password property in the configuration file. Default values for keystore parameters are found by looking up the keystore property in the configuration file. Sometimes it is not defined there, in which case the FIPS or non-FIPS default is used, depending on the case.

The following sections present `bin/arcsight keytool` command lines that are exactly formed to perform the task mentioned in the section. Use only those options to perform the documented tasks.

The command `jre/bin/keytool` can also be used for the SSL certificate tasks. For details on `jre/bin/keytool`, see online vendor documentation. Various vendors have their own version of `keytool`. One reference is <http://docs.oracle.com/javase/8/docs/technotes/tools/windows/keytool.htm>. Note that if you use `keytool -h` to view Help you will see options that are not covered in this documentation. The `keytool` examples presented in this guide do not display all possible `keytool` options. Refer to the `keytool` documentation online for the meaning of parameters.

Export a Key Pair

You can use `bin/arcsight keytool` to export a key pair.

Use of `bin/arcsight keytool` (which runs from the command line in a terminal window) is recommended. Using the `keytoolgui` interface requires that the X Window system be installed on your system. Also, the X Window system is not present on ESM on an appliance.

Exporting a Key Pair Using `bin/arcsight keytool`

An example of a `bin/arcsight keytool` command line is provided. Use this example as a basis to form the command line you need. Normally done to import it into a browser or connector.

To export a key pair with the alias `admin` into a file named `admin.p12` from the client keystore :

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store clientkeys -importkeystore -  
destkeystore admin.p12 -deststoretype PKCS12 -srcalias admin
```

Exporting a Key Pair Using `keytoolgui`

To use `keytoolgui`:

1. Start `keytoolgui` from the component from which you want to export the key pair. To do so, run the following command from the component's `<ARCSIGHT_HOME>/bin` directory:

```
./arcsight keytoolgui
```
2. Click **File->Open keystore** and navigate to the component's keystore.
3. Enter the password for the keystore when prompted. For the default password see

["Keystore password" on page 92.](#)

4. Right-click the key pair and select **Export**.
5. Select **Private Key and Certificates** radio button and click **OK**.
6. Enter the password for the key pair when prompted. For the default password see ["Keystore password" on page 92.](#)
7. Navigate to the location on your machine to where you want to export the key pair.
8. Enter a name for the key pair with a `.pfx` extension in the **Filename** text box and click **Export**. You get an Export Successful message.
9. Click **OK**.

Import a Key Pair

You can use `keytool` to import a key pair.

Use of `bin/arcsight keytool` (which runs from the command line in a terminal window) is recommended. Using the `keytoolgui` interface requires that the X Window system be installed on your system. Also, the X Window system is not present on ESM on an appliance.

Importing a Key Pair Using `bin/arcsight keytool`

An example of a `keytool` command line is provided. Use this example as a basis to form the command line you need.

To import a key pair with the alias `admin` from a file named `admin.p12` into the client keystore:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store clientkeys -importkeystore -  
srckeystore admin.p12 -srcstoretype PKCS12 -srcalias admin
```

Importing a Key Pair Using `keytoolgui`

1. Start `keytoolgui` from the component to which you want to import the key pair. To do so, run the following command from the component's `<ARCSIGHT_HOME>/bin` directory.

```
./arcsight keytoolgui
```
2. Select **File->Open keystore** and navigate to your component's keystore.
3. Enter the keystore password when prompted. For the default password see ["Keystore password" on page 92.](#)
4. Select **Tools->Import Key Pair** and navigate to the location of the key pair file, select it and click **Choose**.

5. Enter the password for the key pair file when prompted and click **OK**. For the default password see ["Keystore password" on page 92](#).
6. Select the key pair and click **Import**.
7. Enter an alias for the key pair and click **OK**.
8. Enter a new password for the key pair file to be imported, confirm it, and click **OK**. You see a message saying Key Pair Import Successful.
9. Click **OK**.
10. Select **File->Save keystore** to save the changes to the keystore and exit the keytoolgui.

Export a Certificate

You can use `bin/arcsight keytool` to export a certificate.

Use of `bin/arcsight keytool` (which runs from the command line in a terminal window) is recommended. Using the `keytoolgui` interface requires that the X Window system be installed on your system. Also, the X Window system is not present on ESM on an appliance.

Exporting a Certificate Using `bin/arcsight keytool`

An example of a `bin/arcsight keytool` command line is provided. Use this example as a basis to form the command line you need. The example shown below is that of exporting a certificate associated with a key, which applies in most in ESM use cases.

Note that if the alias points to a trusted certificate, the output is that certificate. Also, if the alias points to a key entry, the output is the first certificate from key's certificate chain.

For example:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store clientkeys -exportcert -alias  
admin -file admin.cer -rfc
```

Exporting a Certificate Using `keytoolgui`

1. Start `keytoolgui` from the component from which you want to export the certificate. To do so, run the following command from the component's `<ARCSIGHT_HOME>/bin` directory.

```
./arcsight keytoolgui
```
2. Select **File->Open keystore** and navigate to your component's truststore.
3. Enter the truststore password when prompted. For the default password see ["Truststore password" on page 93](#).

4. Right-click the certificate and select **Export**.
 - a. Select **Head Certificate** as Export Type and **DER Encoded** as the Export Format and click **OK**:
 - b. Navigate to the location where you want to export the certificate, and enter a name for the certificate with a `.cer` extension and click **Export**.
 - c. You see the **Export Successful** message
5. If the component into which you want to import this certificate resides on a different machine than the machine from which you exported the certificate (the current machine), copy this certificate to the to the other machine.

Import a Certificate

You can use `bin/arcsight keytool` to import a certificate.

Use of `bin/arcsight keytool` (which runs from the command line in a terminal window) is recommended. Using the `keytoolgui` interface requires that the X Window system be installed on your system. Also, the X Window system is not present on ESM on an appliance.

Importing a Certificate Using `bin/arcsight keytool`

An example of a `bin/arcsight keytool` command line is provided. Use this example as a basis to form the command line you need. Certificates should always be imported into `clientcerts` or `managercerts`. `bin/arcsight keytool` will ask if you want to trust this certificate; answer **Yes**.

For example:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store managercerts -importcert -alias  
admin -file admin.cer
```

Importing a Certificate Using `keytoolgui`

1. Start `keytoolgui` from the component into which you want to import the certificate. To do so, run the following command from the component's `<ARCSIGHT_HOME>/bin` directory.

```
./arcsight keytoolgui
```
2. Click **File->Open keystore** and navigate to the truststore (`<ARCSIGHT_HOME>/jre/lib/security`) of the component.
3. Select the store named `cacerts` and click **Open**.
4. Enter the password for the truststore when prompted. For the default password see ["Truststore password" on page 93](#).

5. Click **Tools->Import Trusted Certificate** and navigate to the location of the certificate that you want to import.
6. Click **Import**.
7. You see the message
Could not establish a trust path for the certificate. The certificate information will now be displayed after which you may confirm whether or not you trust the certificate.
Click **OK**.
8. The Certificate details are displayed. Click **OK**.
9. You see the message **Do you want to accept the certificate as trusted?.** Click **Yes**.
10. Enter an alias for the Trusted Certificate you just imported and click **OK**.
Typically, the alias Name is same as the fully qualified host name (for example *devgroup.topco.com*).
11. You see the message **Trusted Certificate Import Successful..** Click **OK**.
12. Save the truststore file.

Creating a Keystore

You can use `/jre/bin/keytool` to create a keystore.

Using the `keytoolgui` is not preferred, and the interface requires that the X Window system be installed on your system. Also, the X Window system is not present on ESM on an appliance.

Note: Generally, you will only need to create the non-FIPS keystore for a client. Also, keystores are created automatically when you generate a keypair to add to a keystore. If a keystore does not exist, it gets created automatically when the first item is put into it.

Creating a Keystore Using `jre/bin/keytool`

An example of a `jre/bin/keytool` command line is provided below. Use this example as a basis to form the command line you need. Note that this command does not use the `bin/arcsight keytool` wrapper and requires more options be specified than some other `keytool` commands.

The abbreviations in the command below denote the following fields: `cn` = Common Name, `ou` = Organizational Unit, `o` = Organization, and `c` = Country.

The command generates a new self-signed certificate with `ALIAS_NAME` in the specified keystore `PATH_TO_KEYSTORE`.

Example for a new keystore:

```
<ARCSIGHT_HOME>/jre/bin/keytool -genkeypair -keystore config/keystore.client  
-storetype JKS -storepass password -dname "cn=John Smith, ou=ArcSight, o=MF,  
c=US" -alias testKey -validity 365
```

Specify all the options in the above example using the appropriate values for your installation.

As a separate operation, either before or after you run the `genkeypair` command, you have to set the values for the keystore location, keystore type, and password in the `client.properties` file. This file is in `<ARCSIGHT_HOME>/config`. The Console uses this file to access the keystore during authentication.

The `client.properties` file works as an override for the `client.defaults.properties` file. (You do not edit the default properties file because it is overwritten at upgrade time.) Set these properties in `client.properties`, as follows:

- **ssl.keystore.path**= Set this value if it differs from the default in `client.defaults.properties`. It must be the same as the path specified in the `-keystore` option in the command example, above.
- **ssl.keystore.type**= Set this value if it differs from the default in `client.defaults.properties`. It must be the same as the path specified in the `-storetype` option in the command example, above.
- **ssl.keystore.password**=Set this value if it differs from the default in `client.defaults.properties`. It must be the same as the password specified in the `-storepass` option in the command example, above. The default is blank (no password), but having a password is recommended.

However, if you plan to encrypt the password (also recommended), there is no need to set it manually in this file. You specify it and encrypt it using the `changepassword` command, next.

To set an encrypted password, run the following command:

```
arcsight changepassword -f config/client.properties -p ssl.keystore.password
```

This command prompts you for the actual password, adds it to the `client.properties` file, and encrypts it. It must be the same as the password specified in the `-storepass` option in the command example, above.

Creating a Keystore Using `keytoolgui`

1. Start `keytoolgui` from the component into which you want to import the certificate. To do so, run the following command from the component's `<ARCSIGHT_HOME>/bin` directory.

```
./arcsight keytoolgui
```

2. Click **File->New keystore**.
3. Select **JKS** and click **OK**.
4. Click **File->Save keystore**.

Generating a Key Pair

You can use `bin/arcsight keytool` to generate a key pair.

Use of `bin/arcsight keytool` (which runs from the command line in a terminal window) is recommended. Using the `keytoolgui` interface requires that the X Window system be installed on your system. Also, the X Window system is not present on ESM on an appliance.

Generating a Key Pair Using `bin/arcsight keytool`

The abbreviations in the command below denote the following fields: `cn` = Common Name, `ou` = Organizational Unit, `o` = Organization, and `c` = Country.

To generate a key for client authorization, make sure that `sssl.keystore.path` is set in `client.properties`, and then run the command shown below:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store clientkeys -genkeypair -dname "cn=John Smith, ou=ArcSight, o=MF, c=US" -keyalg rsa -keysize 2048 -alias admin -startdate -1d -validity 366
```

This creates a key valid starting yesterday (to avoid problems with clock skew between servers), and expiring in about one year. Make sure the `cn` value matches that of the External ID of the user you log in as.

Note: We strongly recommend the use of RSA keys with a keysize of 2048 as client keys. Some browsers have known issues with elliptic curve keys.

About the only time you should need to change the manager key is if you change the hostname of the manager. You should never need to create a manager key in non-FIPS mode; `managersetup` will take care of that. Instructions for creating a manager key for FIPS mode are given below.

Before adding a new manager key, be sure to delete the old one. It has the alias `mykey`.

For FIPS 140-2 create a RSA 2048 key:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store managerkeys -genkeypair -dname "cn=myhost.mydomain.com, ou=ArcSight, o=MF, c=US" -keyalg rsa -keysize 2048 -alias mykey -startdate -1d -validity 366
```

The `cn` value must be the ESM hostname for all manager keys - regardless of the type.

Note: FIPS Suite B requires elliptic curve keys. The minimum length for 128 bit is 256 bits, and for 192 bits it is 384 bits. Some browsers will not work with elliptic curve keys longer than 384 bits. So 384 bits, as shown below is a good choice for FIPS Suite B.

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store managerkeys -genkeypair -dname "cn=myhost.mydomain.com, ou=ArcSight, o=MF, c=US" -keyalg ec -keysize 384 -alias mykey -startdate -1d -validity 366
```

Verifying Whether a Key Pair Has Been Successfully Generated

To verify whether the key pair has been successfully created in the keystore, run the following from the component's <ARCSIGHT_HOME>/bin directory:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store managerkeys -list
```

Setting the Expiration Date of a Certificate

To set the expiry date of the certificate, do it when generating the key pair. After you have generated the key pair, you cannot change the expiration date on the certificate and the certificate expires in three months by default.

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store clientkeys -genkeypair -dname "cn=John Smith, ou=ArcSight, o=MF, c=US" -keyalg rsa -keysize 2048 -alias admin -startdate -1d -validity 366
```

You specify the validity of the certificate with the `-validity <number_of_days>` option. The value that you provide with `-validity` calculates the number of days that the certificate is valid starting from the current time.

Generating a Key Pair Using `keytoolgui`

1. Start `keytoolgui` from the component into which you want to import the certificate. To do so, run the following command from the component's <ARCSIGHT_HOME>/bin directory.

```
./arcsight keytoolgui
```
2. Click **File->Open keystore** and navigate to your keystore.
3. Click **Tools->Generate Key Pair** and fill in the fields in the General Certificate dialog and click **OK**.
4. Enter an alias for the newly created key pair and click **OK**.
5. Save the keystore by clicking **File->Save keystore**.

View Certificate Details From the Store

You can use `bin/arcsight keytool` to view certificate details from the keystore (list the entries in a keystore).

Use of `bin/arcsight keytool` (which runs from the command line in a terminal window) is recommended. Using the `keytoolgui` interface requires that the X Window system be installed on your system. Also, the X Window system is not present on ESM on an appliance.

Viewing a Certificate Details from the Store Using `bin/arcsight keytool`

An example of a `bin/arcsight keytool` command line is provided. Use this example as a basis to form the command line you need.

By default, `clientcerts` has 100 or so certificates in it. The `-v` option lists details about each certificate, so the total output will be approximately 1,000 lines. If you do not use `-v`, the command will return one line per certificate. Add the option `-alias mycert` to only see details about the certificate with the alias `mycert`.

To list details about all keys:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store clientcerts -list -v
```

To print details for the key with the specified alias:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store managerkeys -list -v -alias mykey
```

Viewing a Certificate Details from the Store Using `keytoolgui`

For certificates in the keystore or truststore use the `keytoolgui` command to see certificate information.

1. Start `keytoolgui` from the component from which you want to export the certificate. To do so, run the following command from the component's `<ARCSIGHT_HOME>/bin` directory.

```
./arcsight keytoolgui
```
2. Select **File->Open keystore** and navigate to your component's truststore.
3. Enter the truststore password when prompted. For the default password see ["Truststore password" on page 93](#).
4. Double-click the certificate whose details you want to view. Details include valid date range, and other information about the certificate.

For the Manager certificate you can also use `tempca -i` command.

Delete a Certificate

You can use `bin/arcsight keytool` to delete a certificate from the keystore.

Use of `bin/arcsight keytool` (which runs from the command line in a terminal window) is recommended. Using the `keytoolgui` interface requires that the X Window system be installed on your system. Also, the X Window system is not present on ESM on an appliance.

Deleting a Certificate Using `bin/arcsight keytool`

An example of a `keytool` command line is provided. Use this example as a basis to form the command line you need.

To remove a third party trusted certificate with alias `rootCA`:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store managercerts -delete -alias  
rootCA
```

Deleting a Certificate Using `keytoolgui`

To delete a certificate from the `truststore`, start `keytoolgui` and navigate to the certificate, right-click on the certificate, and select **Delete**.

Changing Keystore/Truststore Passwords

It is a good security practice to change the keystore and truststore passwords after installing ESM or ESM console. In addition to changing the keystore password, you need to separately change the value that ESM uses for this password, so that ESM can continue to access the keystore. FIPS has a single shared keystore/truststore, so the keystore and truststore passwords must be the same. Changing passwords using `bin/arcsight changepassword` is recommended since this program will encrypt the passwords in the configuration file.

Note: Key pairs also have passwords. ESM expects that these passwords will be the same as the keystore passwords, so both must be changed.

Below is an example of how to change the passwords on the Manager keystore.

Note: These steps must be performed in the order given.

1. `/etc/init.d/arcsight_services stop manager`
2. `bin/arcsight keytool -store managerkeys -keypasswd -alias mykey`

The command `keytool` will prompt for the new password.

3. `bin/arcsight keytool -store managerkeys -storepasswd`

The command `keytool` will prompt for the new password. Enter the same password as for step 2.

4. `bin/arcsight changepassword -f config/esm.properties -p server.privatekey.password`

The command `changepassword` will prompt for the new password. Enter the same password as for step 2.

5. `/etc/init.d/arcsight_services start all`

Here is an example of how to change the password on a Console truststore to match that of the console keystore. This can be needed to convert a default mode installation (with separate keystore/truststore) to FIPS mode with a single keystore/truststore. The console should not be running. Note that no `keytool -keypasswd` command is needed, as there are no keys in the truststore.

1. `bin/arcsight keytool -store clientcerts -storepasswd`

The command `keytool` will prompt for the new password. Enter the password for the `clientcerts` keystore.

2. `bin/arcsight changepassword -f config/client.properties -p ssl.truststore.password`

The command `changepassword` will prompt for the new password. Enter the password for the `clientcerts` keystore.

Using a Self-Signed Certificate

When dealing with certificate based identification and encryption, components fall into one of two categories: servers and clients. Signed certificates enable these components to verify the validity of communications with the other components. You can use either a self-signed certificate or a CA-signed certificate when setting up SSL authentication on your ESM components.

The procedure you follow depends on the number of Managers with which your clients communicate, because each Manager will have its own self-signed certificate, and any client that has to communicate with different Managers has to be configured to accept all those Manager's certificates.

When Clients Communicate With One Manager

To use a self-signed certificate for deployments in which clients communicate with only one Manager, perform these steps:

1. On the Manager, create a self-signed key pair:

Note: Steps to create a self-signed key pair may be different for a new Manager installation as the Configuration Wizard is launched automatically during the installation process.

- a. In `<ARCSIGHT_HOME>/bin`, run the following command:

```
./arcsight managersetup
```
- b. In the Manager Configuration Wizard, select **Replace with new Self-Signed key pair** and click **Next**.
- c. Enter information about the SSL certificate and click **Next**.
- d. Enter the SSL keystore password for the certificate and click **Next**. Remember this password. You will use it to open the keystore.
- e. Continue through the Configuration Wizard.

The Configuration Wizard does these three SSL-related things:

- It replaces the Manager's keystore at, `<ARCSIGHT_HOME>/config/jetty/keystore`, with the one created using this procedure.
- It generates the `selfsigned.cer` certificate file in the `<ARCSIGHT_HOME>/config/jetty` directory.
- It overwrites the existing Manager truststore file, `<ARCSIGHT_HOME>/jre/lib/security/cacerts`, with one containing the new self-signed certificate to the Manager's truststore file.

The new `cacerts` file contains the information about the Trusted Certificate Authority (CA) that signed your self-signed certificate.

The self-signed certificate does not take effect until the Manager and clients are restarted later in this procedure.

2. If you are running ESM in distributed mode, complete the following:

- a. Run the following command:

```
./arcsight certadmin -list submitted
```

Review the output to verify that the Manager certificate is listed. To view the certificate details, use the `-v` option.

- b. After you confirm that the Manager certificate is correct, run the following command:

```
./arcsight certadmin -approveall
```

Specify the cluster administration password that was provided when you installed ESM on the persistor node.

3. Export the Manager's certificate from `<ARCSIGHT_HOME>/jre/lib/security/cacerts`.

4. Copy the Manager's certificate to each machine from which clients connect to the Manager.
5. On those clients, import the Manager's certificate to the `<ARCSIGHT_HOME>/jre/lib/security/cacerts` directory. See ["Import a Certificate" on page 101](#).

Note: Make sure you have imported the Manager's certificate to all existing clients before proceeding further. Otherwise, after you perform the next steps, only clients with the new Manager's certificate can connect to the Manager.

6. Stop and start all services so that the new self-signed certificate will take effect:

```
/etc/init.d/arcsight_services stop
```

```
/etc/init.d/arcsight_services start
```

7. Restart all clients.
8. When installing a new client, repeat Steps 2-4 of this procedure.
9. Optionally, if SSL client-side authentication is needed, on the ArcSight Console, perform the steps listed in section ["Setting up SSL Client-Side Authentication on ArcSight Console- Self-Signed Certificate" on page 122](#)

When Clients Communicate With Multiple Managers

This procedure is for using a self-signed certificate where clients communicate with more than one Manager. In this procedure you get the self-signed certificate files from each manager, copy them to a client, import them all into that client, then copy that client cacerts file to all your other clients.

1. On the Manager, create a self-signed key pair:

Note: Steps to create a self-signed key pair may be different for a new Manager installation as the Configuration Wizard is launched automatically during the installation process.

- a. In `<ARCSIGHT_HOME>/bin`, run the following command:

```
./arcsight managersetup
```
- b. In the Manager Configuration Wizard, select **Replace with new Self-Signed key pair** and click **Next**.
- c. Enter information about the SSL certificate and click **Next**.
- d. Enter the SSL keystore password for the certificate and click **Next**. Remember this password. You will use it to open the keystore.
- e. Continue through the Configuration Wizard.

The Configuration Wizard does these three SSL-related things:

- It replaces the Manager's keystore at, <ARCSIGHT_HOME>/config/jetty/keystore, with the one created using this procedure.
- It generates the `selfsigned.cer` certificate file in the <ARCSIGHT_HOME>/config/jetty directory.
- It overwrites the existing Manager truststore file, <ARCSIGHT_HOME>/jre/lib/security/cacerts, with one containing the new self-signed certificate to the Manager's truststore file.

The new `cacerts` file contains the information about the Trusted Certificate Authority (CA) that signed your self-signed certificate.

The self-signed certificate does not take effect until the Manager and clients are restarted later in this procedure.

2. If you are running ESM in distributed mode, complete the following:

a. Run the following command:

```
./arcsight certadmin -list submitted
```

Review the output to verify that the Manager certificate is listed. To view the certificate details, use the `-v` option.

b. After you confirm that the Manager certificate is correct, run the following command:

```
./arcsight certadmin -approveall
```

Specify the cluster administration password that was provided when you installed ESM on the persistor node.

3. Copy the `selfsigned.cer` file from each Manager to the <ARCSIGHT_HOME>/jre/lib/security directory on one of your clients.

The certificate files all have the same name. Rename each one so they do not overwrite another on the client. For example, rename the certificate file from `ManagerA` to `SelfSigned_MgrA.cer`.

4. On that client, use the `keytool` or `keytoolgui` command to import certificates into the truststore (`cacerts`):

The `keytool` command is preferred. Using the `keytoolgui` interface requires that the X Window system be installed on your system. Note that using the X Window system is not preferred, but if you have it installed and want to use it, you can use `keytoolgui`. The X Window system is not present on ESM on an appliance. See ["Import a Certificate" on page 101](#) for details on using `keytool`.

To use the `keytoolgui` command:

a. In <ARCSIGHT_HOME>/bin, run this command:

```
./arcsight keytoolgui
```

b. Click **File->Open keystore**.

- c. In `<ARCSIGHT_HOME>/jre/lib/security`, select the store named `cacerts`.
- d. Click **Tools->Import Trusted Certificate**:
 - i. Select the self-signed certificate for a Manager and click **Import**.
 - ii. You see the message:
Could not establish a trust path for the certificate. The certificate information will now be displayed after which you may confirm whether or not you trust the certificate.
Click **OK**.
The Certificate details are displayed. Click **OK**.
 - iii. When asked if you want to accept the certificate as trusted, click **OK**.
 - iv. Enter an alias for the Trusted Certificate you just imported and click **OK**.
Typically, the alias Name is same as the fully qualified host name.
 - v. You see the message **Trusted Certificate Import Successful..** Click **OK**.
 - vi. Save the truststore file (`cacerts`).
 - vii. Repeat Steps i through vi for all self-signed certificates you copied.
- e. On the client, enter this command in `<ARCSIGHT_HOME>/bin` to stop the client from using the Demo certificate:

```
./arcsight tempca -rc
```

For SmartConnectors, run:

```
./arcsight agent tempca -rc
```

5. Stop and start all services so that the new self-signed certificate will take effect:

```
/etc/init.d/arcsight_services stop
```

```
/etc/init.d/arcsight_services start
```

6. Restart the client.
7. Copy the `cacerts` file to all your other clients and restart them. If you install a new client, copy the `cacerts` file to it as well.

Using a CA-Signed SSL Certificate

Using a certificate signed by a Certificate Authority means replacing your demo or self-signed certificate. Follow the procedures described in this section to obtain and import the certificate into the Manager.

Obtaining and deploying a CA-signed certificate involves these steps:

1. ["Create a Key Pair for a CA-Signed Certificate" below.](#)
2. ["Send for the CA-Signed Certificate" on the next page.](#)
3. [" Import the CA Root Certificate" on page 115.](#)
4. ["Import the CA-Signed Certificate" on page 115.](#)
5. ["Start the Manager Again \(Restart the Manager\)" on page 118.](#)
6. ["Using CA-Signed Certificates with Additional Components" on page 119.](#)
7. Optionally, if SSL client-side authentication is needed, on the ArcSight Console, perform the steps listed in section ["Setting up SSL Client-Side Authentication on ArcSight Console- Self-Signed Certificate" on page 122](#)

Create a Key Pair for a CA-Signed Certificate

To create a key pair, the `keytool` command is preferred. Using the `keytoolgui` interface requires that the X Window system be installed on your system. Note that using the X Window system is not preferred, but if you have it installed and want to use it, you can use `keytoolgui`. The X Window system is not present on ESM on an appliance. See ["Generating a Key Pair" on page 104](#) for details on using `keytool`.

To use the `keytoolgui` command:

1. On the Manager machine, run this command to launch `keytoolgui` in `<ARCSIGHT_HOME>/bin`:

```
./arcsight keytoolgui
```
2. Click **File->New keystore** to create a new keystore.
3. Select **JKS** for the keystore Type, it supports Java keystore:
4. Click **Tools->Generate Key Pair** to create the key pair. This can take some time.
5. Enter key pair information such as the length of time for its validity (in days). Click **OK**.

For **Common Name (CN)**, enter the fully qualified domain name of the Manager. Ensure that DNS servers, used by the clients connecting to this host, can resolve this host name.

For **Email(E)**, provide a valid e-mail address as the CAs typically send an e-mail to this address to renew the certificate.

When you click **OK** it asks you for a new password. Use the password of your existing keystore to save this one. The Manager may fail to start if the password of the Key pair does not match the password of the keystore encrypted in `esm.properties`. If you do not remember the password, run the Manager setup Wizard and change the password of your existing keystore before you proceed. You reuse this file after receiving the reply from the CA.

6. Specify an alias name of *mykey* for referring to the new key pair.
7. Click **File->Save as** and save the keystore with a name such as *keystore.request*.

Send for the CA-Signed Certificate

To send for the CA-signed certificate, first create a certificate signing request (CSR).

You can use `keytool` to send for a CA-signed certificate. Use of `keytool` (which runs from the command line in a terminal window) is preferred. Using the `keytoolgui` interface requires that the X Window system be installed on your system. Note that using the X Window system is not preferred, but if you have it installed and want to use it, you can use `keytoolgui`. The X Window system is not present on ESM on an appliance.

Sending a CA-Signed Certificate Using `keytool`

An example of a `keytool` command line is provided. Use this example as a basis to form the command line you need.

For example:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -certreq -store managerkeys -alias  
testkey -file config/testkey.csr
```

The command creates signing request using the PKCS#10 format for a certificate with alias `<ALIAS_NAME>` from `keystore_path`. Here `<storepass>` is keystore password, and `<keypass>` is a password for the specified alias. No need to be specified for empty values. As a result the command creates a file `<testkey.csr>` that should be sent to certificate authority (CA).

After verifying the information you sent, the CA electronically signs the certificate using its private key and replies with a certification response containing the signed certificate (`cer-file`).

Sending a CA-Signed Certificate Using `keytoolgui`

1. In `keytoolgui`, right-click the *mykey* alias name and select **Generate CSR** to create a Certificate Signing Request.
2. Choose a path and filename, and click **Generate**.
After you enter a file name, the CSR file is generated in the current working directory.
3. Send the CSR to the selected Certificate Authority (CA).

After verifying the information you sent, the CA electronically signs the certificate using its private key and replies with a certification response containing the signed certificate.

Import the CA Root Certificate

When you get the response from the certificate authority, it should include instructions for getting the root CA certificate. You can skip this step if renewing a CA-signed certificate issued by the same root certificate authority. You import the CA root certificate into the truststore file.

To create a key pair, the `keytool` command is preferred. Using the `keytoolgui` interface requires that the X Window system be installed on your system. Note that using the X Window system is not preferred, but if you have it installed and want to use it, you can use `keytoolgui`. The X Window system is not present on ESM on an appliance. See ["Import a Certificate" on page 101](#) for details on using `keytool`.

1. Save the Root CA certificate as a file `rootca.cer`.
2. Repeat the following procedure on all the machines where the Manager is installed:
 - a. Launch `keytoolgui` on the Manager machine.
 - b. Click **File > Open keystore**.
 - c. Select the Truststore file located at `<ARCSIGHT_HOME>/jre/lib/security/cacerts`. Use the default password to open `cacerts`.
 - d. Click **Tools > Import Trusted Certificate**, and pick the `rootca.cer` file.
 - e. You see the following warning message:
"Could not establish a trust path for the certificate. The certificate information will now be displayed after which you may confirm whether or not you trust the certificate."
 - f. Click **OK** to finish.

Note: Hints on importing the CA root certificate:

- If the CA root certificate has a chain, follow the same procedure to import all intermediate CA certificates into the Truststore.
- Update the CA root certificate on other ESM components, as well.
 - Repeat step 2 of the procedure on one of the Consoles.
 - Copy the updated `cacerts` to any Logger, and other machines with Consoles or Connectors.
- Restart all services after the new `cacerts` is copied.

Import the CA-Signed Certificate

When the CA has processed your request, it sends you a file with the signed certificate. You import this certificate into the Manager's keystore.

The SSL certificate you receive from the Certificate Authority must be a 128-bit X.509 Version 3 certificate. The type of certificate is the same one that is used for common web servers. The signed certificate must be returned by the CA in base64 encoded format. It looks similar to this:

```
-----BEGIN CERTIFICATE-----
MIICjTCCAfagAwIBAgIDWnWvMA0GCSqGSIb3DQEBAUAMIGHMQswCQYDVQQGEwJaQTEiMCAGA1UEC
BMZRk9SIFRFU1RJTkcgUFVSUE9TRVMgT05MWTEdMBSGA1UEChMUVGhhd3RlIENlcnRpZmlyYXRpb2
4xZzAVBgNVBAsTDlRFU1QgVEVTVCBURVNUMRwwGgYDVQQDEwNUaGF3dGUgVGZzdCBDQSBSb290MBA4
XDTAyMDkyNzIzMzI0MVoXDTAyMTAxODIzMzI0MVoWDELMAkGA1UEBhMCMVxDTALBgNVBAgTBGJs
YWgxDTALBgNVBAcTBGJsYWgxDTALBgNVBAoTBGJsYWgxDTALBgNVBAsTBGJsYWgxHTAbBgNVBAMTF
HppZXIuc3YuYXJjc2lnaHQyY29tMIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCZRGNVfQwG1b
+BgABd/p8UhsaNov5AjaagAoBmouJCwgW2vwN4JViC
CSBkDpiqVF7K11Sx4ZVSXX4+VQ6k4gT5G0kDNvQeN05wWkzEMygmB+ZBnYqPA/XtWRZtjxvH
MoqS+JEqHruiMLITC6q0reUB/txby6+S9zNo/fUG1pkIcQIDAQABoyUwIzATBgNVHSUEDDAKBggrB
gEFBQcDATAMBgNVHRMBAg8EAjAAMA0GCSqGSIb3DQEBAUAA4GBAFY37E60+P4b3zTLnaG7EVM57G
tkED6PwCIi1B6ixjvNL4MNGRubPa8kyaZp5fEDoNUPQVQxnpABjzTa1RfYgjNFJ61tI6ZKjB05kim
9UBeCnKiNNzhIyDyFwbHXOPB/JaLIV+jGugYNS7hf/ay0BXXK1fue007EgjhhB/mQFs2JB
-----END CERTIFICATE-----
```

Before proceeding, make sure the name of the issuer that signed your certificate exists as a Trusted CA in cacerts.

Follow these steps to import the signed certificate:

1. If the returned file has the .CER or .CRT file extension, save it to the <ARCSIGHT_HOME>/config/jetty directory and skip to Step 4.
2. If it has a different extension, use a text editor to copy and paste the text string to a file. Include the lines "-----BEGIN CERTIFICATE-----" and "-----END CERTIFICATE-----", and make sure there are no extra spaces before or after the string.
3. Make a backup of the existing keystore by renaming it:
Rename <ARCSIGHT_HOME>/config/jetty/keystore to <ARCSIGHT_HOME>/config/jetty/keystore.old.
If, for any reason, the new keystore does not work properly, you can revert back to the demo keystore you saved as keystore.old.
4. Save it to a file named ca_reply.txt on the Manager in the <ARCSIGHT_HOME>/config/jetty directory.
5. On the Manager machine, run this command in <ARCSIGHT_HOME>/bin:
bin/arcsight keytool -store managerkeys -importcert -alias mykey -file config/jetty/ca_reply.txt
or use keytoolgui:
./arcsight keytoolgui
6. Click **File->Open keystore** and select the keystore (**keystore.request**) you saved in

Step 7 of "[Create a Key Pair for a CA-Signed Certificate](#)" on page 113. Provide the password you used to save the keystore in that step.

7. Right-click the key pair you created at the beginning of the process and named *mykey* in Step 6 of "[Create a Key Pair for a CA-Signed Certificate](#)" on page 113.
8. Select **Import CA Reply** from the menu.
9. Select the CA reply certificate file you saved in <ARCSIGHT_HOME>/config/jetty and click **Import**.

If the CA reply file contains a chain of certificates, `keytool` tries to match the reply's root CA to an existing Trusted Certificate in your cacerts truststore. If this operation fails, the Certificate Details dialog appears for manual verification. Acknowledge the certificate by clicking **OK** and answering **Yes** to the subsequent challenge. Answer **No** if the certificate is not trustworthy for some reason.

After the key pair you generated has been updated to reflect the content of the CA reply, the keystore named `keystore.request` contains both the private key and the signed certificate (in the alias *mykey*).

10. Select **File > Save**. The keystore is now ready for use by the Manager.
11. Copy <ARCSIGHT_HOME>/config/jetty/keystore.request to <ARCSIGHT_HOME>/config/jetty/keystore.
12. For successful reconfiguration and Manager startup, enter the keystore passwords into the appropriate properties file.

Enter the password into the `esm.properties` file for the Manager using the following command (all on one line):

```
arcsight changepassword  
-f <ARCSIGHT_HOME>/config/esm.properties  
-p server.privatekey.password
```

After entering this command, the system displays the previous password as asterisks and asks you to enter and then confirm your new password. These commands enter the password into the properties file in an encrypted format.

13. If your Manager clients trust the CA that signed your server certificate, go to "[Start the Manager Again \(Restart the Manager\)](#)" on the next page.

Otherwise, perform these steps to update the client's cacerts (truststore):

Note: Also perform these steps on the Manager to update the Manager's cacerts so that Manager clients such as the `archive` command can work.

- a. Obtain a root certificate from the CA that signed your server certificate and copy it to your client machine. (you got this in "[Import the CA Root Certificate](#)" on page 115.)
- b. For one client, use `keytoolgui` to import the certificate into the truststore (cacerts):

- i. In `<ARCSIGHT_HOME>/bin`, run this command:

```
./arcsight keytoolgui
```
 - ii. Click **File->Open keystore**.
 - iii. Select the store named `cacerts`. Use the default password to open `cacerts`.
 - iv. Click **Tools->Import Trusted Certificate** and select the certificate you copied earlier in this procedure.
 - v. You see the message:
Could not establish a trust path for the certificate. The certificate information will now be displayed after which you may confirm whether or not you trust the certificate.
Click **OK**.
 - vi. Enter an alias for the Trusted Certificate you just imported and click **OK**.
 - vii. Right-click the alias `ca` in the truststore and choose **Delete** from the menu.
 - viii. Save the keystore.
- c. Copy the `<ARCSIGHT_HOME>/jre/lib/security/cacerts` file from the client in the previous step to all other clients.
14. Import the new certificate into the client truststore on the manager. This is necessary so that manager utilities will continue to work.

Delete the existing manager certificate from the manager's client truststore. To delete a certificate from the truststore, start `keytoolgui` and navigate to the certificate, right-click on the certificate, and select **Delete**.

For `bin/arcsight keytool`:

```
bin/arcsight keytool -store clientcerts -delete -alias <hostname>
```

Then add the new certificate by exporting it and importing it. See ["Export a Certificate" on page 100](#), subtopic "Exporting a Certificate using `keytoolgui`", and ["Import a Certificate" on page 101](#), subtopic "Importing a Certificate Using `keytoolgui`".

Or, the commands for `bin/arcsight keytool`:

```
bin/arcsight keytool -store managerkeys -exportcert -alias mykey -file mykey.cer
```

```
bin/arcsight keytool -store clientcerts -importcert -alias <hostname> -file mykey.cer
```

Start the Manager Again (Restart the Manager)

When you start the Manager again, clients cannot communicate with it until their keystores are populated with the new certificate.

1. Start the Manager.

First, stop the Manager:

```
/etc/init.d/arcsight_services stop manager
```

Then start all services:

```
/etc/init.d/arcsight_services start all
```

The Manager may fail to start if the password of the Key pair does not match the password of the keystore, which is encrypted in `esm.properties`. If you do not remember the keystore password, run the Manager setup wizard and change the password of your existing keystore.

2. Restart all clients.

3. To verify that the new certificate is in use:

- a. From the command line navigate to `<ARCSIGHT_HOME>` and enter the command:

```
arcsight tempca -i
```

The output shows which CA issuer signed the SSL CA-signed certificate, certificate type, status of a validation of the certificate, and so on.
- b. Point a web browser to `https://<manager_hostname>:8443`. to test it.

Using CA-Signed Certificates with Additional Components

Perform these extra steps to use CA-signed certificates with additional ESM components such as the ArcSight Console, or SmartConnectors.

- Adding additional Managers

You do not need to add the CA root certificate to the Truststore-cacerts file again. Just copy the cacerts file from the existing Manager to the new Manager.

- Other ArcSight Components (Console and SmartConnectors).

When installing a new Console, copy the cacerts file from an existing Console to the new Console.

Removing a Demo Certificate

You can remove the demo certificate by using the `tempca` script located in `<ARCSIGHT_HOME>/bin`. Issue the following command on all Manager and Console installations:

```
./arcsight tempca -rc
```

For SmartConnectors, run the `tempca` script using the following command:

```
./arcsight agent tempca -rc
```

Replacing an Expired Certificate

When a certificate in your truststore/cacerts expires, replace it with a new one as follows.

To delete an expired certificate, the `keytool` command is preferred. Using the `keytoolgui` interface requires that the X Window system be installed on your system. Note that using the X Window system is not preferred, but if you have it installed and want to use it, you can use `keytoolgui`. The X Window system is not present on ESM on an appliance. To replace an expired certificate, you must delete the current certificate and import a new one. See ["Delete a Certificate" on page 107](#) and ["Import a Certificate" on page 101](#) for details on using `keytool`.

1. Delete the expired certificate from the truststore/cacerts.

To delete a certificate from the truststore/cacerts, start `keytoolgui` and navigate to the certificate, right-click on the certificate, and select **Delete**.

2. Replace the certificate by importing the new certificate into truststore/cacerts. Use `keytoolgui` to import the new certificate into the truststore/cacerts. See ["Using a Self-Signed Certificate" on page 108](#), or ["Using a CA-Signed SSL Certificate" on page 112](#) section (depending on the type of certificate you are importing) for steps on how to import the certificate.

Since the common name (CN) for the new certificate is the same as the old certificate, you cannot have both of them in the truststore, cacerts.

Establishing SSL Client Authentication

This section describes the required steps for enabling client-authentication for ArcSight Console.

All communications between ESM and Console are performed over SSL connections. Which protocols and cipher suites to use for SSL connection is decided in the very beginning, during the initial SSL handshake. SSL handshake always validates that server could be trusted by reviewing and challenging its certificate. Optionally SSL handshake could validate client's certificate to ensure that connection was requested from a legitimate client. For that purpose the client provides SSL certificate and SSL handshake verifies that the client owns the corresponding private key.

Depending on the selected authentication mode the described below configuration steps might have effect on overall user authentication. These are the implications of the various modes:

1. **Password Based Authentication:** No impact
2. **Password Based and SSL Client Based Authentication:** In this mode, the client sends the SSL certificate and password-based credentials. Both of them should identify exactly the same user.
3. **Password Based or SSL Client Based Authentication:** In this mode, the result depends on your choice. For this authentication mode Console's login dialog provides two buttons: "Login" and "SSL Client Login" to send either the username and password or the SSL certificate.
4. **SSL Client Only Authentication:** In this mode, authentication is performed based on SSL certificate only.
5. **OSP Client Only Authentication:** Select this mode if ESM will use OSP Client Only Authentication or External SAML2 Client Only Authentication.

OSP client only authentication allows ESM to use an existing One SSO Provider (OSP) (for example, from ArcSight Fusion) for authentication. For more information, see [OSP Client Only Authentication](#).

External SAML2 client only authentication configures the SAML2 client that is embedded in ESM to establish a trust relationship with your own external identity provider. For more information, see [External SAML2 Client Only Authentication](#).

Note: The OSP authentication method is not supported if ESM is running in FIPS mode.

Unless it's PKCS#11 login in the modes 2 and 4 described above with configured client-side authentication, SSL Login will always be performed under the same user, because the login dialog will always use the same client certificate.

For PKCS#11 logins the authentication process uses the certificates from PKCS#11 token, so the result will depend on the provided token.

Regardless of PKCS#11 mode, SSL login authentication is performed on server-side in two steps by validating SSL certificate and then by looking up the ArcSight user with the external ID that matches CN (Common Name) from the provided certificate.

Note: Client-side authentication could be helpful when you want to establish connection from a client to ESM always under the same user account. That eliminates the need to provide username/password. If it's what you need use the following instructions and once the client certificate is created, select "SSL Client Only Authentication" mode for that client, and create ArcSight User (in ESM) with externalID matching CN from client certificate. Do not forget to secure access to this certificate. If keystore with the certificate is stolen, it could be used to access ESM from other clients.

Setting up SSL Client-Side Authentication on ArcSight Console- Self-Signed Certificate

To enable client-side authentication for ArcSight Console running in default mode, perform these steps in addition to the ones you perform for setting up server authentication:

1. Set the External ID of the ArcSight user to the Common Name (CN) of the certificate that you will create when you generate a new key pair in a subsequent step. It is easiest to set the External ID to the user name. Note the External ID because you will need it shortly:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store clientkeys -genkeypair -dname "cn=admin, ou=ArcSight, o=MF, c=US" -keyalg rsa -keysize 2048 -alias testkey -startdate -1d -validity 366
```

2. Run `consolesetup` to select the desired client authentication method. See the section "Reconfiguring the ArcSight Console" in the *ESM Installation Guide*.

3. Export the Console's certificate:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -exportcert -store clientkeys -alias testkey -file console.cer
```

4. Copy the Console's certificate to the manager machine, and import it into the Manager's truststore:

```
bin/arcsight keytool -importcert -store managercerts -alias testkey -file console.cer
```

5. Stop the Manager as user *arcsight* by running:

```
/etc/init.d/arcsight_services stop manager
```

6. From the `<ARCSIGHT_HOME>/bin` directory, run:

```
./arcsight managersetup
```

7. Change the SSL selection to the appropriate setting.

8. Start all services so the Manager can start using the self-signed certificate. Run the following commands to do so:

```
/etc/init.d/arcsight_services start all
```

Setting up SSL Client-Side Authentication on ArcSight Console- CA-Signed Certificate

To enable client-side authentication for ArcSight Console running in default mode, perform these steps in addition to the ones you perform for setting up server authentication:

1. Set the External ID of the ArcSight user to the Common Name (CN) of the certificate that you will create when you generate a new key pair in a subsequent step. It is easiest to set the External ID to the user name. Note the External ID because you will need it shortly.
2. Run `consolesetup` to select the desired client authentication method. See the section "Reconfiguring the ArcSight Console" in the *ESM Installation Guide*.
3. On each Console generate a key pair, making sure to set the Common Name (CN) to the External ID of the user that you updated above:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -store clientkeys -genkeypair -dname "cn=admin, ou=ArcSight, o=MF, c=US" -keyalg rsa -keysize 2048 -alias testkey -startdate -1d -validity 366
```
4. Create a Signing Request by following the steps in ["Send for the CA-Signed Certificate" on page 114](#) and ["Import the CA-Signed Certificate" on page 115](#).
5. Send a request to the certificate authorities. Example for `keytool` command line to create a certificate request:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -certreq -store clientkeys -alias testkey -file config/testkey.csr
```
6. Follow the steps in ["Import the CA Root Certificate" on page 115](#). Import the CA Root Certificate into both the Console's and the Manager truststore.
7. After receiving a response, enter it into the client keystore. Example for `keytool` command line:

```
<ARCSIGHT_HOME>/bin/arcsight keytool -importcert -store clientkeys -alias testkey -file /tmp/signedcert.cer
```
8. Stop the Manager as user `arcsight` by running:

```
/etc/init.d/arcsight_services stop manager
```
9. From the `<ARCSIGHT_HOME>/bin` directory, run:

```
./arcsight managersetup
```
10. Change the SSL selection to the appropriate setting.
11. Start all services so the Manager can start using the self-signed certificate. Run the following commands to do so:

```
/etc/init.d/arcsight_services start all
```

Setting Up Client-Side Authentication for ArcSight Command Center

To set up client-side authentication for the Arcsight Command Center, you must export the Console's private key into a `p12-file`, and then import that file into the browser's internal truststore.

1. Export the Console's private key:

```
bin/arcsight keytool -importkeystore -store clientkeys -destkeystore  
config/consolekey.p12 -deststoretype PKCS12 -srcalias consolekey
```

The above command creates a new file `config/consolekey.p12` with keystore of the type **PKCS12** and stores there a private key for alias **consolekey** from client's keystore file `config/keystore.client`.

2. Use keystore `config/consolekey.p12` that contains Console's private key to import the certificate into internal browser's keystore.

Setting Up Client-Side Authentication on SmartConnectors

In order to enable client-side authentication on SmartConnectors running in default mode:

1. Create a new keypair in the `config/keystore.client` of the SmartConnector. Note that when you create the keypair, a keystore is created as well, if it does not already exist.

An example of a `keytool` command line:

```
jre/bin/keytool -genkeypair -keystore config/keystore.client -storetype  
JKS -storepass password -dname "cn=John Smith, ou=ArcSight, o=MF, c=US" -  
alias testKey -validity 365
```

2. Create a client SSL configuration text file in the `user/agent` directory and name it `agent.properties` for a connector. The contents of this file (whether client or agent) should be as follows:

```
auth.null=true  
ssl.client.auth=true  
cac.login.on=false  
ssl.keystore.path=config/keystore.client  
ssl.keystore.password=<client.keystore_password>
```

Note: Make sure that this password is identical to the password that you set for `/config/keystore.client` when creating it.

3. Export the SmartConnector's (the client's) certificate:

An example of a `keytool` command line:

```
bin/arcsight agent keytool -exportcert -store clientkeys -alias testkey -  
file /tmp/agent-certificate.cer
```

4. Import the CA's certificate of the client's certificate (in case you are using CA-signed certificate) or the client's certificate itself (in case you are using a self-signed certificate) into the Manager's truststore, `/config/jetty/truststore`.

Example command for `keytool` command line:

```
bin/arcsight keytool -importcert -store managercerts -alias testkey -file /tmp/agent-certificate.cer
```

5. Stop the Manager:

```
/etc/init.d/arcsight_services stop manager
```

Then start all services:

```
/etc/init.d/arcsight_services start all
```

6. Restart the SmartConnector.

Setting Up Client-Side Authentication for Utilities on the ESM Server

Some ArcSight commands, such as `arcsight managerinventory` require you to log into the server. To support these commands, you must set up SSL client-side authentication on the console. This requires you to set up a SSL keypair for these commands in order to use them in the tasks described in the subtopics "Password Based and SSL Client Based Authentication" and "SSL Client Only Authentication" in ["Password-Based Authentication" on page 136](#).

Ensure that the client truststore contains the manager certificate. This normally happens automatically, but there are two cases in which you must copy the certificate in manually. These are for a CA signed manager certificate, and for FIPS mode installations. When SSL authentication is not required (in the tasks described in the subtopics "Password Based Authentication" or "Password Based or SSL Client Based Authentication") no additional configuration is required.

The type of authentication used by utilities is controlled by settings in `config/client.properties`. `managersetup` will put the correct values for client authentication in `client.properties` when it first creates it, but, to preserve any custom modifications you have made in that file, it will not subsequently modify that file. If you need to change your settings after initial installation, you need to remove the `client.properties` before running `managersetup`, or edit `client.properties` as shown below.

- For "Password Based Authentication" remove the two properties: `auth.null` and `ssl.client.auth` from the `client.properties` file.
- For "SSL Client Only Authentication" set the two properties: `auth.null` and `ssl.client.auth` in the `client.properties` file to **true**.

- For "Password and SSL Client Based Authentication" set the properties in the `client.properties` file: `auth.null` to **false** and `ssl.client.auth` to **true**.
- For "Password or SSL Client Based Authentication" set the properties in the `client.properties` file: `auth.null` to **false** and `ssl.client.auth` to **optional**.

To support SSL Client Authentication for ArcSight commands:

1. Set the external ID of the user that tools will use to a known value. For example, set the admin external ID to admin.

2. Stop the Manager:

```
/etc/init.d/arcsight_services stop manager
```

3. Run `managersetup` to select the appropriate value for SSL Authentication.

4. Create a keypair in the client keystore:

```
bin/arcsight keytool -store clientkeys -genkeypair -dname "cn=admin" -  
keyalg rsa -keysize 2048 -alias admin -startdate -1d -validity 366
```

5. Export the certificate

```
bin/arcsight keytool -store clientkeys -exportcert -alias admin -file  
admin.cer
```

6. Import the certificate into the `managercerts` truststore.

```
bin/arcsight keytool -store managercerts -importcert -alias admin -file  
admin.cer
```

7. Start all services:

```
/etc/init.d/arcsight_services start all
```

8. Additionally, in a FIPS installation, you must import the manager certificate into the client keystore. This is done automatically for you in default mode.

Export the certificate with this command:

```
bin/arcsight keytool -store managerkeys -exportcert -alias mykey -file  
mykey.cer
```

Then import the certificate:

```
bin/arcsight keytool -store clientcerts -importcert -alias mykey -file  
mykey.cer
```

SSL Authentication - Migrating Certificate Types

When you migrate from one certificate type to another on the Manager, update all Consoles, and SmartConnectors.

Migrating from Demo to Self-Signed

To migrate from a demo to self-signed certificate:

1. Follow the steps described in ["Using a Self-Signed Certificate" on page 108](#).
2. Follow the instructions in ["Verifying SSL Certificate Use" below](#) to ensure that a self-signed certificate is in use.

Migrating from Demo to CA-Signed

To migrate from a demo to CA-Signed certificate:

1. Follow the steps described in ["Using a CA-Signed SSL Certificate" on page 112](#).
2. Follow the instructions in ["Verifying SSL Certificate Use" below](#) to ensure that CA-signed certificate is in use.

Migrating from Self-Signed to CA-Signed

To migrate from a self-signed to CA-signed certificate:

1. Follow the steps described in ["Using a CA-Signed SSL Certificate" on page 112](#).
2. Follow the instructions in ["Verifying SSL Certificate Use" below](#) to ensure that CA-signed certificate is in use.

Verifying SSL Certificate Use

After the migration, run this command in `<ARCSIGHT_HOME>/bin` on the client to ensure the certificate type you intended is in use:

```
./arcsight tempca -i
```

In the resulting output, a sample of which is available below, do the following:

1. Review the value of the line: Demo CA trusted.
The value should be "no."
If the value is "yes," the demo certificate is still in use. Follow these steps to stop using the demo certificate:
 - a. In `<ARCSIGHT_HOME>/bin`, enter the following command to make the client stop using the currently in use demo certificate:

```
./arcsight tempca -rc
```

For SmartConnectors, run:

```
./arcsight agent tempca -rc
```

b. Restart the client.

2. Verify that the Certificate Authority that signed your certificate is listed in the output. For a self-signed certificate, the Trusted CA is the name of the machine on which you created the certificate

Sample Output for Verifying SSL Certificate Use

This is a sample output of the `arcsight tempca -i` command run from a Console's `bin` directory:

```
ArcSight TempCA starting...

SSL Client
truststore C:\arcsight\Console\current\jre\lib\security\cacerts
  Type           JKS
  Demo CA trusted no
  Trusted CA     DigiCert Assured ID Root CA
[digicertassuredidrootca]
  Trusted CA     TC TrustCenter Class 2 CA II
[trustcenterclass2caii]
.
.
.
Demo CA
  keystore       C:\arcsight\Console\current\config\keystore.tempca
Exiting...
```

Using Certificates to Authenticate Users to the Manager

Instead of using a user name and password to authenticate a user to the Manager, you can configure these systems to use a digitally-signed user certificate. This section tells you how to do that. This capability is useful in environments that make use of Public Key Infrastructure (PKI) for user authentication.

The Manager accepts login calls with empty passwords and use the Subject CN (Common Name) from the user's certificate to identify the user.

Note: Before you enable client-side authentication, make sure that you log in to the Console and create a new user or modify an existing user such that you set the user's `external_id` to the one specified in the certificate created on the Console. The external id should be set to the users name set as the CN (Common Name) setting when creating the certificate.

You must enable SSL client authentication as described in the previous section to use digitally-signed user certificates for user authentication.

To configure the Manager to use user certificates, do the following:

1. On the Console, make sure that External ID field in the User Editor for every user is set to a value that matches the CN in their user certificate.
2. Restart the system you are configuring.
3. Restart the Consoles.

When you start the Console, the user name and password fields are grayed out. Simply select the Manager to which you want to connect and click **OK** to log in.

Using the Certificate Revocation List (CRL)

ESM supports the use of CRL to revoke a CA-signed certificate which has been invalidated. The CA that issued the certificates also issues a CRL file which contains a signed list of certificates which it had previously issued that it now considers invalid. The Manager checks the client certificates against the list of certificates listed in the CRL and denies access to clients whose certificates appear in the CRL.

Be sure these conditions exist for the CRL functionality to work properly:

- Your certificates are issued and signed by a valid Certificate Authority or an authority with an ability to revoke certificates.
- The CA's certificate is present in the Manager's trust store.
In the case of client-side authentication, the Manager validates the authenticity of the client certificate using the certificate of the signing CA.
- You have a current CRL file provided by your CA.
The CA updates the CRL file periodically when subsequent certificates are invalidated.

How CRL works:

1. CRL verification is performed by the SSL handshake. When started, ESM reads the value of the property `auth.crl.dir` (default value `config/jetty/crls`) and starts monitoring for any changes to the files with the `.crl` extension in the specified folder.

2. When there are changes (for example, a new CRL file), ESM reloads the full content of that folder and updates the current set of CRLs.
3. The current set of CRLs is used by ESM each time it initializes SSL Context.

Further considerations:

Be sure that the property `auth.crl.dir` points to the appropriate folder and have that property in the corresponding properties file (for example, `console.properties` for the ArcSight Console).

When a component starts, it reads the CRL files from the specified folder and includes them into the SSL Context. The only difference between ESM and the clients is that CRL files are read once, during startup; after the ArcSight Console started you cannot change the CRL list. The ArcSight Console does not monitor changes in the CRL folder. For example, to add an additional CRL to the ArcSight Console, you need to restart the ArcSight Console after copying the CRL file to the designated folder.

To use the CRL functionality:

1. For components other than ESM, log out of those components.
2. Copy the CA-provided CRL file into the folder specified in the property `auth.crl.dir` of that component (the property is set in the corresponding properties file).
3. Restart the components (Console, or ESM utilities) so that the current set of provided CRL files is read. For ESM, restart is not required. After adding the CRL file, it takes about a minute for the Manager to be updated with the current CRL files.

Chapter 5: Running the Manager Configuration Wizard

After you have installed and configured your system, you can change some configuration parameters by running the `managersetup -i console` command in a terminal window to launch the Manager Configuration Wizard. Running the command in console mode is the preferred way of launching the wizard. Using the X Window system to run the wizard in graphical user interface mode is not preferred, but if you have the X Window system installed and want to use it, you can run the `managersetup` command without options to launch the wizard. The X Window system is not present on an appliance.

If issues occur while running the Manager Configuration Wizard, this command logs troubleshooting information in a log file:

```
/opt/arcsight/var/logs/manager/default/serverwizard.log.
```

Important: When you run `managersetup`, you will see:

Do you want to run ESM in Compact or Distributed mode?

In this case, accept the default and keep moving through the wizard. You cannot change the ESM mode in the wizard after the initial installation of ESM. If you want to change ESM from compact to distributed mode, see the topic "Converting Compact Mode to Distributed Correlation Mode" in the *ESM Upgrade Guide* and follow that conversion process. Note that conversion from distributed correlation mode to compact mode is not supported.

Running the Wizard

Run the wizard as user *arcsight*. Before you run the Manager Configuration Wizard, stop your Manager by running the following command:

```
/etc/init.d/arcsight_services stop manager
```

Verify that the Manager has stopped by running the following command (as user *arcsight*):

```
/etc/init.d/arcsight_services status all
```

To start the wizard, run the following from `<ARCSIGHT_HOME>/bin` directory:

```
./arcsight managersetup -i console
```

The Manager Configuration Wizard establishes parameters required for the Manager to start up when you reboot.

1. Select either **Run manager in default mode** or **Run manager in FIPS mode**. For information on FIPS, see "[Configuration Changes Related to FIPS](#) " on page 230
2. You can enter **Manager Host Name**, **Manager Port**, and **Physical Location**. To change the hostname or IP address for your Manager host, enter the new one. The Manager host name that you enter appears on the Manager certificate. If you change the host name, be sure to regenerate the Manager's certificate in by selecting **Replace with the new Self-Signed key pair** in the screen that allows you to select key pair options (make a note of this if you change your host name). We recommend that you do not change the Manager Port number.

The **IP Version** selection (IPv4 or IPv6) appears if you have a dual-stack machine, such as an appliance. If you see this option, your selection has the following effects:

- It controls what IP Address is used by third party software if a hostname is given. for example, the e-mail server in Manager Setup.
 - It sets the preferred IP version to choose if there are multiple IP addresses available for the different IP versions.
 - It controls which IP Address is tried on the peering page if a hostname is specified.
 - It controls whether an IPv4 or IPv6 Address is chosen for the manager asset.
3. If you would like to replace your license file with a new one, select **Replace current license file**. Otherwise, accept the default option of **Keep the current license file**.
If you selected **Replace the current license file**. you are prompted for the new one.
 4. Select the Java Heap memory size. The Java Heap memory size is the amount of memory that ESM allocates for its heap. (Besides the heap memory, the Manager also uses some additional system memory.)
 5. Select a key pair option. The Manager controls SSL certificate type for communications with the Console, so the wizard prompts you to select the type of SSL certificate that the Manager is using. If you changed the Manager host name in the first or second step above, select **Replace with new Self-Signed key pair**, otherwise select **Do not change anything**.
If you selected **Replace with new Self-Signed key pair**, you are prompted to enter the password for the SSL key store and then details about the new SSL certificate to be issued.
 6. Accept the **Logger JDBC URL** and **Database Password** defaults.
 7. If Transformation Hub is part of your ESM implementation, select whether to set up the connection to it.
If you select to set up a connection, provide the following information:
 - a. Specify the host name and port information for the nodes in Transformation Hub. Include the host and port information for all worker nodes. Use a comma-

separated list (for example: <host>:<port>,<host>:<port>).

Note: You must specify the host name and *not* the IP address.

Transformation Hub can only accept IPv4 connections from ESM.

If the Kafka cluster is configured to use SASL/PLAIN authentication, ensure that you specify the port configured in the cluster for the SASL_SSL listener.

- b. Specify the topics in Transformation Hub from which you want to read. These topics determine the data source.

For more information, see the [Transformation Hub Administration Guide](#).

Note: You can specify up to 25 topics using a comma-separated list (for example: topic1,topic2).

- c. To import the Transformation Hub root certificate to ESM's client truststore, complete the following:

- i. On the Transformation Hub master node, run the following command to generate the `ca.crt` root certificate in the `/tmp` folder:

```
/opt/arcsight/kubernetes/scripts/cdf-updateRE.sh > /tmp/ca.crt
```

- ii. Copy the Transformation Hub root certificate to a local folder on the ESM server.

- iii. In the `managersetup` wizard, provide the path to the certificate.

The wizard imports the Transformation Hub root certificate into ESM's client truststore.

Note: If you need to manually import the certificate in FIPS mode, use the following command on the ESM server:

```
/opt/arcsight/manager/bin/arcsight keytool -store clientcerts -importcert -file <path to ca.crt>/ca.crt -alias <aliasname>
```

- d. If the Kafka cluster is *not* configured to use SASL/PLAIN authentication, leave the authentication type as None. If the Kafka cluster is configured to use SASL/PLAIN authentication, select SASL/PLAIN as the authentication type.
- e. If you selected SASL/PLAIN as the client authentication type, specify the user name and password for authenticating to Kafka.

The wizard validates the connection to Transformation Hub. If there are any issues, you will receive an error or warning message. If the wizard does not generate error or warning messages and you are able to advance to the next screen, the connection is valid.

8. Select whether to set up integration with ArcSight Investigate. Select **Yes** to enable the integration; select **No** to continue. If you select **Yes**, specify the **Search URL** for the ArcSight Investigate implementation.
9. Select the desired authentication method as described in "[Establishing SSL Client Authentication](#)" on page 120.
10. Select the method for authenticating the users. See "[Authentication Details](#)" on the [next page](#) for more details on each of these options.
11. Accept the default (**Internal SMTP server**) or configure a different email server for notification.

Caution: You must set up notification and specify notification recipients in order to receive system warnings. The importance of this step is sometimes overlooked, leading to preventable system failures.

If you choose External SMTP Server, additional options are requested, to which the following steps apply:

- a. Enter the name of the outbound **SMTP Server** to use for notifications.
- b. Enter the **From Address** that the Manager is to place in the From field of outgoing emails.
- c. Enter the **Error Notification Recipients** as a comma-separated list of email addresses to which the Manager should send error notifications.

Emails are sent when the system detects the following occurrences:

- The subsystem status is changed. The email shows the change and who did it.
 - The report has been successfully archived.
 - The account password has been reset.
 - The Archive report generation fails.
 - There is too many notifications received by a destination.
 - The event archive location has reached the cap space. It will ask you to free up some space by moving the event archives to some other place.
 - The user elects to email the ArcSight Console settings.
 - The user sends partition archival command.
 - An archive fails because there is not enough space.
 - The Connection to the database failed.
- d. Select **Use my server for notification acknowledgements**.
 - e. Enter the SMTP server and account information. This includes the incoming email server and the server protocol, and the username and password for the email account to be used.

12. The Manager can automatically create an asset when it receives an event with a new sensor or device information. The default, **Enable Sensor Asset Creation**, ensures that assets are automatically created. If you want to disable this feature, select **Disable Sensor Asset Creation**.
13. You have completed the Manager setup program. Start all services by running the following as user *arcsight*:

```
/etc/init.d/arcsight_services start all
```

Authentication Details

The authentication options enable you to select the type of authentication to use when logging into the Manager.

Caution: In order to use PKCS#11 authentication, note that PKCS#11 authentication is not supported with Radius, LDAP and Active Directory authentication methods.

By default, the system uses its own, built-in authentication, but you can specify third party, external authentication mechanisms, such as RADIUS Authentication, Microsoft Active Directory, or LDAP.

How External Authentication Works

The Manager uses the external authentication mechanism for authentication only, and not for authorization or access control. That is, the external authenticator only validates the information that users enter when they connect to the Manager by doing these checks:

- The password entered for a user name is valid.
- If groups are applicable to the mechanism in use, the user name is present in the groups that are allowed to access ArcSight Manager.

Users who pass these checks are authenticated.

After you select an external authentication mechanism, all user accounts, including the *admin* account, are authenticated through it.

Guidelines for Setting Up External Authentication

Follow these guidelines when setting up an external authentication mechanism:

- Users connecting to the Manager must exist on the Manager.
- User accounts, including *admin*, must map to accounts on the external authenticator. If the accounts do not map literally, you must configure internal to external ID mappings in the Manager.
- Users do not need to be configured in groups on the Manager even if they are configured in groups on the external authenticator.
- If user groups are configured on the Manager, they do not need to map to the group structure configured on the external authenticator.
- Information entered to set up external authentication is *not* case sensitive.
- To restrict information users can access, set up Access Control Lists (ACLs) on the Manager.

Password-Based Authentication

Password-based authentication requires users to enter their User ID and Password when logging in. You can select the built-in authentication or external authentication.

Built-In Authentication

This is the default authentication when you do not specify a third party external authentication method.

If you selected this option, you are done.

Setting up RADIUS Authentication

To configure ArcSight Manager for RADIUS Authentication, choose **RADIUS Authentication** and supply the following parameter values:

Parameter	Description
Authentication Protocol	Which authentication protocol is configured on your RADIUS server: PAP, CHAP, MSCHAP, or MSCHAP2.
RADIUS Server Host	Host name of the RADIUS server. To specify multiple RADIUS servers for failover, enter comma-separated names of those servers in this field. For example, server1, server2, server3. If server1 is unavailable, server2 is contacted, and if server2 is also unavailable, server3 is contacted.

Parameter	Description
RADIUS Server Type	Type of RADIUS server: <ul style="list-style-type: none"> • RSA Authentication Manager • Generic RADIUS Server • Safeword PremierAccess
RADIUS Server Port	Specify the port on which the RADIUS server is running. The default is 1812.
RADIUS Shared Secret	Specify the RADIUS shared secret string used to verify the authenticity and integrity of the messages exchanged between the Manager and the RADIUS server.

Setting up Active Directory User Authentication

To authenticate users using a Microsoft Active Directory authentication server, choose **Microsoft Active Directory**. Communication with the Active Directory server uses LDAP and optionally SSL.

The next panel prompts you for this information.

Parameter	Description
Active Directory Server	Host name of the Active Directory Server.
Enable SSL	Whether the Active Directory Server is using SSL. The default is True (SSL enabled on the AD server). No further SSL configuration is required for the AD server. Whether you selected SSL earlier for communications with the Console is irrelevant. Certificate type is set on the AD server side, not the manager.
Active Directory Port	Specify the port to use for the Active Directory Server. If the AD server is using SSL (Enable SSL=true), use port 636. If SSL is not enabled on the AD server, use port 389.
Search Base	Search base of the Active Directory domain; for example, DC=company, DC=com.
User DN	Distinguished Name (DN) of an existing, valid user with read access to the Active Directory. For example, CN=John Doe, CN=Users, DC=company, DC=com. The CN of the user is the "Full Name," not the user name.
Password	Domain password of the user specified earlier.
Allowed User Groups	Comma-separated list of Active Directory group names. Only users belonging to the groups listed here will be allowed to log in. You can enter group names with spaces.

Specify any user who exists in AD to test the server connection.

Specify the user name used to log in to the Manager and the External ID name to which it is mapped on the AD server.

Configuring AD SSL

If you are using SSL between the Manager and your authentication server, you must ensure that the server's certificate is trusted in the Manager's trust store <ARCSIGHT_HOME>/jre/lib/security/cacerts, whether the authentication server is using self-signed or CA certificates. For CA certificates, if the Certificate Authority (CA) that signed your server's certificate is already listed in cacerts, you do not need to do anything. Otherwise, obtain a root certificate from the CA and import it in your Manager's cacerts using the `keytoolgui` command.

Setting up LDAP Authentication

The ArcSight Manager binds with an LDAP server using a simple bind. To authenticate users using an LDAP authentication server, choose **Simple LDAP Bind** and click **Next**. The next panel prompts you for this information.

Parameter	Description
LDAP Server Host	Specify the host name of the LDAP Server.
Enable SSL	Whether the LDAP Server is using SSL. The default is True (SSL enabled on the LDAP server). No further SSL configuration is required for the LDAP server. Whether you selected SSL earlier for communications with the Console is irrelevant. Certificate type is set on the LDAP server side, not the manager.
LDAP Server Port	Specify the port to use for the LDAP Server. If the LDAP server is using SSL (Enable SSL=true), use port 636. If SSL is not enabled on the LDAP server, use port 389.

Specify any user who exists in LDAP to test the server connection.

Enter a valid Distinguished Name (DN) of a user (and that user's password) that exists on the LDAP server; for example, CN=John Doe, OU= Engineering, O=YourCompany. This information is used to establish a connection to the LDAP server to test the validity of the information you entered in the previous panel.

Note: LDAP groups are not supported. Therefore, you cannot allow or restrict logging into the Manager based on LDAP groups.

If you configure your Manager to use LDAP authentication, ensure that you create users on the Manager with their Distinguished Name (DN) information in the external ID field. For example, CN=John Doe, OU= Engineering, O=YourCompany.

Specify the user name used to log in to the Manager and the External ID name to which it is mapped on the LDAP server.

Configuring LDAP SSL

If you are using SSL between the Manager and your authentication server, you must ensure that the server's certificate is trusted in the Manager's trust store <ARCSIGHT_HOME>/jre/lib/security/cacerts, whether the authentication server is using self-signed or CA certificates. For CA certificates, if the Certificate Authority (CA) that signed your server's certificate is already listed in cacerts, you do not need to do anything. Otherwise, obtain a root certificate from the CA and import it in your Manager's cacerts using the `keytoolgui` command.

Password Based and SSL Client Based Authentication

Your authentication will be based both upon the username and password combination as well as the authentication of the client certificate by the Manager.

Note: Using PKCS#11 provider as your SSL Client Based authentication method within this option is not currently supported.

Password Based or SSL Client Based Authentication

You can either use the username/password combination or the authentication of the client certificate by the Manager (for example PKCS#11 token) to login if you select this option.

For more detail on SSL authentication for browser logins, see "Login in with SSL Authentication" in the chapter "Starting the Command Center" in the ArcSight Command Center Guide.

SSL Client Only Authentication

You must manually set up the authentication of the client certificate by the Manager. You can either use a PKCS#11 Token or a client keystore to authenticate.

One SSO Provider (OSP) Authentication

This section describes the OSP authentication methods that are available with ESM.

OSP Client Only Authentication

OSP client only authentication allows ESM to use an existing One SSO Provider (OSP) (for example, from ArcSight Fusion) for authentication.

Note: This authentication method is not supported if ESM is running in FIPS mode.

With this authentication method, you need to provide the following information:

- Location and port of the OSP server
- Host name and port of the Arcsight Command Center
Typically, this is the same as the ArcSight Manager host name and manager port that was previously provided.
- Tenant name that you specified for the OSP

With this authentication method, when a user logs in to the ArcSight Console, the **OSP Client Login** button is automatically enabled. The user is redirected to the OSP to authenticate to the console. When a user logs in to ArcSight Command Center, they are automatically redirected to the OSP to authenticate to the console. When authentication is complete, the consoles use the provided user information to communicate with the ArcSight Manager.

Note: Micro Focus recommends running only one console instance per workstation at a time. Otherwise, authentication will not work as expected. For example:

- If a console instance is running and you have authenticated to the console, if you open a second console instance the same user that is logged in to the first instance is logged in to the second instance.
- If you open two instances of the console before you authenticate, two browser sessions are opened for OSP login and the login attempt will fail. You must close all consoles and browser windows before you attempt to log in again.

External SAML2 Client Only Authentication

External SAML2 client only authentication configures the SAML2 client that is embedded in ESM to establish a trust relationship with your own external identity provider.

Note: This authentication method is not supported if ESM is running in FIPS mode.

With this authentication method, you need to provide the following information:

- Host name and port of the ArcSight Command Center
Typically, this is the same as the ArcSight Manager host name and manager port that was previously provided.
- Either the SAML2 metadata URL or the location of the SAML2 metadata file to be uploaded
- Location of the certificate that the external identity provider uses for signing SAML2 requests

Note: The certificate must be in PEM format.

With this authentication method, when a user logs in to the ArcSight Console, the **OSP Client Login** button is automatically enabled. Upon clicking the button, the user is redirected to the OSP to authenticate to the console. The OSP is configured to require authentication with a third-party SAML2 identity provider, so the user is redirected to authenticate with the external provider. Your external identity provider validates whether the user is authenticated and, optionally, redirects the user to your configured second-factor authentication, if required. Once authentication is established the user is redirected back to the OSP with a SAML2 assertion, and the OSP returns the authentication information to the console for login.

When a user logs in to ArcSight Command Center, ESM validates whether the user is authenticated. If not, the user is redirected to the OSP to authenticate to the console. The OSP is configured to require authentication with a third-party SAML2 identity provider, so the user is redirected to authenticate with the external provider. Your external identity provider validates whether the user is authenticated and, optionally, redirects the user to your configured second-factor authentication, if required. Once authentication is established the user is redirected back to the OSP with a SAML2 assertion, and the OSP returns the authentication information to the console for login.

When authentication is complete, the consoles use the provided user information to communicate with the ArcSight Manager.

Note: Micro Focus recommends running only one console instance per workstation at a time. Otherwise, authentication will not work as expected. For example:

- If a console instance is running and you have authenticated to the console, if you open a second console instance the same user that is logged in to the first instance is logged in to the second instance.
- If you open two instances of the console before you authenticate, two browser

sessions are opened for OSP login and the login attempt will fail. You must close all consoles and browser windows before you attempt to log in again.

Appendix A: Administrative Commands

ArcSight_Services Command - Compact Mode

The `arcsight_services` command syntax and options are described below.

Note: Do not start or stop services that are listed in the category Background Component Services. They are listed for information only.

Description	This command manages component services.
Applies to	All components
Syntax	<code>/etc/init.d/arcsight_services <action> <component></code>

Service Actions	start	Start the specified component, and any components it depends on.	
	stop	Stop the specified component and any components that depend on it.	
	status	This provides a service status value:	
		Status Value	Description
		initializing	Preparing to provide service.
		available_and_initializing	Providing some service while coming up.
available		Providing service	
available_unresponsive		Service is running, but is not responding to the service request.	
stopping		Service is in the process of exiting.	
unavailable		Service is not running.	
unknown	Status of service could not be determined.		
mixed_status	Service is partially up.		
help	Provides command usage (no component). The information displayed depends on where you are in ESM when you run the command with the help action.		
version	Print the complete version numbers of all components.		
Component Services	all	This is the default if no component is specified.	

	logger_httpd logger_servers logger_web manager mysqld	Logger Apache httpd service Logger service Logger Web service ESM Manager MySQL database
Background Component Services (for information only)	aps postgresql execprocsvc	ArcSight Platform Services; functions in background to perform configuration tasks; you can start this service, but do not stop it unless you are stopping all services Open source database, which functions in the background; you can start this service, but do not stop it unless you are stopping all services Helper service for the Manager; actions not supported on this service
Examples	<pre>/etc/init.d/arcsight_services stop /etc/init.d/arcsight_services start /etc/init.d/arcsight_services stop manager /etc/init.d/arcsight_services start manager /etc/init.d/arcsight_services status all</pre>	

ArcSight_Services Command - Distributed Correlation Mode

The `arcsight_services` command syntax and options are described below for use in the context of distributed correlation only.

Note: Do not start or stop services that are listed in the category Background Component Services. They are listed for information only.

Description	This command manages ArcSight services processes.
Applies to	All components
Syntax	<code>/etc/init.d/arcsight_services <action> <distributed - correlation - service></code>

Service Actions	start	Start the specified component, and any components it depends on.																
	stop	Stop the specified component and any components that depend on it.																
	status	Provides a service status value: <table border="1" data-bbox="898 430 1416 1039"> <thead> <tr> <th>Status Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>initializing</td> <td>Preparing to provide service.</td> </tr> <tr> <td>available_and_initializing</td> <td>Providing some service while coming up.</td> </tr> <tr> <td>available</td> <td>Providing service</td> </tr> <tr> <td>available_unresponsive</td> <td>Service is running, but is not responding to the service request.</td> </tr> <tr> <td>stopping</td> <td>Service is in the process of exiting.</td> </tr> <tr> <td>unavailable</td> <td>Service is not running.</td> </tr> <tr> <td>cannot_determine</td> <td>Repository is down or is up but not available, so the status of cluster node service is not known.</td> </tr> </tbody> </table>	Status Value	Description	initializing	Preparing to provide service.	available_and_initializing	Providing some service while coming up.	available	Providing service	available_unresponsive	Service is running, but is not responding to the service request.	stopping	Service is in the process of exiting.	unavailable	Service is not running.	cannot_determine	Repository is down or is up but not available, so the status of cluster node service is not known.
	Status Value	Description																
	initializing	Preparing to provide service.																
	available_and_initializing	Providing some service while coming up.																
	available	Providing service																
	available_unresponsive	Service is running, but is not responding to the service request.																
stopping	Service is in the process of exiting.																	
unavailable	Service is not running.																	
cannot_determine	Repository is down or is up but not available, so the status of cluster node service is not known.																	
statusByNode	Provides status of a node in a cluster, with all of the services running on it. Use <code>service</code> to see the hostname of the node. If no node is specified, the status of all nodes (and all services grouped by node) are displayed.																	
checkSshSetup	Checks for key-based passwordless ssh; use <code>service</code> to specify the hostname. If a hostname is not specified, performs this task on all of the hosts in the cluster.																	
sshSetup	Sets up key-based passwordless ssh; use <code>service</code> to specify the hostname. If a hostname is not specified, performs this task on all of the hosts in the cluster. Prompts for ArcSight password for each host.																	
help	Provides command usage. The information displayed depends on where you are in ESM when you run the command with the <code>help</code> action.																	
version	Print the complete version numbers of all components. Provides version information grouped by node for distributed correlation.																	

Component Services	all	This is the default if no component is specified.
	logger_httpd logger_servers logger_web manager mysqld	Logger Apache httpd service Logger service Logger Web service ESM Manager MySQL database
Background Component Services (for information only)	aps	ArcSight Platform Services; functions in background to perform configuration tasks; you can start this service, but do not stop it unless you are stopping all services
	postgresql	Open source database, which functions in the background; you can start this service, but do not stop it unless you are stopping all services
	execprocsvc	Helper service for the Manager; actions not supported on this service
Distributed Correlation Component Services	correlator# aggregator# mbus_data# mbus_control# repo# dcache# The # after the service type indicates there can be more than one of these services running at a time. Specify a instance ID (for example, <i>correlator2</i>) to work with a specific service. If no instance ID is provided, the command applies to all services of the type specified.	Service performing correlation in a cluster Services performing aggregation in a cluster Manages message bus data in a cluster Manages message bus data in a cluster Repository of cluster service information Distributed cache for a cluster
Examples	<pre>/etc/init.d/arcsight_services stop /etc/init.d/arcsight_services start /etc/init.d/arcsight_services stop manager /etc/init.d/arcsight_services start manager /etc/init.d/arcsight_services status all</pre>	

ArcSight Commands

To run an ArcSight command script on a component, open a command window and switch to the <ARCSIGHT_HOME> directory. The arcsight commands run using the file (on Windows) or arcsight.sh in <ARCSIGHT_HOME>\bin. The general syntax is as follows:

```
bin/arcsight <command_name> [parameters]
```

In general, commands that accept a path, accept either a path that is absolute or relative to <ARCSIGHT_HOME>. Running the command from <ARCSIGHT_HOME> and prefixing it with `bin/` enables you to use the shell's capabilities in looking for relative paths.

Not all parameters are required. For example, username and password may be a parameter for certain commands, such as the Manager and Package commands, but the username and password are only required if the command is being run from a host that does not also host the Manager.

ACLReportGen

This command generates a report on ACLs either at the group level or at the user level. By default, the generated report is placed in the <ARCSIGHT_HOME>/ACLReports directory.

Distributed Correlation Cluster Only: You can only run this command on a persistor node.

ACLReportGen

Applies to	Manager	
Syntax	ACLReportGen [parameters]	
Parameters	Optional:	
	-config <config>	The primary configuration file (config/server.defaults.properties).
	-locale	The locale under which to run the command.
	-mode <mode>	Mode in which this tool is run to generate the ACLs report. Supported modes are <code>grouplevel</code> and <code>userlevel</code> . The default value is <code>grouplevel</code> .
	-pc <privateConfig>	The name of the override configuration file (config/server.properties).
	-h	Help
	Example	<code>arcsight ACLReportGen</code>

agent logfu

This command runs a graphical SmartConnector log file analyzer.

agent logfu

Applies to	SmartConnectors
Syntax	<code>agent logfu -a [parameters]</code>

agent logfu, continued

Parameters	-a	SmartConnector log, which is required. For other parameters, see the description of the <code>logfu</code> command for the Manager.
Example	<code>arcsight agent logfu -a</code>	

agent tempca

This command allows you to inspect and manage temporary certificates for a SmartConnector host machine.

agent tempca

Applies to	SmartConnectors	
Syntax	<code>agent tempca</code>	
Parameters		For parameters, see the description of the <code>tempca</code> command for the Manager.
Example	<code>arcsight agent tempca</code>	

agentcommand

This command allows you to send a command to SmartConnectors.

agentcommand

Applies to	SmartConnectors	
Syntax	<code>agentcommand -c (restart status terminate)</code>	
Parameters	-c	Valid parameters are <code>restart</code> , <code>status</code> , or <code>terminate</code> .
Examples	To retrieve status properties from the SmartConnector: <code>arcsight agentcommand -c status</code> To terminate the SmartConnectorprocess: <code>arcsight agentcommand -c terminate</code> To restart the SmartConnectorprocess: <code>arcsight agentcommand -c restart</code>	

agents

This command runs all installed ArcSight SmartConnector on the host as a standalone application.

agents

Applies to	SmartConnectors	
Syntax	agents	
Parameters	None	
Example	arcsight agents	

agentsvc

This command installs an ArcSight SmartConnector as a service.

agentsvc

Applies to	SmartConnectors	
Syntax	agentsvc -i -u <user>	
Parameters	-i	Install the service.
	-u <user>	Run service as specified user.
Example	arcsight agentsvc	

agentup

This command allows you to verify the current state of a SmartConnector. It returns 0 if the SmartConnector is running and accessible, and returns 1 if it is not.

agentup

Applies to	SmartConnectors	
Syntax	agentup	
Parameters	None	
Example	arcsight agentup	

aggregatorthreaddump

This command applies to ESM in distributed mode. Run this command on the node where the aggregator is configured.

The command runs a script to dump current threads from a specified aggregator instance. The threads go into
`/opt/arcsight/var/logs/aggregator<instanceid>/aggregator.std.log.`

aggregatorthreaddump

Applies to	aggregator	
Syntax	aggregatorthreaddump	
Parameters	<aggregator instanceid>	
Examples	To run: bin/arcsight aggregatorthreaddump aggregator1	

arcdt

This command allows you to run diagnostic utilities such as session wait times, and thread dumps about your system, which can help Customer Support analyze performance issues on your components.

Distributed Correlation Cluster Only: You can only run this command on a persistor node.

arcdt

Applies to	Manager	
Syntax	arcdt diagnostic_utility utility_Parameters	
Parameters	diagnostic_utility	<p>Utilities you can run are:</p> <p>runsql—Run SQL commands contained in a file that is specified as a parameter of this command.</p> <p>Required Parameter:</p> <p>-f <sqlfile> —The file containing the sql statements to be executed.</p> <p>Optional Parameters:</p> <p>-fmt <format> —The format the output should be displayed in (where relevant), choices are html or text</p> <p>-o <outputfile> —File name to save output to. ()</p> <p>-rc <row_count> —The number of rows to be shown as a result of a select. (10000)</p>

arcdt, continued

	<p>-se <sessionEnd>— if type is EndTime or mrt, value is like yyyy-MM-dd-HH-mm-ss-SSS-zzz; if type is EventId, value is a positive integer indicating the end of eventId. (2011-06-30-01-00-00-000-GMT)</p> <p>-sr <start_row> —The row number from which you want data to be shown (0)</p> <p>-ss <sessionStart> —if type is StartTime or mrt, value is like yyyy-MM-dd-HH-mm-ss-SSS-zzz; if type is EventId, value is a positive integer indicating the start of eventId. (2011-06-30-00-00-00-000-GMT)</p> <p>-t <terminator> —The character that separates SQL statements in the input file. (;)</p> <p>-type <type> —Session type for sql query: EndTime, mrt, or EventId (EndTime)</p> <p>-cmt — Flag indicating whether all inserts and updates should be committed before exiting.</p> <p>-sp — Flag specifying whether output should be saved to disk or not.</p> <p>Required Parameter:</p> <p>-sp — Flag specifying whether output should be saved to disk or not.</p> <p>Optional Parameters:</p> <p>-c <count> — The number of times we want to query the various session tables. (5)</p> <p>-f <frequency> — The time interval (in seconds) between queries to the session tables. (20)</p>
	<p>-fmt <format> — The format the output should be displayed in (where relevant), choices are: html/text (text)</p> <p>-o <outputfile> — File name to save output to. ()</p>

arcdt, continued

		<p>thread-dumps—Obtain thread dumps from the Manager. Optional parameters which can be specified</p> <p>-c <count>— The number of thread dumps to request. (3)</p> <p>-f <frequency>—The interval in SECONDS between each thread dump request. (10)</p> <p>-od <outputdir>— The output directory into which the requested thread dumps have to be placed. ()</p>
	<p>help</p> <p>help commands</p> <p>help <command></p>	<p>Use these help Parameters (no dash) to see the Parameters, a list of commands, or help for a specific command.</p>
Examples	<p>To find out the number of cases in your database:</p> <ol style="list-style-type: none"> 1. Create a file called <code>sample.txt</code> in <code><ARCSIGHT_HOME>/temp</code> on the Manager with this SQL command: <pre>select count(*) from arc_resource where resource_type=7;</pre> 2. Run this command in <code><ARCSIGHT_HOME>/bin</code>: <pre>arcsight arcdt runsql -f temp/sample.txt</pre> 	

If not done correctly, you might get no result querying the `ArcSight.events` table from `arcdt`. For example, to run SQL to query events for a specific time period, follow the steps below:

1. Create a file such as `1.sql` in `/tmp/` containing this SQL:

```
"select * from arcsight.events where arc_deviceHostName = 'host_name'
limit 2;"
```

2. Run `arcdt` and pass the created SQL file as parameter, and also specify the time period to examine.

```
./arcsight arcdt runsql -f /tmp/1.sql -type EndTime -ss <start time> -se
<end time>
```

The result will be empty if there are no events in the specified time period.

archive

This command imports or exports resources (users, rules, and so on) to or from one or more XML files. Generally, there is no need to use this command. The Packages feature in the ArcSight Console is more robust and easier to use for managing resources.

Distributed Correlation Cluster Only: You can run this command on a cluster node that is running instances of aggregators and correlators.

archive

Description		
Applies to	Manager, Console	
Syntax	archive -f <archivefile> [Parameters]	
Required Parameter	-f <archivefile>	The input (import) or the output (export) file specification. File name paths can be absolute or relative. Relative paths are relative to <ARCSIGHT_HOME>, not the current directory.
Optional Parameters	-action <action>	Possible actions include: diff, export, i18nsync, import, list, merge, sort, and upgrade. Default: export.
	-all	Export all resources in the system (not including events).
	-autorepair	Check ARL for expressions that operate directly on resource URI's.
	-base <basefile>	The basefile when creating a migration archive. The new archive file is specified with -source (the result file is specified with -f).
	-config <file>	Configuration file to use. Default: config/server.defaults.properties
	-conflict <conflictpolicy>	The policy to use for conflicts resolution. Possible policies are: default: Prompts user to resolve import conflicts. force: Conflicts are resolved by the new overwriting the old. overwrite: Merges resources, but does not perform any union of relationships.
	-exportaction <exportaction>	The action to assign to each resource object exported. Export actions are: insert: Insert the new resource if it doesn't exist (this is the default). update: Update a resource if it exists. remove: Remove a resource if it exists.

archive, continued

-format <fmt>	Specifies the format of the archive. If you specify nothing, the default is <code>default</code> . <code>default</code> : Prompts user to resolve import conflicts. <code>preferarchive</code> : if there is a conflict, it prefers the information that is coming in over what is there. <code>install</code> : Use this for the first time. <code>update</code> : Merges the archive with the existing content. <code>overwrite</code> : Overwrites any existing content.
-h	Get help for this command.
-i	(Synonym for <code>-action import</code> .)
-m <manager>	The Manager to communicate with.
-o	Overwrite any existing files.
-p <password>	Password with which to log in to the Manager.
-param <archiveparamsfile>	The source file for parameters used for archiving. Any parameters in the named file can be overridden by command line values.
-pc <configfile>	Private configuration file to override <code>-config</code> . Default: <code>config/server.properties</code>
-pkcs11	Use this option when authenticating with a PKCS#11 provider. For example, <code>arcsight archive -m <hostname> -pkcs11 -f <file path></code>
-port <port>	The port to use for Manager communication. Default: 8443
-q	Quiet: do not output progress information while archiving
-source <sourcefile>	The source file. This is used for all commands that use the <code>-f</code> to specify an output file and use a separate file as the input.
-standalone	Operate directly on the Database, not the Manager. Warning: Do not run archive in <code>-standalone</code> mode when the Manager is running; database corruption could result.
-u <username>	The user name to log in to the Manager

archive, continued

	-uri <includeURIs>	The URIs to export. No effect during import. All dependent resources are exported, as well—for example, all children of a group. Separate multiple URIs (such as "/All Filters/Geographic/West Cost ") with a space, or repeat the -uri switch
	-urichildren <includeURIchildren>	The URIs to export (there is no effect during import). All child resources of the specified resources are exported. A parent of a specified resource is only exported if the specified resource is dependent on it.
	-xrefids	Exclude reference IDs. This option determines whether to include reference IDs during export. This is intended only to keep changes to a minimum between exports. Do not use this option without a complete understanding of its implications.
	-xtype <excludeTypes>	The types to exclude during export. No effect during import. Exclude types must be valid type names, such as Group, Asset, or ActiveChannel.
	-xtyperef <excludeTypeRefs>	The types to exclude during export (there is no effect during import). This is the same as -xtype, except it also excludes all references of the given type. These must include only valid type names such as Group, Asset, and ActiveChannel.

archive, continued

	-xuri <excludeURIs>	The URIs to exclude during export. No effect during import. Resources for which all possible URIs are explicitly excluded are not exported. Resources which can still be reached by a URI that is not excluded are still exported.
	-xurichildren <excludeURIchildren>	The URIs to exclude during export (there is no effect during import). These exclusions are such that all URIs for the children objects must be included in the set before the object will be excluded. In other words, they can still be exported if they can be reached through any path that is not excluded.
Examples	<p>To import resources from an XML file (on a Unix host):</p> <pre>arcsight archive -action import -f /user/subdir/resfile.xml -u admin -m mgrName -p pwd</pre> <p>To export certain resources (the program displays available resources):</p> <pre>arcsight archive -f resfile.xml -u admin -m mgrName -p pwd</pre> <p>To export all resources to an XML file in quiet, batch mode:</p> <pre>arcsight archive -all -q -f resfile.xml -u admin -m mgrName -p password</pre> <p>To export a specific resource:</p> <pre>arcsight archive -uri "/All Filters/Geographic/West Coast" -f resfile.xml -u admin -m mgrName -p pwd</pre> <p>Manual import (program prompts for password):</p> <pre>arcsight archive -i -format preferarchive -f resfile.xml -u admin -m mgrName</pre> <p>Scheduled or batch importing:</p> <pre>arcsight archive -i -q -format preferarchive -f resfile.xml -u admin -m mgrName -p password</pre> <p>Scheduled or batch exporting:</p> <pre>arcsight archive -f resfile.xml -u admin -m mgrName -p password -uri "/All Filters/Geographic/East Coast" -uri "/All Filters/Geographic/South"</pre>	

Make sure that the archive tool client can trust the Manager's SSL certificate. See "[SSL Authentication](#)" on [page 91](#) for information on managing certificates.

From the <ARCSIGHT_HOME>/bin/directory, you can enter the command, arcsight archive -h to get help.

Archive Command Details

Note: Ordinarily, you should use the packages feature to archive and import

resources. For more information about packages and how to use them, see the "Managing Packages" topic in ArcSight Console Online Help. Also, see the `packages` command.

You can use the `archive` command line tool to import and export resources. It is useful for managing configuration information, for example, importing asset information collected from throughout your enterprise. You can also use this tool to archive resources so you can restore it after installing new versions of this system.

The `archive` command automatically creates the archive files you specify, saving resource objects in XML format. This documentation does not provide details on the structure of archive files and the XML schema used to store resource objects for re-import into the system. Generally it is easier to use packages.

This command displays a resource in the archive menu list of resources only if the user running the utility has top-level access to the resource. Access is different for each mode.

Remote Mode

In remote mode, you can import or export from either a Manager or ArcSight Console installation and can perform archive operations while the Manager is running.

```
arcsight archive -u Username -m Manager [-p Password] -f Filename  
[-i | -sort] [-q] ...
```

Caution: The `cacerts` file on the Manager host must trust the Manager's certificate. You may have to update `cacerts` if you are using demo certificates by running:

```
arcsight tempca -ac
```

You do not need to run the above command if you run the `archive` command from the Console.

When you run the archive utility in the remote mode, it runs as the user specified in the command line. However, even users with the highest privilege level (administrator) do not have top level access to, for example, the user resource (All Users). Thus, the User resource does not show up in the list of resources. You can export users with the `-uri` option, but if you want to use the `-u` option, use the Standalone mode.

To export user resources, you can use the `-uri` option and specify a user resource to which you have direct access. For example:

```
arcsight archive -u <username> -m <manager_hostname> -format exportuser -f  
exportusers.xml -uri "/All Users/Administrators/John
```

Standalone Mode

In standalone mode, from the computer where the Manager is installed, you can connect directly to the database to import or export resource information, however, the Manager must be shut down before you perform archive operations.

Caution: Do not run the archive tool in standalone mode against a database currently in use by a Manager as it is possible to corrupt the database.

The basic syntax for the `archive` command in standalone mode is the following:

```
arcsight archive -standalone -f Filename [-i | -sort] [-q] ...
```

Note: Both remote and standalone `archive` commands support the same optional arguments.

Note that the standalone mode only works from the archive command found in the Manager installation, and does not work remotely. For example:

```
arcsight archive -standalone -format exportuser -f exportusers.xml
```

Exporting Resources to an Archive

1. Make sure the archive tool client can trust the Manager's SSL certificate. Refer to ["SSL Authentication" on page 91](#) for information on managing certificates.

From the `<ARCSIGHT_HOME>/bin` directory, you can enter the command, `arcsight archive -h` to get help.

2. From the `<ARCSIGHT_HOME>/bin` directory, enter the `arcsight archive` command along with any parameters you want to specify.

This command logs into the Manager then displays a list of Resources available for archiving.

Note: If the Manager is running, you must specify archive commands in remote mode, entering your user name, password, and Manager name to connect to the Manager. To run the archive command in standalone mode, accessing resources directly from the ArcSight Database, enter `-standalone` rather than `-u <username> -p <password> -m <manager>`.

3. Enter the number of the resource type to archive.

The `archive` command displays a list of options that let you choose which resource or group within the resource type that you want to archive.

4. Choose the resource or group to archive.

After making your selection, you are prompted whether you want to add more resources to the archive.

5. You can continue adding additional resources to the archive list. When you've

finished, answer no to the prompt

Would you like to add more values to the archive? (Y/N)

After it is finished writing the archive file, you are returned to the command prompt.

Importing Resources from an Archive

1. Make sure the archive tool client can trust the Manager's SSL certificate.
2. From the <ARCSIGHT_HOME>/bin directory, type `arcsight archive` with its parameters and attach `-i` for import.

Note: If the Manager is running, you must specify archive commands in remote mode, entering your user name, password, and Manager name to connect to the Manager. To run the archive command in standalone mode, accessing resources directly from the database, enter `-standalone` rather than `-u <username> -p <password> -m <manager>`.

3. Select one of the listed options if there is a conflict.

Importing is complete when the screen displays `Import Complete`.

Syntax for Performing Common Archive Tasks

For manual importing, run this command in <ARCSIGHT_HOME>/bin:

```
arcsight archive -i -format preferarchive -f <file name>  
-u <user> -m <manager hostname>
```

Before performing the import operation, you are prompted for a password to log in to the Manager.

For exporting:

```
arcsight archive -f <file name>  
-u <user> -m <manager hostname>
```

Before performing the import operation, you are prompted for a password to log in to the Manager and use a series of text menus to pick which Resources are archived.

For scheduled/batch importing:

```
arcsight archive -i -q -format preferarchive  
-f <file name> -u <user>  
-p <password> -m <manager hostname>
```

For scheduled/batch exporting:

```
arcsight archive -u admin -p password -m arcsightserver
```



```
-f somefile.xml -uri "/All Filters/Geographic Zones/West Coast"
-uri "/All Filters/Geographic Zones/East Coast"
```

Note: You can specify multiple URI resources with the URI parameter keyword by separating each resource with a space character, or you can repeat the URI keyword with each resource entry.

archivefilter

This command changes the contents of the archive. The archivefilter command takes a source archive xml file as input, applies the filter specified and writes the output to the target file.

Distributed Correlation Cluster Only: You can run this command on a cluster node that is running instances of aggregators and correlators.

archivefilter

Applies to	Manager	
Syntax	archivefilter -source <sourcefile> -f <archivefile> [Parameters]	
Parameters	-a <action>	Action to perform can be insert, or remove}. if you specify nothing, no action is performed.
	-e <element_list>	Elements to process (Default: '*' which denotes all elements)
	-extid <regex>	Regular expression to represent all of the external IDs to include. This is the external ID of the archival object. (Default: none)
	-f <file>	Target file (required). If a file with an identical name already exists in the location where you want to create your target file, the existing file is overwritten. If you would like to receive a prompt before this file gets overwritten, use the -o option
	-o	Overwrite existing target file without prompting (Default: false)
	-relateduri <regex>	Regular expression to get all of the URIs found in references to include. This checks all attribute lists that have references and if any of them have a URI that matches any of the expressions, that object is included
	-source <file>	Source file (required)
	-uri <regex>	Regular expression to represent all of the URIs to include. This is the URI of the archival object
	-xe <element_list>	Elements to exclude
	-xextid <regex>	Regular expression to represent all of the external IDs to exclude
	-xgroup <types>	The group types to exclude.

archivefilter, continued

	-xuri <regex>	Regular expression to represent all of the URIs to exclude
	-h	Help
Examples	<p>To include any resources, for example all Active Channels, whose attributes contain the URI specified by the <code>-relateduri</code> option:</p> <pre>arcsight archivefilter -source allchannels.xml -f t0.xml -relateduri "/All Active Channels/ArcSight Administration/"</pre> <p>To include any resources whose parent URI matches the URI specified by the <code>-uri</code> option:</p> <pre>arcsight archivefilter -source allchannels.xml -f t0.xml -uri "/All Active Channels/ArcSight Administration/*.*</pre> <p>To exclude resources whose parent URI matches the URI specified by the <code>-xuri</code> option:</p> <pre>arcsight archivefilter -source allchannels.xml -f t0.xml -xuri "/All Active Channels/*.*</pre> <p>To include all the resources that contain either URIs specified by the two <code>-relateduri</code> Parameters:</p> <pre>arcsight archivefilter -source allchannelsFilter.xml -f t0.xml -relateduri "/All Active Channels/ArcSight Administration/" -relateduri .*Monitor.*</pre>	

bleep

This command is an unsupported stress test to supply a Manager with security events from replay files (see `replayfilegen`). Replay files containing more than 30,000 events require a lot of memory on the bleep host.

Do not run `bleep` on the Manager host. Install the Manager on the bleep host and cancel the configuration wizard when it asks for the Manager's host name.

Run `arcsight tempca -ac` on the bleep host if the Manager under test is using a demo certificate.

Create the file `config/bleep.properties` using the descriptions in `bleep.defaults.properties`.

bleep

Applies to	Manager	
Syntax	<code>bleep [-c <file>] [-D <key>=<value> [<key>=<value>...]]</code>	
Parameters	<code>-c file</code>	Alternate configuration file (default: <code>config/bleep.properties</code>)
	<code>-D <key>=<value></code>	Override definition of configuration properties
	<code>-m <n></code>	Maximum number of events to send. (Default: -1)

bleep, continued

	-n <host>	Manager host name
	-p <password>	Manager password
	-t <port>	Manager port (Default: 8443)
	-u <username>	Manager user name
	-h	Help
Examples	To run: arcsight bleep	

bleepsetup

This command runs a wizard to create the `bleep.properties` file.

bleepsetup

Applies to	Manager	
Syntax	bleepsetup	
Parameters	-f	Properties file (silent mode)
	-i	Mode: {swing, console, recorderui, silent} Default: swing
	-g	Generate sample properties file
Examples	To run: arcsight bleepsetup	

changepassword

This command changes obfuscated passwords in properties files. The utility prompts for the new password at the command line.

Distributed Correlation Cluster Only: You can run this command on a cluster node that is running instances of aggregators and correlators.

changepassword

Description	This command changes obfuscated passwords in properties files. The utility prompts for the new password at the command line.	
Applies to	Manager	
Syntax	changepassword -f <file> -p <property_name>	

changepassword, continued

Parameters	-f <file>	Properties file, such as config/server.properties
	-p <property_name>	Password property to change, such as server.privatekey.password
Examples	To run: arcsight changepassword	

checklist

This command is the ArcSight Environment Check. Used internally by the installer to see if you have the correct JRE and a supported operating system.

This can run from the Manager.

certadmin

This command only applies to ESM in distributed correlation mode. Run this command on the persistor node.

The command starts a tool that shows all certificates based on status: Submitted, Approved, or Revoked. Use this command to approve or revoke SSL certificates that are used by the aggregator, correlator, and persistor services. The command uses an internal certificate administrator user that you configure in the First Boot Wizard when you start a cluster.

You can run this command from the command line or use the `-i` option to run it interactively.

The following table describes the command syntax and parameters:

certadmin

Syntax	certadmin [-approve -approveall -changepassword -count -help -i -importcert <file> -init -list <status> -recoverpassword -revert -revoke] [-id <id>] [-v]	
Parameters	-approve -id <id>	Approves the certificate of the specified ID and then changes the status from Submitted to Approved
	-approveall	Approves all certificates currently in Submitted status
	-changepassword	Changes the certificate administrator password
		Note: This is different from the Manager-level changepassword command that applies to obfuscated passwords in ESM properties files.

certadmin, continued

-count	Displays the number of certificates in the repository by status For example: Submitted: 2 Approved: 0 Revoked: 0
-help	Displays all possible commands and options with short descriptions for the certadmin tool
-i	Specifies to use interactive mode to run the command
-id	ID of the entry to process
-importcert <file>	Imports an SSL certificate into the information repository (repo) and approves it The command output displays the ID and alias of the newly-approved certificate. After the certificate is imported, it is included in the list of approved certificates that the -list command returns. Note: You might need to restart the cluster in order for ESM to recognize the certificate.
-init	Initializes certificate administration Use this only when converting from compact to distributed mode, or when upgrading from a previous ESM release (assumed to be in compact mode) to ESM in distributed mode. Note: If ESM is installed in distributed mode from the start, certificate administration is automatically initialized.
-list <status>	Lists certificates by the specified status, where <status> can be submitted, approved, or revoked You can also use the -id parameter to limit the list.
-recoverpassword	Generates a Technical Support password recovery request
-revert -id <id>	Reverts any previously approved or revoked entry for the specified ID
-revoke -id <id>	Revokes one submitted entry for the specified ID
-v	Specifies to display verbose output

configure-event-forwarding

This command only applies to ESM in distributed correlation mode. Run this command on the persistor node.

The command starts a tool that allows you to validate and commit to the information repository the properties that you set for distributed event forwarding. For more information about configuring distributed event forwarding, see [Forwarding Events in Distributed Correlation Mode](#).

The following table describes the command syntax and parameters:

configure-event-forwarding

Syntax	configure-event-forwarding [-backpressure <enable disable> -commit <file> -disable -enable -h -print -validate <file>]	
Parameters	-backpressure <enable disable>	If enabled (default), internal processing lag while forwarding messages will pause ESM's ingestion of events and might reduce EPS. If disabled, ingestion is not paused and ESM might not forward some filtered events.
	-commit <file>	Commits the settings in the specified properties file <ARCSIGHT_HOME>/config/forwarding.properties is the default file value.
	-disable	Disables event forwarding
	-enable	Enables event forwarding
	-h, --help	Prints this message
	-print, --printCommitted	Prints the current committed configuration of event forwarding
	-validate <file>	Validates the settings in the specified properties file <ARCSIGHT_HOME>/config/forwarding.properties is the default file value.

console

This command runs the ArcSight Console.

console

Applies to	Console
Syntax	console [-i] [parameters]

console, continued

Parameters	-ast <file>	
	-debug	
	-i	
	-imageeditor	
	-laf <style>	Look and feel style: metal, plastic, plastic3d. The default style is plastic3d.
	-p <password>	Password
	-port	Port to connect to Manager (default: 8443)
	-redirect	
	-relogin	
	-server	Manager host name
	-slideshow	
	-timezone <tz>	Timezone: such as "GMT" or "GMT-8:00"
	-trace	Log all Manager calls
	-u <name>	User name
Examples	To run the console: ArcSight Console	

consolesetup

This command runs the ArcSight Console Configuration Wizard to reconfigure an existing installation.

consolesetup

Applies to	Console	
Syntax	consolesetup [-i <mode>] [-f <file>] [-g]	
Parameters	-i <mode>	Mode: console, silent, recorderui, swing
	-f <file>	Log file name (properties file in -i silent mode)
	-g	Generate sample properties file for -i silent mode
Examples	To change some console configuration parameters: ArcSight Consolesetup	

correlationsetup

This command applies to ESM in distributed mode. It allows you to setup correlation and aggregation services in a distributed correlation environment. This command is described in the topic "[Configuring Services in a Distributed Correlation Cluster](#)" on [page 60](#), "Configuring Correlators and Aggregators".

correlatorthreaddump

This command applies to ESM in distributed mode. Run this command on the node where the correlator is configured.

The command runs a script to dump current threads from a specified correlator instance. The threads go into

`/opt/arcsight/var/logs/correlator<instanceid>/correlator.std.log.`

correlatorthreaddump

Applies to	correlator	
Syntax	correlatorthreaddump	
Parameters	<correlatorinstanceid>	
Examples	To run: <code>bin/arcsight correlatorthreaddump correlator1</code>	

dcachesetup

This command applies to ESM in distributed mode. It allows you to setup distributed cache services in a distributed correlation environment. This command is described in the topic "[Configuring Services in a Distributed Correlation Cluster](#)" on [page 60](#), "Configuring a Distributed Cache".

deploylicense

This command is the recommended method for replacing or adding ArcSight Manager licenses. After the command successfully completes, ESM immediately applies the changes and you do not have to restart the Manager.

deploylicense

Applies to	Manager
Syntax	<code>deploylicense [options]</code>

deploylicense, continued

Parameters	<p><code>-i <mode></code></p> <p>If you do not specify this parameter, ESM will use swing mode.</p>	<ul style="list-style-type: none"> • <code>console</code> <p>Console is the preferred mode of operation and the only mode available on an appliance. In console mode, you answer configuration questions in a terminal window. Do not specify other options with console mode.</p> <p>If you receive the following error message after you run <code>deploylicense</code>, use console mode:</p> <pre>Could not initialize class sun.awt.X11GraphicsEnvironment</pre> <ul style="list-style-type: none"> • <code>swing</code> <p>Swing mode allows you to answer configuration questions in a graphical user interface. Do not specify other options with swing mode.</p> <ul style="list-style-type: none"> • <code>silent</code> <p>Specify <code>silent</code> with the <code>-f</code> option to allow ESM to read the configuration from a file that you created with the <code>recordui</code> mode or the <code>-g</code> option. Only specify the <code>-f</code> option with silent mode.</p> <ul style="list-style-type: none"> • <code>recordui</code> <p>Use this mode to answer configuration questions in a graphical user interface and record the configuration in a file for use with silent mode on other systems. Do not specify other options with <code>recordui</code> mode.</p>
	<code>-f <file></code>	Name of the file to use with silent mode
	<code>-g</code>	<p>Generates sample properties for use with silent mode</p> <p>ESM sends the sample properties to <code>stdout</code>, but you can redirect the output to a file. You can then edit the output file with your own configuration information for use with the <code>-f</code> option in silent mode.</p>
	<code>-h</code>	Displays the Help for the <code>deploylicense</code> command
Examples	<ul style="list-style-type: none"> • To run in console mode: <code>arcsight deploylicense -i console</code> • To run in swing mode: <code>arcsight deploylicense</code> • To generate a sample properties file for editing: <code>arcsight deploylicense -g > /opt/mysetup.file</code> • To configure the system from a file: <code>arcsight deploylicense -i silent -f /opt/mysetup.file</code> 	

downloadcertificate

This command runs a wizard for importing certificates from another server.

Note: This command is not supported in FIPS mode. To import or export certificates in FIPS mode, use the keytool command. For more information, see ["keytool" on page 174](#).

Distributed Correlation Cluster Only: You can run this command on a cluster node that is running instances of aggregators and correlators.

This command is primarily for downloading a certificate from one ESM server to a Console to facilitate communications between them. When you run this command it prompts you for:

- Host name (or IP address) of the server to download from
- Port number
- Path to the keystore to which to download the certificate
This is typically <ARCSIGHT_HOME>/jre/lib/security/cacerts

For FIPS mode, It needs to go to

<ARCSIGHT_HOME>/config/jetty/keystore.bcfks.

- Keystore password
- A new alias (name) for the certificate you are downloading

downloadcertificate

Applies to	Manager	
Syntax	downloadcertificate	
Parameters	-i <mode>	Mode: console, silent, recorderui, swing
	-f <file>	Log file name (properties file in -i silent mode)
	-g	Generate sample properties file for -i silent mode
Examples	To run: arcsight downloadcertificate	

exceptions

This command allows you to search for logged exceptions in ArcSight log files.

Distributed Correlation Cluster Only: You can run this command on a cluster node that is running instances of aggregators and correlators.

exceptions

Applies to	Manager, Console, SmartConnectors	
Syntax	<code>exceptions logfile_list [parameters] [path to the log file]</code> The path to the log file must be specified relative to the current working directory.	
Parameters	-x	Exclude exceptions/errors that contain the given string. Use @filename to load a list from a file.
	-i	Include exceptions/errors that contain the given string. Use @filename to load a list from a file.
	-r	Exclude errors.
	-q	Quiet mode. Does not display exceptions/errors on the screen.
	-e	Send exceptions/errors to the given email address.
	-s	Use a non-default SMTP server. Default is <code>bynari.sv.arcsight.com</code> .
	-u	Specify a mail subject line addition, that is, details in the log.
	-n	Group exceptions for readability.
	-l	Show only exceptions that have no explanation.
	-p	Suppress the explanations for the exceptions.
Example	To run: <code>arcsight exceptions /opt/home/arcsight/manager/logs/default/server.log*</code>	

export_system_tables

This command exports your database tables. On completion, the command generates two files: a temporary parameter file and the actual database dump file, which is placed in `<ARCSIGHT_HOME>/tmp`.

For best results stop the Manager before running this command.

Distributed Correlation Cluster Only: You can only run this command on a persistor node.

export_system_tables

Applies to	Manager	
Syntax	<code>export_system_tables <username> <password> <DBname> [-s]</code>	
Parameters	<username>	CORR-Engine username

export_system_tables, continued

	<password>	Password for the CORR-Engine user
	<DBname>	Name of the CORR-Engine from which you are exporting the system tables
	-s	Include session list tables
Examples	<p>To run:</p> <pre>/etc/init.d/arcsight_services stop manager</pre> <pre>arcsight export_system_tables <DB username> <password> <DBname></pre> <p>Trend resources are exported, but not trend data from running them. After you import, re-run the trends to generate new data.</p> <p>When you are done.</p> <p>Stop the Manager:</p> <pre>/etc/init.d/arcsight_services stop manager</pre> <p>Start all services:</p> <pre>/etc/init.d/arcsight_services start all</pre>	

flexagentwizard

This command generates simple ArcSight FlexConnectors.

flexagentwizard

Applies to	SmartConnectors	
Syntax	flexagentwizard	
Parameters	None	
Examples	<p>To run:</p> <pre>arcsight flexagentwizard</pre>	

groupconflictingassets

This command groups asset resources with common attribute values (the Group Conflicting Attribute Assets Tool). Assets can have conflicting IP addresses or host names within a zone.

groupconflictingassets

Applies to	Manager
Syntax	groupconflictingassets

groupconflictingassets, continued

Parameters	-c	Clean (delete the contents of) the group to receive links to assets before starting. (Default: false)
	-m <host>	Manager host name or address
	-o <name>	Name for group to receive links to assets which have conflicting attributes. (Default: "CONFLICTING ASSETS")
	-p <password>	Password
	-port <n>	Port to connect to Manager (Default: 8443)
	-prot <string>	Protocol; only use https (Default: https)
	-u <name>	User name
	-h	Help
Examples	To run: arcsight groupconflictingassets	

import_system_tables

This command imports database tables. The file you import from must be the one that export_system_tables utility created. This utility looks for the dump file you specify in <ARCSIGHT_HOME>/tmp/.

For best results stop the Manager before running this command.

import_system_tables

Applies to	Manager	
Syntax	import_system_tables <arcsight_user> <password> <DBname> <dump_file_name>	
Parameters	<arcsight_user>	The database username, as set when you ran the first-boot wizard.
	<password>	Password for the database, as set when you ran the first-boot wizard.

import_system_tables, continued

	<DBname>	This is the name of the CORR-Engine, and it is always arcsight.
	<dump_file_name>	Use arcsight_dump_system_tables.sql, which is the name the system gave this dump file when you exported it. If you specify no path, the file is located in <ARCSIGHT_HOME>/tmp/. To specify a different path, use an absolute path. Do not specify a relative path.
Examples	<pre>/etc/init.d/arcsight_services stop manager</pre> <pre>arcsight import_system_tables dbuser mxyzptlk arcsight arcsight_dump_system_tables.sql</pre> <p>Note: Trend resources are exported, but not trend data from running them. After you import, re-run the trends to generate new data.</p> <p>When you are done.</p> <p>Stop the Manager:</p> <pre>/etc/init.d/arcsight_services stop manager</pre> <p>Start all services:</p> <pre>/etc/init.d/arcsight_services start all</pre>	

keytool

This command runs the Java Runtime Environment keytool utility to manage key stores.

Distributed Correlation Cluster Only: You can run this command on a cluster node that is running instances of aggregators and correlators.

keytool

Applies to	Manager, Console, SmartConnectors
Syntax	keytool -store <name>

keytool, continued

Parameters	-store <name>	<p>(Required) The specific store can be managerkeys, managercerts, clientkeys, clientcerts, ldapkeys, or ldapcerts.</p> <p>(original parameters) All parameters supported by the JRE keytool utility are passed along. Use arcsight keytool</p>
	-help	For a list of parameters and arguments. Also, use the command <code>keytool</code> without arguments or the <code>arcsight</code> prefix for more-detailed help.
Examples	<p>To view Console key store:</p> <pre>arcsight keytool -store clientkeys</pre> <p>The parameters for this command are actually sub-commands and many of them have their own sub-commands or parameters. To see all the possible sub-commands use <code>-help</code> followed by the sub-command for which you want to see all sub sub commands or parameters.</p> <p>For example, if you have a keystore called "managercerts," you could type <code>keytool -help -store managercerts</code> to see a list of all 16 additional subcommands. Then you could run:</p> <pre>keytool -help -store managercerts -list</pre> <p>to get additional help with the sub sub-command <code>-list</code>.</p>	

keytoolgui

This command runs a graphical user interface command for manipulating key stores and certificates. It is recommended that you use `bin/arcsight keytool`.

Distributed Correlation Cluster Only: You can run this command on a cluster node that is running instances of aggregators and correlators, as well as on the persistor node.

Note: Using `keytoolgui` requires that the X Window System be installed on your system. The X Window System is not present on ESM on an appliance. Also, `keytoolgui` is not supported on FIPS.

keytoolgui

Applies to	Manager, Console	
Syntax	<code>keytoolgui</code>	
Parameters	None	
Examples	<p>To run:</p> <pre>arcsight keytoolgui</pre>	

kickbleep

This command runs a simple, standardized test using the bleep utility..

kickbleep

Applies to	Manager	
Syntax	kickbleep	
Parameters	-f	Properties file (silent mode)
	-g	Generate sample properties file
	-i	Mode: {swing, console, recorderui, silent} Default: swing
Examples	To run: <code>arcsight kickbleep</code>	

listsubjectdns

This command displays subject distinguished names (DN) from a key store.

Distributed Correlation Cluster Only: You can run this command on a cluster node that is running instances of aggregators and correlators.

listsubjectdns

Applies to	Manager, SmartConnectors	
Syntax	listsubjectdns	
Parameters	-store name	Specific store { managerkeys managercerts clientkeys clientcerts ldapkeys ldapcerts } (Default: clientkeys.)
Examples	To list Distinguished Names in the Console key store: <code>arcsight listsubjectdns</code>	

logfu

This command runs a graphical tool for analyzing log files. It generates an HTML report (`logfu.html`) and, in SmartConnector mode, includes a graphic view of time-based log data. Logfu pinpoints the time of the problem and often the cause as well.

Note: Using `logfu` requires that the X Window system be installed on your system. The X Window system is not present on ESM on an appliance.

Running logfu

1. Open a command or shell window in `<ARCSIGHT_HOME>/logs/default`. This refers to the logs directory under the ArcSight installation directory. Note that `logfu` requires an X Windows server on Unix platforms.
2. It is recommended that you increase the log size before executing the `logfu` command by running:

```
export ARCSIGHT_JVM_OPTIONS=" -Xms1024m -Xmx1024m -Djava.security.policy=$ARCSIGHT_HOME/config/agent/agent.policy"
```
3. Run `logfu` for the type of log to analyze:
For Manager logs, run: `../../bin/arcsight logfu -m`
For SmartConnector logs, run: `../../bin/arcsight agent logfu -a`
4. Right-click in the grid and select **Show Plot/Event Window** from the context menu.
5. Check at least one attribute (such as **Events Processed**) to be displayed.

Working with logfu

The initial `logfu` display is an empty grid; loading large log files can take a few minutes (a 100 MB log might take 5 or 10 minutes). After log files are scanned, their information is cached (in files named `data.*`), which makes subsequent log file loading faster. If something about the log changes, however, you must manually delete the cache files to force `logfu` to reprocess the log.

Right-click the grid and choose **Show Plot/Event Window** from the context menu. Select what to show on the grid from the **Plot/Event Window** that appears.

The tree of possible items to display is divided into **Plot** (attributes that can be plotted over time, like events per second) and **Event** (one-time events, like exceptions, which are shown as vertical lines). Check items to display.

Because SmartConnectors can talk to multiple Managers and each can be configured to use multiple threads for events, the **Plot** hierarchy includes nodes for each SmartConnector and each Manager. Within the SmartConnector, threads are named E0, E1, and so on. Each SmartConnector has one heartbeat thread (H0) as well. Different types of SmartConnector (firewall log SmartConnector, IDS SNMP SmartConnector, and so on) have different attributes to be plotted.

The interactive Chart uses sliders to change the view. Hovering over a data point displays detailed information. There are two horizontal sliders; one at the top of the grid, one underneath. The slider at the top indicates the time scale. Drag it to the right to zoom in, or widen the distance between time intervals (vertical lines). The slider at the bottom changes the interval between lines—anywhere from 1 second at the far left to 1 day at the far right. The time shown in the grid is listed below the bottom slider:

Showing YY/MM/DD HH:MM:SS - YY/MM/DD HH:MM:SS (Interval= X)

Click anywhere in the grid area and drag a green rectangle to zoom in, changing both the vertical and horizontal scales. Hold the `Ctrl` key as you drag to pan the window in

the vertical or horizontal direction, and hold both the `Shift` and `Ctrl` keys as you drag to constrain the pan to either vertical or horizontal movement. When you are panning, only sampled data is shown, but when you stop moving, the complete data fills in. You can change this by unchecking **Enable reduced data point rendering** by right-clicking and selecting **Preferences**. You can change the rendering back to the default behavior by right-clicking and selecting **Enable fast rendering**.

For each attribute being plotted, a colored, vertical slider appears on the right of the grid. This slider adjusts the vertical (value) scale of the attribute plotted.

By default, data points are connected by lines. When data is missing, these lines can be misleading. To turn off lines, uncheck **Connect dots** by right-clicking and selecting **Preferences**.

After you have specified attributes of interest, scaled the values, centered and zoomed the display to show exactly the information of concern, right-click and select **Save as JPG** to create a snapshot of the grid display that you can print or email. The size of the output image is the same as the grid window, so maximize the window to create a detailed snapshot, or reduce the window size to create a thumbnail.

To return to a previous data view, right-click and select **Bring to Front**, **Send to Back**, **Undo Zoom**, or **Zoom out**, depending on context. Use **Auto Scale** to fit data into the grid. **Go to** allows you to display the specific log file line that corresponds to a data point. **Reset** clears all checked attributes and empties the grid.

logfu Analysis Example - Peak Event Volume

In this example, a SmartConnector is sending 10 events per second (EPS) to the Manager, but is later sending 100, then 500, then 1000 EPS before dropping back down to 10. Logfu lets you plot the SmartConnector's EPS over time; in this case the result is a peak in event volume.

When you plot the Manager's receipt of these events, you might see that it keeps up with the SmartConnector until 450 EPS or so. You notice that the Manager continues consuming 450 EPS even as the SmartConnector's EPS falls off. This is because the Manager is consuming events that were automatically cached.

By plotting the estimated cache size, you can see that the SmartConnector experienced a peak event volume and the cache stepped in to make sure that the Manager did not lose events, even when it could not keep up with the SmartConnector.

Use the vertical sliders on the right to give each attribute a different scale to keep the peak EPS from the SmartConnector from obscuring the plot of the Manager's EPS.

logfu Analysis Example - SmartConnector Down

In this example, a Check Point SmartConnector that was down for almost seven days. Logfu plotted the event stream from the SmartConnector and it was flat during the seven days, pinpointing the outage as well as the time that the event flow resumed. By

overlaying Check Point Log Rotation events on the grid, it became clear that the event outage started with a Log Rotation and that event flow resumed coincident with a Log Rotation.

Further investigation revealed what had happened: the first Check Point Log Rotation failed due to lack of disk space, which shut down event flow from the device. When the disk space problem had been resolved, the customer completed the Log Rotation and event flow resumed.

If the Manager suddenly stops seeing events from a SmartConnector, `logfu` can help determine whether the SmartConnector is getting events from the device. Another common complaint is that not all events are getting through. `logfu` has a plot attribute called `ZFilter` (zone filter) that indicates how many raw device events are being filtered by the SmartConnector. Events processed (the number of events sent by the device) minus `ZFilter` should equal **Sent** (the number of events sent to the Manager).

logfu

Applies to	Manager (See also agent logfu.)	
Syntax	<code>logfu {-a -m} [parameters]</code>	
Parameters	<code>-a</code>	Analyze SmartConnector logs
	<code>-f <timestamp></code>	From time
	<code>-i</code>	Display information about the log files to be analyzed
	<code>-l <timespec></code>	Analyze only the specified time (Format: <code><time>{smhd}</code>) Examples: 1d = one day, 4h = four hours
	<code>-m</code>	Analyze Manager logs
	<code>-mempercent <n></code>	Percent of memory messages to consider for plotting. (Default: 100)
	<code>-noex</code>	Skip exception processing
	<code>-noplot</code>	Skip the plotting
	<code>-t <timestamp></code>	To time
Examples	To analyze Manager logs for the last 12 hours: <code>arcsight logfu -m -l 12h</code>	

managerinventory

This command displays configuration information about the installed Manager.

managerinventory

Applies to	Manager
Syntax	<code>managerinventory</code>

managerinventory, continued

Parameters	-a <filter>	Attribute filter. Default: ""
	-f <filter>	Object filter. Default: "Arcsight:*"
	-m <host>	Manager host name or address
	-o <op>	Operation {list, show}. Default is list
	-out <file>	Output filename. Default is stdout
	-p <password>	Password
	-port <n>	Port to connect to Manager (Default: 8443)
	-prot <string>	Protocol; only use https (Default: https)
	-u <name>	User name
	-append	Append to the output file rather than create a new one and overwrite any existing one
	-sanitize	Sanitize the IP addresses and host names
	-h	Help
Examples	To run: arcsight managerinventory	

manager-reload-config

This command loads the `server.defaults.properties` and `server.properties` files on the Manager.

manager-reload-config

Applies to	Manager	
Syntax	arcsight manager-reload-config	
Parameters	-diff	Displays the difference between the properties the Manager is currently using and the properties that this command loads

manager-reload-config, continued

	-as	Forces the command to load properties that can be changed without stopping and starting the Manager. The properties that require a Manager restart are updated in the <code>server.properties</code> but are not effective until the Manager is stopped and all services are started.
	-t <seconds>	Number of seconds after which the <code>manager-reload-config</code> command stops trying to load the updated properties file on the Manager
Examples	To reload config: <pre>arcsight manager-reload-config</pre> To view the differences between the properties the Manager is currently using and the properties that this command loads: <pre>arcsight manager-reload-config -diff</pre>	

managersetup

This command allows you to configure the manager by launching the Manager Configuration Wizard. You can launch the wizard in console mode by using the `-i console` option while running the command in a terminal window. Run it without any option to launch the wizard in the graphical user interface mode when you have the X Window system installed and wish to use it. For more information about using the wizard, see ["Running the Manager Configuration Wizard" on page 131](#). The options are all optional.

Note that using the X Window system (to run the Manager Configuration Wizard) is not preferred, but if you have it installed and want to use it, you do not have to use the `-i console` option. The X Window system is not present on ESM on an appliance.

If issues occur while running the Manager Configuration Wizard, this command logs troubleshooting information in a log file:

```
/opt/arcsight/var/logs/manager/default/serverwizard.log.
```

Distributed Correlation Cluster Only: You can only run this command on a configured persistor node.

Important: When you run `managersetup`, you will see:

Do you want to run ESM in Compact or Distributed mode?

In this case, accept the default and keep moving through the wizard. You cannot change the ESM mode in the wizard after the initial installation of ESM. If you want to change ESM from compact to distributed mode, see the topic "Converting Compact Mode to Distributed Correlation Mode" in the *ESM Upgrade Guide* and follow that conversion process. Note that conversion from distributed correlation mode to compact mode is not supported.

managersetup

Applies to	Manager	
Syntax	managersetup [options]	
Parameters	-i <mode>	<p><code>console</code> -- you answer configuration questions in a terminal window. Use no other options. This is the preferred mode of operation, and the only mode available for ESM on an appliance. Use the <code>-i console</code> mode if you get this error when you attempt to run in the Manager Configuration Wizard: <code>Could not initialize class sun.awt.X11GraphicsEnvironment.</code></p> <p><code>swing</code> -- You answer the same questions in a graphical user interface. Use no other options.</p> <p><code>silent</code> -- Followed by the <code>-f</code> option, the configuration is read from a file that was created by the <code>recordui</code> mode or the <code>-g</code> option. Use no other options besides <code>-f</code>.</p> <p><code>recordui</code> -- You provide a file path and name and then answer questions in GUI mode while configuring this system. Your configuration is recorded in the specified file for use with the <code>silent</code> mode on some other system. Use no other options.</p> <p>Blank (no <code>-i</code> option at all) means Swing mode.</p>
	-f <file>	The name of the file to use when running in <code>-i silent</code> mode.
	-g	<p>Generate sample properties for <code>-i silent</code> mode. The sample properties are sent to <code>stdout</code>, but you can redirect this output to a file. If you edit the file to provide your own configuration information, you can use it as the file in the <code>-f</code> option in <code>silent</code> mode.</p> <p>Use no other options.</p>
Examples	<p>To run:</p> <pre>arcsight managersetup -i console (runs in console mode)</pre> <pre>arcsight managersetup (runs in GUI mode)</pre> <pre>arcsight managersetup -g > /opt/mysetup.file(generate sample to edit)</pre> <pre>arcsight managersetup -i silent -f /opt/mysetup.file (configures system from file)</pre>	

managertreaddump

This command runs a script to dump the Manager's current threads. The threads go into `manager/logs/default/server.std.log`. Do not inadvertently add a space between `manager` and `treaddump`, doing so causes the Manager to restart. Specify this file when running `treaddumps`, which provides a convenient HTML file with links to all the thread dumps in a summary format.

Distributed Correlation Cluster Only: You can only run this command on a persistor node.

managethreaddump

Applies to	Manager
Syntax	managethreaddump
Parameters	None
Examples	To run: arcsight managethreaddump

managerup

This command gets the current state of the Manager. Returns 0 if the Manager is running and reachable. Returns 1 if it is not.

Distributed Correlation Cluster Only: You can only run this command on a persistor node.

managerup

Applies to	Manager
Syntax	managerup
Parameters	None
Examples	To check that the Manager is up, running, and accessible: arcsight managerup

mbussetup

This command applies to ESM in distributed mode. It allows you to setup message bus services in a distributed correlation environment. This command is described in the topic "[Configuring Services in a Distributed Correlation Cluster](#)" on page 60, "Configuring a Message Bus Control and Message Bus Data".

monitor

This command is used with the Network Management Systems.

monitor

Applies to	Manager
Syntax	monitor

monitor, continued

Parameters	-a <filter>	Attribute filter. Default: "*"
	-append	Append to output file instead of overwriting (Default: false)
	-f <filter>	Object filter. Default: "Arcsight:*"
	-m <host>	Manager host name or address
	-o <op>	Operation {list, show}. Default is list
	-out <file>	Output filename for management service information. Default is stdout
	-p <pwd>	Password
	-sanitize	Sanitize IP address and host names (Default: false)
	-u <name>	User name
Examples	To run: <code>arcsight monitor</code>	

netio

This command is a simple network throughput measurement utility.

Distributed Correlation Cluster Only: You can run this command on a cluster node that is running instances of aggregators and correlators.

netio

Applies to	Manager	
Syntax	<code>netio</code>	
Parameters	-c	Client mode (Default: false)
	-n <host>	Host to connect to (Client mode only)
	-p <port>	Port (Default: 9999)
	-s	Server mode
Examples	To run: <code>arcsight netio</code>	

package

This command imports or exports resources (users, rules, and so on) to or from one or more XML files (.arb files).

Distributed Correlation Cluster Only: You can run this command on a cluster node that is running instances of aggregators and correlators.

Use this command instead of the archive command. Refer to the "Managing Packages" topic in the *ArcSight Console User's Guide* for information on performing these and other functions from the ArcSight Console.

package

Applies to	Manager, Database, Console	
Syntax	package -action <action-to-be-taken> -package <package URI> -f <package-file>	
Parameters	- action <action>	Creates a new package based upon one or more packages that you specify. The possible actions include bundle, convertarchives, export, import, install, uninstall. The default is export
	-config <file>	The primary configuration file to use. Default is config/server.defaults.properties
	-convertbaseuri <baseuri>	The base URI for packages that are converted from archives. This option is only used in conjunction with the -actionconvertarchives option
	-f <path>	The location of the package .arb bundle file. File name paths can be absolute or relative. Relative paths are relative to <ARCSIGHT_HOME>
	-m <manager>	The Manager to communicate with
	-p <password>	The password with which to log in to the Manager. A password is not needed and not used in standalone mode, because the connection is made using the stored database account. Password is required otherwise.
	-package <packagerefs>	The URI(s) of the package(s). This option is used in conjunction with -action install and -action uninstall in order to list which packages to operate upon
	-pc <privateConfig>	This configuration file overrides the server.defaults.properties file. The default location is config/server.properties
	-pkcs11	Use this option when authenticating with a PKCS#11 provider. For example, arcsight package -m <hostname> -pkcs11 -f <file path>
	-port <port>	The port to use for communication. The default port used is 8443
	-source <sourcefile>	The source file. This is used in conjunction with the -f command which specifies an output file

package, continued

	-u <username>	The user name used for logging in to the Manager
	-standalone	Operate directly on the Database not the Manager
Examples	<p>To convert a previously archived package:</p> <pre>arcsight package -action convertarchives -convertbaseuri "/All Packages/Personal/Mypackage" -source sourcefile.xml -f packagebundle.arb</pre> <p>To install a package:</p> <pre>arcsight package -action install -package "/All Packages/Personal/Mypackage" -u username -p password -m managename</pre> <p>To uninstall a package:</p> <pre>arcsight package -action uninstall -package "/All Packages/Personal/Mypackage" -standalone -config /config/server.defaults.properties -pc /config/server.properties</pre> <p>To import a package through the Manager:</p> <pre>arcsight package -action import -f packagebundle.arb -u username -p password -m managename</pre> <p>To export a package:</p> <pre>arcsight package -action export -package "/All Packages/Personal/Mypackage" -f packagebundle.arb -u username -p password -m managename</pre> <p>To export multiple packages:</p> <pre>arcsight package -action export -package "/All Packages/Personal/PackageOne" -package "/All Packages/Personal/PackageTwo" -f packagebundle.arb -u username -p password -m managename</pre> <p>To export packages in a standalone mode (directly from the database) Make sure that the Manager is not running:</p> <pre>arcsight package -action export -package "/All Packages/Personal/Mypackage" -f packagebundle.arb -u username -p password -standalone -config server.default.properties -pc server.properties</pre> <p>To combine xml files from multiple packages into one package:</p> <pre>arcsight package -action bundle -f myPkgNew.arb -source chnpkg.xml -source filterpkg.xml -source rulepkg.xml</pre> <p>In the above example, chnpkg.xml, filterpkg.xml, and rulepkg.xml files are extracted from their respective packages and are bundled in one package bundle called myPkgNew.arb.</p>	

portinfo

This command runs a script used by the portinfo tool of the Console. Displays common port usage information for a given port.

portinfo

Applies to	Console	
Syntax	portinfo port	
Parameters	port	Port number
Examples	To run: arcsight portinfo	

reenableuser

This command re-enables a disabled user account.

Distributed Correlation Cluster Only: You can only run this command on a persistor node.

reenableuser

Applies to	Manager	
Syntax	reenableuser <username>	
Parameters	<username>	The name of the user resource to re-enable
Examples	To re-enable a disabled user: arcsight reenableruser <username>	

refcheck

This command is a resource reference checker.

Distributed Correlation Cluster Only: You can only run this command on a persistor node.

refcheck

Applies to	Manager	
Syntax	refcheck	
Parameters	None	
Examples	To run: arcsight refcheck	

regex

This command is a graphical tool for regex-based FlexConnectors.

regex

Applies to	SmartConnectors
Syntax	regex
Parameters	None
Examples	To run: arcsight regex

replayfilegen

This command runs a wizard for creating security event data files ("replay files") that can be run against a Manager for testing, analysis, or demonstration purposes.

Note: This is a client side command only and should be executed from the Console's ARCSIGHT_HOME/bin directory.

replayfilegen

Applies to	Console
Syntax	replayfilegen -m mgr [parameters]
Parameters	-f <file> Log file name (properties file in -i silent mode)
	-g Generate sample properties file for -i silent mode
	-i <mode> Mode: console, silent, recorderui, swing
Examples	Run from the Console's <ARCSIGHT_HOME>/bin directory: arcsight replayfilegen To run in console mode: arcsight replayfilegen -i console

reposestap

This command applies to ESM in distributed mode. It allows you to setup information repository services in a distributed correlation environment. This command is described in the topic "[Configuring Services in a Distributed Correlation Cluster](#)" on page 60, "Configuring a Repository".

resetpwd

This command runs a wizard to reset a user's password and optionally notify the user of the new password by e-mail.

Distributed Correlation Cluster Only: You can only run this command on a persistor node.

resetpwd

Applies to	Manager	
Syntax	resetpwd	
Parameters	-f <file>	Log file name (properties file in -i silent mode)
	-g	Generate sample properties file for -i silent mode
	-i <mode>	Mode: console, silent, recorderui, swing
	-h	Display command help
Examples	To reset a user's password: arcsight resetpwd	

resvalidate

This command checks for whether there are any invalid resources in the database. The utility generates two reports called `validationReport` (with `.xml` and `.html` extensions) that are written to the directory from which you run the `resvalidate` command. Make sure you stop the Manager before you run this command. If you have more than 50,000 actors you should first increase your Java heap size to 8 GB before running this command.

Note: After running the `resvalidate` command, check `/opt/arcsight/var/logs/manager/default/resource-validation.log` to determine resources that were skipped due to incorrect definitions.

Distributed Correlation Cluster Only: You can only run this command on a persistor node.

resvalidate

Applies to	Manager, Database
Syntax	resvalidate

resvalidate, continued

Parameters	-excludeTypes <exclude_resource_names>	Resource type to exclude from being checked; for example, Rule, DataMonitor If specifying multiple resource types to exclude, use comma to separate them. Resource type - Rule,DataMonitor(comma separated)
	-out <output_dir>	Output directory for validation report. If none is specified, the report is placed in the directory from which you run the resvalidate command
	-persist [false true]	If a resource is found to be invalid, whether to mark it invalid or only report it as invalid. For example, a rule depends on a filter that is missing. When you run the resvalidate command and -persist=false, the rule is reported as invalid but not marked invalid. However if -persist=true, the rule is marked as invalid. Default: persist=true.
Examples	<p>In general, if you need to run the resource validation script, run it twice: the first time with '-persist true' (default) to validate and fix invalid resources, and the second time with '-persist false' to generate a correct report:</p> <pre>arcsight resvalidate arcsight resvalidate -persist false</pre>	

searchindex

This command creates or updates the search index for resources.

If you provide the credentials for the Manager, it automatically associates with the newly created or updated index. However, if you do not specify any credentials, manually configure the Manager to use the updated index.

The searchindex command must be deployed on the machine where the ESM Manager is installed.

Distributed Correlation Cluster Only: You can only run this command on a persistor node.

searchindex

Applies to	Manager
Syntax	searchindex -a action

searchindex, continued

Parameters	-a <action>	Possible actions: create, update, or regularupdate. The -a parameter is required. create—Creates a new search index. update—Updates all resources in the index that were touched since the last daily update was run. Although "update" is a scheduled task that runs daily, you can run it manually. regularupdate—Updates all resources in the index that were touched since the last regular update was run. Although "regular update" is a scheduled task that runs every 5 minutes, you can run it manually.
	-t <time>	Time stamp that indicates starting when the resources should be updated
Examples	To run: arcsight searchindex -a <action> For example, arcsight searchindex -a create	

Note: If you get an error in the server log for the `searchindexutility` that says `outofmemoryError`, you can increase the cap on the Java heap size. Go to your environment variables and increase the value for the variable called `ARCSIGHT_SEARCH_INDEX_UTILITY_JVM_OPTIONS`.

Set the variable like the following example:

```
ARCSIGHT_SEARCH_INDEX_UTILITY_JVM_OPTIONS="-Xms512m -Xmx8192m"  
export ARCSIGHT_SEARCH_INDEX_UTILITY_JVM_OPTIONS
```

`Xms` is the initial Java heap size. `Xmx` is the maximum. The above values are the defaults.

When that variable is set, it takes priority over the default settings as well as `ARCSIGHT_JVM_OPTIONS`.

sendlogs

This command runs a wizard to sanitize and save ArcSight log files so that you can send them to customer support for analysis, if they instruct you to do so. Log files are under `/opt/arcsight/var/logs` for both distributed and compact modes. Logs larger than 1GB are skipped in all cases. When listing hostnames to sanitize, the hostnames must be comma-separated.

You can run `sendlogs` from the ArcSight Console, the Manager command line, or in the context of log retrieval in the ArcSight Command Center.

The `sendlogs local` log collects log files for:

- ArcSight Console
- Connectors (when selected and running)
- Manager (from all nodes if running distributed mode)
- DBMS (CORR-Engine)
- Analytic files (for example, thread dumps)
- ArcSight Command Center (Manager logs and Logger configuration and log files)

The `sendlogs Local Logs Only` functionality is available if you run `sendlogs` from command line or from the ArcSight Console as an administrative user; it is not available if you run `sendlogs` as a non-administrative user. Also, if you are not logged in as a non-administrative user, you can choose only the sanitizer mode.

Note: The `sendlogs` command does not send the log files.

Note: For ESM in **distributed mode**, run this command only on the **persistor** node. Running `sendlogs` from a non-persistor node in a cluster is not supported.

For distributed mode (and when run from the persistor node) the `sendlogs local` log will collect all cluster logs, except for:

- DBMS (CORR-Engine)
- Analytic files (for example, thread dumps)

sendlogs

Applies to	Manager, Database, Console	
Syntax	<code>sendlogs</code>	
Parameters	<code>-n <num></code>	Incident number (Quick mode)
Examples	<code>arcsight sendlogs</code>	

syncpreferip

This command synchronizes the IP preference of cluster nodes with the choice you make on the persistor node when you run `managersetup`. See ["Changing the Internet Protocol Version in a Distributed Correlation Environment" on page 85](#) for details on changing the IP preference.

Distributed Correlation Cluster Only: You can only run this command as user `arcsight` on a persistor node.

syncpreferip

Applies to	Manager
Syntax	syncpreferip [no parameters]
Example	arcsight syncpreferip

tee

This command displays the output of a program and simultaneously writes that output to a file.

Distributed Correlation Cluster Only: You can run this command on a cluster node that is running instances of aggregators and correlators.

tee

Applies to	Manager	
Syntax	-f <filename>	
Parameters	-a	Append to the existing file
Examples	To run: arcsight tempca -i arcsight tee -f sslinfo.txt	

tempca

This command allows you to inspect and manage demo certificates.

tempca

Applies to	Console	
Syntax	tempca	
Parameters	-a <alias>	Key store alias of the private key to dump
	-ac	Add the demo CA's certificate to the client truststore
	-ap	Create demo SSL key pair and add it to the Manager key store
	-dc	Dump/export the demo CA's certificate to a file (demo.crt) for browser import
	-dpriv	Dump private key from the Manager key store
	-f <file>	Filename to write the demo CA's certificate to
	-i	Display summary of current SSL settings
	-k <n>	Key store: Manager (1)

tempca, continued

	-n <host>	Host name of the Manager (opt for the creation of a demo key pair)
	-nc	No chain: Do not include certificate chain (option for creation of a demo key pair)
	-rc	Reconfigure not to trust demo certificates. Removes the demo CA's certificate from the client truststore
	-rp	Remove pair's current key pair from the Manager key store
	-v <days>	Validity of the new demo certificate in days (Default: 365)
Examples	To run: arcsight tempca	

threaddumps

This command extracts and reformats thread dumps from the file to which you wrote the thread dumps in the `managertreaddump` command (manager/logs/default/server.std.log). The output is an html file in the bin directory from which you run this command. It provides a list of links to all the thread dumps in a summary format.

Distributed Correlation Cluster Only: You can run this command on a cluster node that is running instances of aggregators and correlators.

threaddumps

Applies to	Manager	
Syntax	threaddumps <file>	
Parameters	<filename>	Specify the name of the thread-dump file.
	-h	Display command help
Examples	To run: arcsight threaddumps	

tproc

This command is a standalone Velocity template processor.

tproc

Applies to	Manager	
Syntax	tproc	
Parameters	-d <file>	Definitions file

tproc, continued

	-Dname=value	Defines
	-h	Display command help
	-l	Keep log file
	-o <file>	Output file
	-p <file>	Properties file
	-t <file>	Template file
	-v	Verbose mode
Examples	To run: arcsight tproc	

updaterepohostconfig

This command applies to distributed correlation mode only.

This command prompts for new hostnames corresponding to old hostnames in the distributed correlation cluster and makes changes to the configuration.

updaterepohostconfig

Applies to	Manager	
Syntax	consolesetup [-i <mode>]	
Parameters	-i <mode>	Mode: console, silent, recorderui, swing

whois

This command is used by the `whois` command of the console

whois

Applies to	Console	
Syntax	whois [-p <port>] [-s <host>] <target>	
Parameters	-p <port>	Server port
	-s <host>	Name or address of 'whois' server
	<target>	Name or address to lookup
Examples	To run: arcsight whois	

zoneUpdate

This command updates IPv4 and IPV6 address allocations and dark space information that are provided in the periodic Zone Update Subscription Package, contained in the `Zone_Updates_<version>.zip` file. Then, at the command line, run the `zoneUpdate` command to apply the zone updates. Use of this command is optional. You can use `zoneUpdate` after a successful Manager installation or upgrade. This command is available from the command line only, and has no GUI functionality.

Distributed Correlation Cluster Only: Run this command from the persistor node **only**.

Running `zoneUpdate` requires an ESM administrator login and password. While the process is running, do not use the same administrator account to access the ArcSight Console or ArcSight Command Center for other administrative tasks. Allow up to 50 minutes or longer for a first-time zone update, depending on the manager workload and the number of assets assigned to the global network. Subsequent incremental updates should not take as long. While `zoneUpdate` is running, other ESM administrators and users may access the Console or Command Center.

`zoneUpdate` performs these actions in the Global network:

- Inventories affected assets
- Removes old zones
- Installs and updates zones
- Auto-zones assets that appeared in the inventory of affected assets in the Global network

`zoneUpdate` updates zones in the Global network only. Local zones are not updated by this command. The behavior of `zoneUpdate` is the same for both dynamic and static zones.

Best Practices for Importing Packages

If you need to perform zone updates and/or operate under high loads, disable the ability to move resources before updating zones. This makes large packages more likely to import successfully. Verify the property `esm.manager.disable.resource.move` is set to `true` in `server.properties`.

To set the `resource.move` property to `true`, add this statement to `server.properties`:

```
esm.manager.disable.resource.move=true
```

Refer to ["Editing Properties Files" on page 16](#) for details on editing the `server.properties` file.

Recommendations

- Allocate assets to the local network only and that the Global network does not contain assets. Also, zones that have categories assigned to them, and then are removed and reinstalled as part of the zone update process lose the category assignments. Do not assign categories to the system zones.
- Perform a full system database table backup (`export_system_tables`) and export the current ArcSight Network package before using `zoneUpdate`, to ensure that you have a usable snapshot of your network model. If the zone update process is interrupted or a problem occurs and you must revert your data, be sure to use this backup to restore your ArcSight resources before attempting to run `zoneUpdate` again.
- Run `zoneUpdate` during non-peak system time.

Running `zoneUpdate`

Note: Zone Groups belonging to Regional Internet Registries (RIR) that contain more than 1000 zones will place their corresponding zones in subgroups, each group containing up to 950 zones, to enable you to better manage those zones, and content related to them, from within the ArcSight Console.

1. Log in as user *arcsight* on the system where Manager runs.
2. Verify that the Manager is running.
3. Extract the `Zone_Updates_<version>.zip` file into a directory. The directory can be of your choice. The zipped files extract into the folder `ArcSight_Networks_<version>`, which contains the files `ArcSight_Networks.arb` and `Zone_Removal_Tool.xml`. Do not change the name of this folder or the names of the extracted files.
4. Verify that the user *arcsight* has write permissions to the directory into which you extracted `Zone_Updates_<version>.zip`.
5. As user *arcsight*, run this command:

```
<ARCSIGHT_HOME>/bin/arcsight zoneUpdate -m<Manager hostname or IP address>  
-u<ESM user with administrative privileges> -f<folder where zip file was  
extracted>
```

You are prompted for the ESM user's password. **Be sure to enter the correct password.** `zoneUpdate` uses the entered password several times, and temporarily locks you out if you use the wrong password. If this happens, you can reenab the user or wait for the user to be reenabled automatically.

Running `zoneUpdate` can take longer than 50 minutes, depending on Manager workload and the number of assets assigned to the Global network. `zoneUpdate` will print its progress while it processes zone data. The information is also written to `zoneUpdate.log` in Manager's log directory.

Warning: Do not interrupt or kill `zoneUpdate` after the process starts. Allow `zoneUpdate` to complete, and then make a determination of the condition of your

zones and whether to install another version of the Zone Update Subscription package.

Recovery and Troubleshooting

Zone Updates Not Applied

If zoneUpdate runs with errors, and does not apply the zone updates from the Zone Update Subscription Package, follow these steps:

1. Stop the Manager and then start all services. Run the following commands to do so:

```
/etc/init.d/arcsight_services stop manager
```

```
/etc/init.d/arcsight_services start all
```

2. Run zoneUpdate again.
3. If the above steps do not work, and you encounter the same errors as before, import the full system database table backup (export_system_tables) and the current ArcSight Network package that you exported before initially running zoneUpdate.
4. Run zoneUpdate again.

Package Exists Error When Applying the Zone Update Subscription Package

If you encounter these messages when running zoneUpdate:

```
Reading bundle 'Common Bundle Alias' Done. 0 min 0 sec 41 ms
Importing 1 packages
Importing package 1/1 '/All Packages/ArcSight System/ArcSight Networks'
Parsing archive 'ArcSight Networks.xml'... Done. 0 min 1 sec 19 ms
Package Already Exists with Newer Content
```

```
-----
Package '/All Packages/ArcSight System/ArcSight Networks' already exists in
the system with newer content
```

- 1: Leave newer package
- 2: Never override newer packages
- 3: Update package
- 4: Always update Packages
- 5: Abort

```
-----
Choose option 3, Update Package.
```

Asset Zoning

Assets that were zoned in the Global network before you run zoneUpdate will be zoned after the command completes.

Asset Ranges

Asset ranges are not auto-zoned by zoneUpdate. Asset ranges will be unzoned by the running of the zoneUpdate; you must manually rezone asset ranges after you run zoneUpdate if you had asset ranges in the Global network.

For example, if you had an asset range in Zone A in a previous version of ESM, the asset range is unzoned after you run zoneUpdate. For this example, suppose Zone A was split into two zones, Zone A and Zone B, and after upgrade your asset range spans the last part of Zone A and first part of Zone B. In this case, the asset range becomes unzoned. To recover zoning, you must open each unzoned asset range resource and map it to the correct zone, or split it into two asset ranges that map to the new Zones A and B.

zoneUpdate

Applies to	Manager	
Syntax	<code><ARCSIGHT_HOME>/bin/arcsight zoneUpdate -m <Manager hostname or IP address> -u <user with administrative privileges> -f <folder where zip file was extracted></code>	
Parameters	<code>-m <manager></code>	The Manager hostname or IP address. Use of a hostname or an IP address depends on whether your Manager was configured using a hostname or an IP address.
	<code>-u <username></code>	The name of a user with administrative privileges. For example, admin1 or admin2.
	<code>-f <folder></code>	The path to the folder or directory that contains the unzipped Zone Update Subscription package. For example, /opt/arcsight/manager. Extract the file Zone_Updates_<version>.zip into this folder, and give the folder write permission.
	<code>-h</code>	Help
Example	To update zones: <code><ARCSIGHT_HOME>/bin/arcsight zoneUpdate -m 192.0.2.0 -u admin2 -f <ARCSIGHT_HOME></code>	

CORR-Engine ArcSight Commands

These commands are used to manage data in the CORR-Engine. They are located in /opt/arcsight/logger/current/arcsight/logger/bin.

To run a CORR-Engine ArcSight command script, open a command window and switch to the /opt/arcsight/logger/current/arcsight/logger/bin directory. These arcsight commands run using the file arcsight.sh in that location. The general syntax is as follows:

```
arcsight <command_name> [parameters]
```

configbackup

Description	The <code>configbackup</code> command backs up certain essential configuration information such as search settings and the configuration of archives (not the archives themselves). It places this backup in a file called <code>configs.tar.gz</code> which you can find in <code>/opt/arcsight/logger/current/arcsight/logger/tmp/configs</code> .	
Applies to	CORR-Engine	
Syntax	<code>arcsight configbackup</code>	
Parameters	none	
Example	To run: <code>/opt/arcsight/logger/current/arcsight/logger/bin/arcsight configbackup</code>	

Make sure you are familiar with these guidelines before you create a backup file:

The `configbackup` command creates the `configs.tar.gz` file, which you must then copy to a safe location.

Make a note of the following, which must match exactly on the machine to which you restore:

- Operating system and version
- Path to the archive locations for each storage group
- ESM version
- MySQL password

disasterrecovery

Description	This command restores the data backed up using the <code>configbackup</code> command.	
Applies to	CORR-Engine	
Syntax	<code>arcsight disasterrecovery start</code>	
Parameters	start	
Example	To run: <code>/etc/init.d/arcsight_services stop logger_servers</code> <code>cp ~/configs.tar.gz /opt/arcsight/logger/current/backups/configs.tar.gz</code> <code>/opt/arcsight/logger/current/arcsight/logger/bin/arcsight disasterrecovery start</code> <code>/etc/init.d/arcsight_services start logger_servers</code>	

Make sure you are familiar with these guidelines before you restore a backup file:

- When you restore this data, the existing data is deleted.

This command restores the specific 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.

This includes event data. The assumption is that you are restoring this configuration to a new, clean installation with no event data, or at least none that needs to be preserved. Restore the content to a machine where the following characteristics are exactly the same as the backup machine:

- The version of ESM must be the same
- The version of the operating system (and the time zone to which it is set) must be the same
- The archive locations for the backed-up storage groups must already exist and be the same
- The MySQL password must be the same

exportdatausage

Description	<p>ESM keeps track of event counts and size from each connector. Use this command to export this event data as a comma-separated values (CSV) file. You can use this information to track the event throughput by connector.</p> <p>Note: This command has to be run from a different location than the other <code>arcsight</code> commands. Run it from: <code>/opt/arcsight/logger/current/arcsight/logger/bin</code></p>	
Applies to	CORR-Engine	
Syntax	<code>./exportdatausage <path/file></code>	
Optional Parameter	<code><path/file></code>	<p>Specify the path and name of the CSV file to which to export the usage data. It can be a relative or absolute path. You do not need to specify the <code>.csv</code> extension.</p> <p>If you do not specify this parameter, the data is displayed on screen.</p>
Examples	<p>To create a file called <code>usagefile.csv</code> in <code>/opt/arcsight</code>, run:</p> <pre>arcsight exportdatausage /opt/arcsight/usagefile</pre>	

Appendix B: Troubleshooting

The following information may help solve problems that occur while operating the ArcSight system. In some cases, the solution can be found here or in specific ArcSight documentation, but Customer Support is available if you need it.

If you intend to have Customer Support guide you through a diagnostic process, prepare to provide specific symptoms and configuration information.

General Troubleshooting

Rule action does not populate variables such as `$message` and `$deviceCustomString(s)` in correlated events

When a rule fires, it passes the correlated event to notifier for email notification. The rule engine automatically populates some of the fields in the correlated event, such the name or agent-related fields. Note that not all fields are populated. To populate fields with base event values, the fields must be aggregated fields.

Note: If the aggregation threshold is greater than 1, the rule will not fire until all conditions, including aggregation condition, are matched.

For example, if you want to populate the `$message` and `$deviceCustomString5` fields, you can modify the rule to add these two fields to the "aggregate only if these fields are identical" list.

Note: You must add the fields to the "aggregate only if these fields are identical" list. If you add them to the rule action notification's subject, they will not be passed into email notification.

If you forward the original correlated base event while including the `setevent` field requirements for `$deviceCustomString1 - $deviceCustomString5`, the forwarded event populates only the base event ID. It does not populate all the fields on the base event. To ensure the forwarded event populates all base event fields, include aggregation settings in the rule.

This applies to variables such as `$message` and `$deviceCustomString(s)`, but the forwarded event populates other variables, such as `$name`, without aggregation.

You changed your File system format from XFS to EXT4 or back and now you have problems.

Note: You cannot change the file system type on ESM.

Both XFS and EXT4 file system formats are supported during installation. However, ESM configures itself to the file system upon which it is first installed; you therefore cannot change file system type after installation, even during an upgrade. Roll your file system back to what it was before.

Your license expired and you cannot start the ArcSight Command Center to specify a new license file.

Run the `arcsight managersetup` command as documented in ["Running the Manager Configuration Wizard" on page 131](#).

Report is empty or missing information

Check that the user running the report has inspect (read) permission for the data being reported.

Running a large report crashes the Manager

A very large report (for example, a 500 MB PDF report) might require so much virtual machine (VM) memory that it can cause the Manager to crash and restart. To prevent this scenario, you can set up the Manager to expose a special report parameter for generating the report in a separate process. The separate process has its own VM and heap, so the report is more likely to generate successfully. Even if the memory allocated is still not enough, the report failure does not crash the Manager.

This option must be set up on the Manager to expose it in the Console report parameters list. The steps are as follows:

1. On the Manager in the `server.properties` file, set `report.canarchivereportinseparateprocess=true`. This makes a new report parameter available on the Console.
2. Save the `server.properties` file.
3. Stop the Manager:

```
/etc/init.d/arcsight_services stop manager
```

Then start all services:

```
/etc/init.d/arcsight_services start all
```

4. On the ArcSight Console, open the report that you want to run in a separate process in the Report Editor, and click the **Parameters** tab. Set the parameter **Generate Report In Separate Process** to `true`.
5. Run the report. The report should run like a normal report, but it does not consume the resources of the Manager VM.

Note: Use this parameter only if you experience a Manager crash when running large reports such as the ones that contain tables with more than 500,000 rows and 4 or 5 columns per row.

Scheduled rules take too long or time out

If you have a system, perhaps one with a high EPS, in which the scheduled rules are not running quickly enough, you can enable them to run in parallel (multi-threading) to speed them up. Add the following property to the `server.properties` file:

```
rules.replay.run.parallel=true
```

You can also set the number of threads to use, as follows (the default if you do not use this property is four threads):

```
rules.replay.numthreads=<number of threads to use>
```

Some Central European, Cyrillic, and Asian language fonts do not display properly when generating reports in PDF

This problem occurs because some Central European, Cyrillic, or Asian language fonts that are TrueType fonts are not supported directly by versions of Adobe Reader earlier than version 8.0. In order to work around this, each TrueType font must be mapped to an OpenType font supported in Adobe Reader 8.0. ArcSight provides this mapping in the `<ARCSIGHT_HOME>/i18n/server/reportpdf_config_<locale>.properties` file. You have the option to change the default mapping of any TrueType font to the OpenType font by modifying the respective font mapping in this file.

To work around this issue:

1. Install a localized Adobe Reader 8.0 depending on the language of your platform on your Manager machine. This version of the Adobe Reader installs the OpenType fonts by default.
2. Edit the `server.properties` file as follows:
 - a. Set `report.font.truetype.path` property to point to the directory that contains the TrueType and OpenType font. Use ":" as a path separator in Unix. On Unix platforms, the TrueType font path may differ depending on the specific Unix platform, but it is typically `/usr/lib/font`. The CIDFont directory is always the same relative to the Adobe Reader installed directory. So, the default directory would be `/usr/lib/font:<adobe_reader_dir>/Resource/CIDFont`.
 - b. Set `report.font.cmap.path` property to point to Adobe Reader's CMap directory. On Unix, the CMap path is relative to the Adobe Reader installation -- `<adobe_reader_dir>/Resource/CMap`.

E-mail notification doesn't happen

If you receive the following error:

```
[2009-12-03 14:31:33,890][WARN ]  
[default.com.arcsight.notification.NotifierBase][send] Unable to send out e-  
mail notification, notifications have not been configured.
```

- Verify the following properties are set in the `server.properties` file:

```
notifications.enable=true
```

- Check `server.properties` file to find which SMTP server is associated with the Manager. Make sure that the SMTP server is up and running.

Review the Notification resource and confirm the e-mail address and other configuration settings.

Notification always escalates

Check the `server.properties` file to find which POP3 or IMAP server is associated with the Manager. Make sure that the POP3 or IMAP server is up and running, in order to process acknowledgements from notification recipients.

Event IDs have negative numbers

Negative event ID numbers can occur, and are normal. Event IDs are 64-bit values. The less-significant 48 bits are assigned to a newly received event by the receiving Manager; these bits uniquely identify the event in the database of that Manager. The more-significant 16 bits are used to store forwarding information. When an event ID with '1' in the topmost bit is represented as Java 'long' value, the event ID value is interpreted as a negative number according to JVM rules. When displayed, such an event ID appears as a decimal number with a sign '-' in front of it.

Rule Recovery Timeout Occurs

Rule recovery can timeout if there is a high EPS on the system, which causes the server to stop loading events from the database for checkpoint. You can modify the `rules.recovery.time-limit` property in `server.properties` to set a higher recovery time limit to attempt to prevent this timeout. The default value for the `rules.recovery.time-limit` property is 120 seconds (two minutes).

Note: The timeout can still occur even after you increase the time limit, due to overall system load, high EPS, or a large number of rules to recover.

For details on editing the `server.properties` file, see ["Editing Properties Files" on page 16](#).

Manager uses decoupled process execution on UNIX-based systems

On UNIX-based systems, Manager uses decoupled process execution to perform specific tasks, for example, to run a very large report. Decoupled process execution uses a stand-alone process executor (instead of using "in process" or "direct process")

execution) and sends commands to be executed via the file system. The process executor uses the <ARCSIGHT_HOME>/tmp directory, so restrict system level access for this directory.

The process executor is used, by default, on all Unix platforms. The Manager scripts ensure that the process executor runs as a daemon before the Manager is started. This has some implications with regards to troubleshooting Manager startup and runtime problems. The Manager, if configured to use the process executor, does not start unless it detects the presence of a running process executor. The process executor runs within its own watchdog, like the Manager, so if the process stops for any reason, it restarts automatically. The process executor is transparent to users regarding how the Manager is started or stopped.

The `stdout` and `stderr` of the executed process are written into the following two files:

```
<ARCSIGHT_HOME>/tmp/[commandfile-name].stdout
```

```
<ARCSIGHT_HOME>/tmp/[commandfile-name].stderr
```

Automatic ArcSight system tasks

These system tasks are scheduled to run automatically one or more times per day, depending on the task.

AUP Updater: This task runs in the manager and pushes to connectors any updated AUP packages it might have.

Dependent Resource Validator: This task runs validations on resources in the system and disables the ones that have problems.

PurgeStaleMarkSimilarConfigs: This task does maintenance work on the 'mark similar' annotation criteria, removing the ones that are stale.

Resource Search Index Updater: This task updates the resource search index.

Sortable Fields Updater: This task keeps sortable event fields synchronized, based on the current indexes in the database.

Table Stats Updater: This task updates statistics on the non-partitioned schema tables, which includes the resource tables.

In a Windows 7 environment, you receive the error message 'findstr' is not recognized as an internal or external command when you attempt to start the ArcSight Console.

Verify that Java is setup correctly in your Windows environment.

Pattern Discovery Performance Troubleshooting

Note: Pattern Discovery is not supported on ESM on an appliance.

Time spread calculations might consume a lot of CPU time, especially if Pattern Discovery has been running for a long time. To determine whether the Time Spread feature is degrading performance, check the `system.log` for the start and end times of the Pattern Discovery process. If the process is taking longer than expected and causing issues in your environment, you can disable the Time Spread feature.

To disable the Time Spread feature, add the property `patterns.timeSpreadCalculation=False` to the Manager's `server.properties` file.

Query and Trend Performance Tuning Troubleshooting

To improve query execution in high-EPS systems, various queries used by the trends in the default ESM system have been optimized. The scheduler allocates two threads for processing system tasks. This alleviates performance issues caused by conflicts between system tasks and user level tasks within the scheduler.

The following sections provide some troubleshooting tips.

server.defaults.properties Entries for Trends

- `trends.query.timeout.seconds=7200`

This is the amount of time that a trend query is allowed to run, in seconds, before the SQL statement times out and the trend query fails. If absent or 0, no time-based timeout is applied.

- `trends.query.timeout.percent=50`

This is the amount of time that a trend query is allowed to run, as a percentage of the query interval for interval trends, before the SQL statement times out and the trend query fails. If absent or 0, no percentage-based timeout is applied.

As an example, with a 50 percent setting, a query covering a start/end time range of 1 hour times out after 30 minutes. A start/end time range covering 1 day would time out after 12 hours.

If both timeouts are specified, the system uses the smaller of the two.

- `trends.query.failures.deactivation.threshold=3`

If this many consecutive "accumulate" (not refresh) runs fail for any reason, the system automatically disables the trend. The check is always performed after any accumulate query run fails. After the threshold is reached, any remaining queries to be executed by this task are skipped. If this setting is absent or 0, the checking mechanism is turned off.

If a trend or query is stopped because of any of the above reasons, an audit event reflects this.

Troubleshooting checklist after stopping the Manager and starting all services

- Use the Console Trend Editor to manually disable any trends that you do not need or that you notice have excessive query times. Disabling these trends helps reduce scheduler and database contention.
- As trend data gathering tasks wake up, the trend attempts to fill in the gaps for missing intervals. Depending on the size of the gaps, this may take some time before the trends catch up.
- A trend does not usually re-run any previously failed runs. If you want to re-run a particular time, you need to manually request it from the Trend Editor.

Disable trend on high-throughput systems

If your system environment typically processes a very large number of events per second (EPS) (such as more than 1000 EPS or 100 million events per day), we recommend that you manually disable the following trend:

```
/All Trends/ArcSight Administration/ESM/User Access/ArcSight User Login Trends - Hourly (Installed by default)
```

How do you know when a trend is caught up?

You can use either of the following techniques, both using the ArcSight Console UI:

- Using the Trend Data Viewer from within the Trends resource tree, you can see at most 2000 rows of data. (Select a trend in the resource tree, right-click, and choose **Data Viewer**.) Sort the trend timestamp column so that the timestamps show newest to oldest and observe when the newest value indicates it has caught up.
- Using the **Refresh...** button in the Trend Editor, set the start time as far back as needed (days or weeks) to see any entries and click Refresh to see which runs show up as available to be refreshed. Only the most recent ones should show first. Note that you should not actually refresh any runs, but only use this technique to see what has been run.

How long does it take for a trend to catch up?

This depends on how long the underlying query interval is, but a trend typically does up to 48 runs, as needed, when it wakes up.

For a trend that queries an entire day and runs once a day, this would allow for more than a month's worth of data to be queried. The data must be present on the system, however, or the query returns no results (but it does not fail).

SmartConnectors Troubleshooting

My device is not one of the listed SmartConnectors

ArcSight offers an optional feature called the FlexConnector Development Kit which may enable you to create a custom SmartConnector for your device.

ArcSight can create a custom SmartConnector. Contact Customer Support.

My device is on the list of supported products, but it does not appear in the SmartConnector Configuration Wizard

Your device is likely served by a Syslog sub-connector of either file, pipe, or daemon type.

Device events are not handled as expected

Check the SmartConnector configuration to make sure that the event filtering and aggregation setup is appropriate for your needs.

SmartConnector not reporting all events

Check that event filtering and aggregation setup is appropriate for your needs.

Some Event fields are not showing up in the Console

Check that the SmartConnector's Turbo Mode and the Turbo Mode of the Manager for the specific SmartConnector resource are compatible. If the Manager is set for a faster Turbo Mode than the SmartConnector, some event details are lost.

SmartConnector not reporting events

Check the SmartConnector log for errors. If the SmartConnector cannot communicate with the Manager, it caches events until its cache is full.

ArcSight Console Troubleshooting

Can't log in with any Console

Check that the Manager is up and running. If the Manager is not running, start it.

If the Manager is running, but you still can't log in, suspect any recent network changes, such as the installation of a firewall that affects communication with the Manager host.

Can't log in with a specific Console

If you can log in from some Console machines but not others, focus on any recent network changes and any configuration changes on the Console host in question.

Console cannot connect to the Manager

If you start an ArcSight Console that could previously connect to the Manager with no trouble, but now it can't, see if the error is similar to:

"Couldn't connect to manager - improper authorization setup between client and manager."

If so, it's likely that the manager has been reconfigured in such a way that it now has a new certificate. Especially if the Console asked you to accept a new certificate when you started it. To fix this, find and delete the certificate that the Console was using before, and then manually import another certificate from the Manager.

Console reports out of memory

If your ArcSight Console is so busy that it runs out of memory, change the memory settings in the `console.bat` or `console.sh` file. This file (for Windows or Linux, respectively) is located in the directory in which you installed the ArcSight Console, in `Console/current/bin/scripts`.

Find the line that starts with `set ARCSIGHT_JVM_OPTIONS=`
Find the parameter `-Xmx512m` (`Xmx` controls the maximum JVM memory).
Change the value to `1024`: `-Xmx1024m`.
Restart the Console for the new setting to take effect.

Acknowledgement button is not enabled

The Acknowledgement button is enabled when there are notifications to be acknowledged and they are associated with a destination that refers to the current user. To enable the button, add the current user to the notification destination.

The grid view of live security events is not visible

To restore the standard grid view of current security events, select **Active Channels** from the Navigator drop-down menu. Double-click **Live**, found at `/Active channels/Shared/All Active channels/ArcSight System/Core/Live`

The Navigator panel is not visible

Press **Ctrl+1** to force the Navigator panel to appear.

The Viewer panel is not visible

Press **Ctrl+2** to force the Viewer panel to appear.

The Inspect/Edit panel is not visible

Press **Ctrl+3** to force the Inspect/Edit panel to appear.

Internal ArcSight events appear

Internal ArcSight events appear to warn users of situations such as low disk space for the ArcSight Database. If you are not sure how to respond to a warning message, contact Customer Support.

The Manager Status Monitor reports an error

The Console monitors the health of the Manager and the ArcSight Database. If a warning or an error occurs, the Console may present sufficient detail for you to solve the problem. If not, report the specific message to Customer Support.

Console logs out by itself

Check the Console log file for any errors. Log in to the Console. If the Console logs out again, report the error to Customer Support.

Duplicate audit events or rule actions after a crash recovery

When you stop ESM, it takes a checkpoint of the rules engine so that it knows where it stopped. If ESM crashes in such a way that it cannot take a checkpoint (power failure, for example), it returns to the last checkpoint when it restarts, and replays events from there. Any actions that occurred between that checkpoint and the ESM crash will therefore be repeated. Repeated actions that generate audit events generate duplicate audit events.

You should investigate repeated actions that do not duplicate well. For example, if an action adds an item to an Active List, that item's counter will be incremented. If the action runs a command, it will run it again, and so on.

You can reduce duplicates by including a rule condition that checks if the relevant entry is already in the active list.

Case data fields appear blank

A number of case fields accept up to 4,000 bytes. However, if you fill too many such fields to the maximum, then you can exceed the limit and the fields can appear blank when you view the case.

This is because of a database limitation on the size of a row (a case, for example), which is about 8k bytes. For large fields, only 768 bytes are stored in the row, along with a 20 byte pointer to the rest, which is stored outside the table. This enables you to have considerably more than 8K of data, but you can still exceed the limit for the database row for a resource.

As a guideline, keep the number of large fields in a case (or other resource with large fields) below ten. The data in the smaller fields contributes to the total, so if you still encounter the problem, consider them as well.

Hostname Shown as IPv6 Address in Dashboard

This can occur due to a mismatch between the system hostname, the network configuration, and your environment's name resolution. Review your system's hosts file and DNS configuration, as well as the addresses found in the DNS for the system hostname.

Manager Troubleshooting

Can't start Manager

The Manager provides information on the command console which may suggest a solution to the problem. Additional information is written to `<ARCSIGHT_HOME>/logs/default/server.std.log`.

Manager shuts down

The Manager stops when it encounters a fatal error. The file `<ARCSIGHT_HOME>/logs/default/server.std.log` has more details about the error condition.

Services do not start after a power failure during "start all"

An unexpected power-off during services startup may result in unavailable postgresql, logger, and manager services. Those services might not start even after rebooting the server.

To resolve the problem, delete the postgresql lock file. The location of the postgresql lock file, is given in the pgsqlog file in `/opt/arcsight/logger/userdata/logs/pgsql/serverlog`. If this problem occurs, the text "could not create lock file" can be written to the server log. To verify, search the server log for instances of the text "could not create lock file".

Reboot the server after removing the postgresql lock file.

Asset aging not working as expected (not all aged assets deleted)

If you are using ESM's asset auto-deletion feature to remove assets from the system, some aged assets may remain in the system. This can occur in environments that have more than 10,000 assets in an asset group. Best practice is not to exceed 10,000 resources for any resource group. If there are more than 10,000 assets in a group, the auto-deletion process slows down and times out without deleting these assets. For details on asset aging see, ["Configuring Asset Aging" on page 44](#).

To solve this problem, you can set certain `asset.aging.task` parameters to gradually delete the unwanted aged assets off of the system. This gradual deletion allows you to delete a relatively small number of assets at a time, keeping the transaction time short and the database load low while the cleanup occurs. This process will gradually delete the aged assets, but can take several days, depending on the number of assets involved and the system load. Stop the manager before making the parameters changes and

restart it when you are done; see ["Restarting the Manager - Stop the Manager and Start All Services" on page 12](#) for details.

To configure the Manager to start the gradual asset deletion process:

1. Add the following properties to the `server.properties` file (in this example,assets will be aged after 4 days and 500 assets will be deleted each hour). See ["Editing Properties Files" on page 16](#) for details.

```
asset.aging.daysbeforedisable=4
asset.aging.task.operation=delete
asset.aging.task.maxassetsprocess=500
asset.aging.task.maxassetsload=500
asset.aging.task.period=Hourly
asset.aging.task.minute=0
```

Notes:

- For the property `asset.aging.daysbeforedisable` note that the default value of -1 means that asset aging is turned off, not that assets will be disabled and deleted. The value for `asset.aging.daysbeforedisable` is expressed in days that define how long an asset is allowed to age before it is disabled and deleted.
 - For the deletion of aged assets to work properly, verify that the `asset.aging.task.operation` property is set to `delete`.
 - Set the properties `asset.aging.task.maxassetsprocess` and `asset.aging.task.maxassetsload` to the same value. The value depends on your hardware and system load. The higher the number specified, the faster the asset deletion will occur. We recommend starting with the value 500 for these two properties, and after the number of assets falls to around 100,000, you can try increasing these properties to 1000.
2. In the `server.defaults.properties` file, verify the value of the property `dbconmanager.provider.logger.pool.maxcheckout`. If the value is less that 3600, add this line to the `server.properties` file:
`dbconmanager.provider.logger.pool.maxcheckout=3600`
 3. Monitor the progress of the asset deletion. When the desired asset limit it reached, stop the process by deleting the properties you added to the `server.properties` file (`asset.aging.daysbeforedisable`, `asset.aging.task.operation`, `asset.aging.task.maxassetsprocess`, `asset.aging.task.maxassetsload`, `asset.aging.task.period`, `asset.aging.task.minute`, and `dbconmanager.provider.logger.pool.maxcheckout`).

The property settings described above are not standard configurations. In the future, monitor the number of assets in groups and do not let them exceed the recommended maximum of 10,000 resources for any resource group.

Switching between daylight savings and standard time can skip a scheduled task

- If the trigger time for a particular scheduled task run happens to fall during the transition time from DST to ST or vice versa, the interval for that particular run gets thrown off. The interval calculation for subsequent scheduled runs are not affected.
- Currently, there are four time zones that are not supported in ESM:
 - Kwajalein
 - Pacific/Kwajalein
 - Pacific/Enderbury
 - Pacific/KiritimatiThese time zones fall in two countries, Marshall Islands and Kiribati.

CORR Engine Troubleshooting

Temporary Sort Space Exceeded

Under some circumstances you can get an error that includes the following:

```
Encountered persistence problem while fetching data: Unable to execute query:  
Temporary sort space limit exceeded
```

Possible solutions include eliminating unnecessary trends, if any, avoid running too many at the same time, and trim queries to return more refined data sets. If the problem persists, try increasing the value of `sort_temp_limit` in `/opt/arcsight/logger/data/mysql/my.cnf`.

For information on creating queries, trends, and reports, refer to the "Building Reports" chapter in the ArcSight Console User's Guide.

If increasing the `sort_temp_limit` is insufficient, and the following circumstance applies, there are two additional remedies.

Excessive temporary file space gets used when Group By (or sorting) is performed on the Event table. If you use Group By (or sorting), use the ArcSight substring function on varchar/string event fields to minimize the data manipulation during grouping. You can use existing local or global variables to achieve this behavior and replace the existing field in the query with the variable. Search in the ArcSight Console User's Guide, in the "Reference Guide" section, for information in variables and substrings.

If the file space usage is still not satisfactory, you can convert the character set automatically to Latin which uses less space. To do so, set the `event.query.charset.conversion` property to 1 in the `/opt/arcsight/manager/config/server.properties` file to convert the existing character set to latin1. Alternatively, set the property to 2 for conversion to binary and

then to Latin (to minimize conversion error for non-English character set). The default value of this property is 0 (zero).

If you use this conversion on multi-byte character sets, it will truncate the characters to single-byte Latin characters, which is likely to render them meaningless. Only use this approach if it's appropriate.

How do I know if my Reactivated Archives are Corrupted?

ESM uses SHA-256 hashing algorithm to create the event data archives.

When the user reactivates the events in an offline archive, ESM validates the hashing of the data in the archive. If the hashes do not match, ESM logs the following error messages in the `logger_server.log` file:

```
A FATAL message "The original archive has: <hash value> , and the files  
have: <different hash value>
```

and/or

```
An ERROR message "supplementalhash computed from data files does not match  
hash in metadata"
```

ESM does not periodically scan for hash mismatches, as the archives may even be moved to external storage, outside of ESM's view. When an archive is moved back and re-activated, it is checked.

SSL Troubleshooting

Cannot connect to the SSL server: IO Exception in the server logs

Causes:

The SSL server may not be running.

- A firewall may be preventing connections to the server.

Resolutions:

- Ensure that the SSL server is running.
- Ensure that a firewall is not blocking connections to the server.

Cannot connect to the SSL server

The hostname to which the client initiates an SSL connection should exactly match the hostname specified in the server SSL certificate that the server sends to the client during the SSL handshake.

Causes:

- You may be specifying Fully Qualified Domain Name (FQDN) when only hostname is expected, or the other way around.
- You may be specifying IP address when hostname is expected.

Resolutions:

- Type exactly what the server reports on startup in `server.std.log` ("Accepting connections at `http://...`")
- For Network Address Translation (NAT) or multi-homed deployments, use `hosts` file to point client to correct IP.

PKIX exchange failed/could not establish trust chain

Cause:

Issuer cannot be found in trust store, the `cacerts` file.

Resolution:

Import issuer's certificate (chain) into the trust store.

Issuer certificate expired

Cause:

The certificate that the SSL server is presenting to the client has expired.

Resolution:

Import the latest issuer's certificate (chain) into the trust store.

Cannot connect to the Manager: exception in the server log

Cause:

If you replaced the Manager's key store, it is likely that the old key store password does not match the new password.

Resolution:

Make sure the password of the new key store matches the old key store. If you do not remember the current key store's password, run the Manager Configuration Wizard on the Manager to set the password of the current key store to match the new key store's password.

Certificate is invalid

Cause:

The timestamp on the client machine might be out of the bounds of the validity range specified on the certificate.

Resolution:

Make sure that the current time on the client machine is within the validity range on the certificate. To check the certificate's valid date range see "[View Certificate Details From the Store](#)" on page 106.

Appendix C: Event Data Transfer Tool

The Event Data Transfer Tool exports ESM events in three formats, cef, csv, and key-value pairs. Getting event data from ESM allows more flexibility to combine analysis with unstructured data in addition to the structured CEF data. Using the Event Data Transfer Tool to export ArcSight ESM events to Hadoop allows you to access Hadoop ecosystem technologies.

ESM and Hadoop - Benefits

After exporting events to Hadoop, you can use Hadoop technologies in the following scenarios:

- Build a security data warehouse using the Hadoop Distributed File System (HDFS), a Hadoop capability that enables inexpensive, long-term storage for Petabytes of data.
- The Event Data Transfer Tool can export ESM events in multiple formats, which enable the usage of Apache Hive to query and manage the security data warehouse. Hive provides HiveQL, a SQL-like language to query and manipulate large events stored on HDFS. See <https://hive.apache.org/>.
- Once the event data is in Hadoop, you can still run searches quickly. There are powerful search engines such as Elasticsearch or Solr open-source technologies. These search engines can provide results in less than a second, which is good for forensics analysis on the data archives. These platforms are built on top of Apache Lucene, the high-performance text engine library. Both technologies also provide visualization platforms.

For more information, refer to the following links:

<https://lucene.apache.org/core/>

<https://www.elastic.co/products/hadoop>

<http://lucene.apache.org/solr/>

- You can detect novel attacks by adopting Machine-learning approaches on the security data. For example, you can identify new botnet activities in your network by using machine-learning algorithms on large security events data. You can run clustering algorithms such as K-means that Apache Mahout provides to discover attack patterns in your big data. Apache Mahout is a suite of machine-learning libraries designed for big data analysis on Hadoop. You can also send your analysis results as alerts back to ESM by using the SmartConnector for ArcSight Common Event Format Hadoop. Refer to the configuration guide for that SmartConnector. This improves the enterprise's attack-detection capabilities. For more information, refer to the following link:

<http://mahout.apache.org/>

- Run your own, homegrown analytics scripts to process and analyze large security data by using Pig Latin scripting language that Apache Pig platform provides. You can send your analysis results as alerts back to ESM by using the SmartConnector for ArcSight Common Event Format Hadoop. Refer to the configuration guide for that SmartConnector. For more information, refer to the following link:
<https://pig.apache.org/>

Setting Up the Event Data Transfer Tool

The Event Data Transfer Tool is provided with ESM. To set up this tool:

1. Under `/opt/arcsight/logger/current/arcsight/logger/config`, create a folder called `hadoop`.
2. On the Hadoop cluster, find the file `core-site.xml` and copy it into the folder `/opt/arcsight/logger/current/arcsight/logger/config/hadoop`.
3. In the file `core-site.xml`, verify that the property `fs.default.name` is set to the correct DNS/IP address of the master Hadoop node.

Using the Event Data Transfer Tool Command

Use the following `arcsight` command from the current Logger installation path (`/opt/arcsight/logger/current/arcsight/logger/bin`) to initiate the data transfer:

Syntax: `arcsight event_transfer [parameters]`

Required Parameters:

- `-dpath <dpath>`
Specify the path and file name to the destination.
 - If the format is CSV, the extension must be `csv`, `gz`, or `bz2`.
 - If the format is keyvalue of CEF, the extension must be `txt`, `gz`, or `bz2`.
 - When the extension is `gz` or `bz2` the file is compressed.
 - When the extension is `csv` or `txt`, the file is not compressed.
- `-dtype <dtype>`
Specify the type of destination. File means a local path and Hadoop means a path to a Hadoop system.

Optional Parameters:

- `-format <format>`
Specify the format as `cef` (common event format), `csv` (comma-separated values), or

keyvalue (key-value pairs).

- The default is keyvalue.
- Use all lower case letters.
- If you use the CSV format, the file name extension in `-dpath` must be `csv`, `gz`, or `bz2`.
- If you use keyvalue or CEF formats, the file extension must be `txt`, `gz`, or `bz2`.
- `-columns <columns>`
List the CEF column names to include in the transfer. Separate column names with spaces. The default is all columns.
- `-start <start>`
Specify the start of the range of events to transfer as a time (mm/dd/yyyy hh:mm:ss) or by event ID. The default is yesterday at this time (`$NOW-1d`). The time format is for a 24-hour clock. That is, hh is 00 - 24.
- `-end <end>`
Specify the end of the range of events to transfer as a time (mm/dd/yyyy hh:mm:ss) or by event ID. The default is the time specified by `$NOW`. The time format is for a 24-hour clock. That is, hh is 00 - 24.
- `-qtype <qtype>`
Specify the type of entries you used in `-start` and `-end`. For times, the parameters can be `EndTime` or `ManagerReceiptTime` (the default). For event IDs use `EventId`.
- `-sg "<storageGroup>"`
Specify one or more storage groups, in double quotes, and separated by a space. If omitted, events in all storage groups are transferred.
- `-threads <threads>`
Specify the number of threads to use for the transfer. The default is 5. See ["Threads" on the next page](#).
- `--h`
Help

Examples:

```
arcsight event_transfer -dpath <***path***> -dtype Hadoop -sg "storage group 1" "storage group 2"
```

```
arcsight event_transfer -dtype Hadoop -dpath <***path***> -format cef -start "05/04/2016 15:45:00" -end "05/04/2016 16:45:00"
```

Note:

- The `-start`, `-end`, and `-qtype` parameters must be of the same type: either event ID or time. If you mix them up, the tool cannot tell and you get unexpected results.
- The command parameters are case sensitive; use them as shown.

Event Data Transfer Tool Usage Notes

File Names

When the data is transferred to Hadoop, the filename you specify in `-dpath`, such as `abc.gz`, has the start and end appended in front of the extension in the form `abc_<start>_<end>.gz`.

More threads generate more files.

The file extension (`.gz`, in this example), specifies the compression used. If you do not want any compression, use a file extension such as `.txt` or `.csv`.

Threads

The number of threads selected for the transfer affects the rate of transfer in events per second (EPS). Tests have shown that increasing the number of threads increases throughput, but at some point, more threads actually reduce EPS due to resource limitations. The point at which this occurs depends on the number of processors and the amount of other work on that machine. The default of five threads should be satisfactory in most cases.

If you increase to 10 threads you need 6 GB of memory available and another 3 GB for each additional 10 threads. The memory used for this process is controlled by a line in a script used by the `Event_Transfer` command. You can edit the `DirectMemorySize` value in the file, `/opt/arcsight/logger/current/arcsight/logger/config/event_transfer/eventdatatransfertool_config.sh`.

To change the memory used:

1. Comment this line:
`#JVM_HEAP_SIZE="-Xms3g -Xmx6g"`
2. Un-comment this line, which sets direct memory size to 9 GB by default:
`DIRECT_MEMORY_SIZE="-XX:MaxDirectMemorySize=9g"`
3. Optionally change the `DirectMemorySize` value as applicable.

The configuration file itself provides suggestions for size based on threads.

Data Compression

You can specify a Hadoop compression codec by using the file name extension or suffix that corresponds to that codec. Compression occurs before the data is transferred. The compression codecs supported by this transfer tool are:

Suffix	Codec
.bz2	Bzip2Codec
.gz	GzipCodec

The Bzip2Codec appears to have better compression, but the GzipCodec appears to provide a higher EPS value for transfer to the Hadoop system. For more information on these compression codecs, refer to your Hadoop documentation. The `event_transfer` command removes the CORR-Engine compression and then applies the Hadoop compression before, but as part of, the transfer.

If you do not specify a codec-specific file extension, the data is not compressed. Without Hadoop compression, the data in Hadoop is larger than the archive size in the CORREngine. This is because the CORR-Engine data is uncompressed when it is transferred to the Hadoop cluster and the Hadoop file format is larger.

This event migration tool does not transfer Binary Large Object (BLOB) or Character Large Object (CLOB) data.

Transfer Failures

If the transfer fails, you must delete all the data that was transferred in the attempt, before you retry the operation.

The number of files transferred depends on the number of threads used.

File names can be identified by their timestamp.

Transfer Performance

Whether transferring data to Hadoop impacts normal ESM performance depends on how many events you transfer, which event columns you opt to transfer, how often you transfer data, and the number of threads used for data transfer.

However, there is no way to recommend settings that will work in all environments. Try various settings until you settle on the ones that work best for you.

Transferring data to Hadoop is somewhat slower than transferring data to the local machine, but the difference is minimal.

Size of Transferred Files

You may notice a difference in file sizes between Hadoop Distributed File System (HDFS) and Linux local file system. This difference appears only when you use the command `ls -h`. Instead, verify file sizes using the basic `ls` command.

Column Names

The column (field) names assigned in Hadoop are the Common Event Format (CEF) names. For a description of the CEF field names, refer to the document entitled *Implementing ArcSight Common Event Format (CEF)*, which is available on <https://community.softwaregrp.com/t5/ArcSight-Product-Documentation/ct-p/productdocs>.

Appendix D: Creating Custom E-mails Using Velocity Templates

ESM supports the use of *velocity templates* or scripts as defined by The Apache Velocity Project. Velocity templates are a means of specifying dynamic or variable inputs to, or outputs from, underlying Java code.

Velocity templates have many potential applications in ESM. This section describes one such application, E-mail Notification Messages, which you can use Velocity templates on your Manager to create custom e-mail messages to suit your needs.

Note: Velocity templates are an advanced user feature:

- Velocity templates can have wide-ranging effects, so misapplication or inappropriate application is possible. Micro Focus cannot assume responsibility for adverse results caused by user-created Velocity templates.
- ESM does not provide error checking or error messaging for user-created velocity expressions. Refer to the Apache Velocity Project web page at <http://velocity.apache.org/engine/devel/user-guide.html> for information.

Notification Velocity Templates - Example

The <ARCSIGHT_HOME>/Manager/config/notifications directory contains the following two Velocity templates for customizing e-mail notifications:

- `Email.vm`: The primary template file that calls secondary template files.
- `Informative.vm`: The default secondary template file.

Velocity Template `#if` statement

The general format of the `#if` statement for string comparison is:

```
#if ($introspector.getDisplayValue($event, ArcSight_Meta_Tag) Comparative_
Operator Compared_Value)
```

The `#if` statement for integer comparison is:

```
#if ($introspector.getValue($event, ArcSight_Meta_Tag).intValue()Comparative_
Operator Compared_Value)
```


You can specify `ArcSight_Meta_Tag`, `Comparative_Operator`, and `Compared_Value` to suit your needs.

`ArcSight_Meta_Tag` is a string when using the `#if` statement for string comparison (for example, `displayProduct`) and is an integer for the `#if` statement for integer comparison (for example, `severity`).

For a complete listing of ArcSight meta tags, see the `Token Mappings` topic in `ArcSight FlexConnector Guide`.

`Comparative_Operator` is `==` for string comparison; `=`, `>`, and `<` for integer comparison.

`Compared_Value` is a string or an integer. For string comparison, enclose the value in double quotes (`" "`).

Using Email.vm and Informative.vm

It is important to understand the commonly used Velocity programming elements in the `Email.vm` and `Informative.vm` files before editing these files. `Email.vm` calls the secondary template file `Informative.vm` (`#parse ("Informative.vm")`). The `Informative.vm` file lists all the non-empty fields of an event in the format `fieldName : fieldValue`.

The default `Email.vm` template file contents are:

```
## This is a velocity macro file...
## The following fields are defined in the velocity macro.
## event == the event which needs to be sent.
## EVENT_URL == root of the event alert.
#parse ("Informative.vm")
```

This message can be acknowledged in any of the following ways:

- 1) Reply to this email. Make sure that the notification ID listed in this message is present in your reply)
- 2) Login to the ArcSight Console and click on the notification button on the status bar

To view the full alert please go to at `${EVENT_URL}`

The default `Informative.vm` template file contents are:

```
=== Event Details ===
#foreach( $field in $introspector.fields )
```

```
#if( $introspector.getDisplayValue($event, $field).length() > 0 )  
${field.fieldDisplayName}: $introspector.getDisplayValue($event, $field)  
  
#end  
  
#end
```

Understanding the Customization Process

If you want to customize the template files to suit your needs, create new secondary templates containing fields that provide information you want to see in an e-mail for a specific condition.

For example, if you want to see complete details for an event (Threat Details, Source Details, Target Details, and any other information) generated by all Snort devices in your network, create a secondary template file called `Snort.vm` in `<ARCSIGHT_HOME>/config/notification`, on your Manager, with the following lines:

```
=== Complete Event Details ===  
  
Threat Details  
Event: $introspector.getDisplayValue($event,"name")  
Description: $introspector.getDisplayValue($event,"message")  
Severity: $introspector.getDisplayValue($event,"severity")  
-----  
  
Source Details  
Source Address: $introspector.getDisplayValue($event,"attackerAddress")  
Source Host Name: $introspector.getDisplayValue($event,"attackerHostName")  
Source Port: $introspector.getDisplayValue($event,"sourcePort")  
Source User Name: $introspector.getDisplayValue($event,"sourceUserName")  
-----  
  
Target Details  
Target Address: $introspector.getDisplayValue($event,"targetAddress")  
Target Host Name: $introspector.getDisplayValue($event,"targetHostName")  
Target Port: $introspector.getDisplayValue($event,"targetPort")  
Target User Name: $introspector.getDisplayValue($event,"targetUserName")  
-----  
  
Extra Information (where applicable)
```

```
Transport Protocol: $introspector.getDisplayValue($event,"transportProtocol")
```

```
Base Event Count: $introspector.getDisplayValue($event,"baseEventCount")
```

```
Template: /home/arcsight/arcsight/Manager/config/notifications/Snort.vm
```

After you have created the secondary templates, you can edit the `Email.vm` template to insert conditions that call those templates.

As shown in the example below, insert a condition to call `Snort.vm` if the `deviceProduct` in the generated event matches "Snort".

```
#if( $introspector.getDisplayValue($event, "deviceProduct") == "Snort" )
```

```
#parse("Snort.vm")
```

```
#else
```

```
#parse("Informative.vm")
```

```
#end
```

Customizing the Template Files

Follow these steps to customize the `Email.vm` and create any other secondary template files to receive customized e-mail notifications:

1. In `<ARCSIGHT_HOME>/config/notifications`, create a new secondary template file, as shown in the `Snort.vm` example in the previous section.
2. Save the file.
3. Edit `Email.vm` to insert the conditions, as shown in the example in the previous section.
4. Save `Email.vm`.

Velocity Template Sample Output

If you use the `Snort.vm` template and modify `Email.vm` as explained in the previous section, here is the output these templates generate:

```
Notification ID: fInjoQwBABCgMJkA-a8Z-Q== Escalation Level: 1
```

```
=== Complete Event Details ===
```

```
Threat Details
```

```
Event: Internal to External Port Scanning
```

Description: Internal to External Port Scanning Activity Detected;
Investigate Business Need for Activity

Severity: 2

Source Details

Source Address: 10.129.26.37

Source Host Name:

Source Port: 0

Source User Name: jdoe

Target Details

Target Address: 161.58.201.13

Target Host Name:

Target Port: 20090

Target User Name:

Extra Information (where applicable)

Transport Protocol: TCP

Base Event Count: 1

Template: /home/arcsight/arcsight/Manager/config/notifications/Snort.vm

How to Respond

This message can be acknowledged in any of the following ways:

- 1) Reply to this email. Make sure that the notification ID listed in this message is present in your reply)
- 2) Login to the ArcSight Console and click on the notification button on the status bar
- 3) Login to myArcSight and go to the My Notifications Acknowledgment page at <https://mymanager.mycompany.com:9443/arcsight/app?service=page/NotifyHome>

View the full alert at:

<https://mymanager.mycompany.com:9443/arcsight/app?service=page/NotifyHome>

Appendix E: Configuration Changes Related to FIPS

This appendix provides information about and instructions for configuring ESM to support Federal Information Processing Standard (FIPS) 140-2, Suite B, and some other configuration changes you can make while in FIPS mode.

FIPS is a standard published by the National Institute of Standards and Technology (NIST) and is used to accredit cryptographic modules in software components. A cryptographic module is either hardware or software or a combination that is used to implement cryptographic logic. The US federal government requires that all IT products dealing with Sensitive, but Unclassified (SBU) information meet the FIPS standard.

- To be compliant with FIPS 140-2, all components, including Connectors and Logger, if present, must be configured in FIPS mode. Connectors and Logger setup are covered in their documentation.
- For information about supported platforms and specifics about FIPS mode architecture for all ESM products, contact Customer Support.

FIPS Encryption Cipher Suites

A cipher suite is a set of authentication, encryption, and data integrity algorithms used for securely exchanging data between an SSL server and a client. Depending on FIPS mode settings, some of the following specific cipher suites are automatically enabled for ESM and its clients.

FIPS 140-2

- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_128_CBC_SHA

Note: These are the same cipher suites as are used for non-FIPS mode.

FIPS Suite B

In 192 bit mode, the following 192-bit cipher suites are supported.

- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384

In 128 bit mode, the following 128-bit cipher suites are supported.

- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256

Key Pair Types Used in FIPS Mode

For FIPS 140-2, RSA keys with 2,048 bits are used (the same as non-FIPS mode). For FIPS Suite B, Elliptic Curve keys must be used. For 128 bit security, keys with at least 256 bits are required. For 192 bit security, keys with at least 384 bits are required. Note that some browsers will not communicate using keys longer than 384 bits, so 384 bits is a good choice for any Suite B key pair.

The type of key pair for FIPS with Suite B is different. The key depends on the level of classification you need to accommodate. FIPS Suite B requires the use of elliptic curve cryptography. The minimum length of keys is:

- 256: for up to secret classifications corresponding to 128-bit encryption
- 384: for up to top secret classifications corresponding to 192-bit encryption

See "[Generating a Key Pair](#)" on page 104 for details on key pair generation.

Enabling Use of CA-Signed Certificates Over Port 9000 in FIPS Mode

ESM communicates over port 9000 for peering. To allow use of CA-signed certificates over port 9000 in FIPS mode, in addition to completing the procedures that are described in [Using a CA-Signed SSL Certificate](#), you must perform some additional steps.

Note: When completing the procedures in [Using a CA-Signed SSL Certificate](#), follow the instructions for using the `keytool` command.

To enable use of CA-signed certificates over port 9000:

1. Import the certificate reply, using the alias `mykey` (you must use this alias):

```
./arcsight keytool -store managerkeys -importcert -alias mykey -file /home/arcsight/careply.cer
```
2. Copy `/opt/arcsight/logger/userdata/logger/user/logger/keystore.bcfks` to `/opt/arcsight/logger/userdata/logger/user/logger/keystore.bcfks.old`.
This creates a backup copy of the `keystore.bcfks` that you can revert to if there are problems with the certificate import.
3. Import the root certificate (and intermediate certificate, if necessary) to the following

location (this example uses the alias rootCA):

```
keytool -import -trustcacerts -alias rootCA -file ca.cert.pem -keystore
/opt/arcsight/logger/userdata/logger/user/logger/keystore.bcfks -storetype
BCFKS -storepass <keystore password> -providerclass
org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider -providerpath
/opt/arcsight/logger/current/arcsight/logger/lib/modules/org.bouncycastle-
bc-fips-1.0.0.jar
```

Configuring ESM to Use TLS 1.1 or 1.0

By default, ESM uses version 1.2 of the TLS protocol. If you have a need to use TLS 1.1 or 1.0 instead, this topic describes how to configure ESM to do so.

You must configure both ESM and SmartConnectors to use the same TLS version. For non-FIPS mode, the SmartConnector version must be at least 7.5. For FIPS mode, the SmartConnector version must be at least 7.7.

Note: In the following procedures, for any property where you specify TLS versions to support, you can specify a single version or multiple versions. Use a comma-separated list to specify multiple versions. For example:

```
ssl.protocols=TLSv1.1,TLSv1
```

To configure SmartConnectors to use TLS 1.1 or 1.0:

1. If this is a newly-installed connector, modify the properties that are listed below in `<installation_directory>/current/user/agent/default.agent.properties`. If the connector is already running, modify the properties that are listed below in `<installation_directory>/current/user/agent/agent.properties`:
 - `remote.management.ssl.default.protocol=TLS`
 - `remote.management.ssl.enabled.protocols=<TLS versions to support>`
 - `remote.management.ssl.fips.enabled.protocols=<TLS versions to support>`
 - `ssl.protocols=<TLS versions to support>`
2. If this is a newly-installed connector, run the script `runagentsetup.sh` to configure the connector.
3. If this is a newly-installed connector, to ensure that ESM preserves the properties modifications when you upgrade the connector, repeat the modifications in `<installation_directory>/current/user/agent/agent.properties`.

4. If the connector is already running, restart the connector for the changes to take effect.

To configure ESM to use TLS 1.1 or 1.0:

1. On the Manager, add the following entries to `/opt/arcsight/manager/config/server.properties`:
 - `ssl.protocols=<TLS versions to support>`
 - `ssl.protocols.nonfips=SSLv2Hello,<TLS versions to support>`
2. On the Manager, in `/opt/arcsight/logger/current/local/apache/conf/httpd.conf`, modify the line that starts with "SSLProtocol" to read:
`SSLProtocol <TLS versions to support>`

Note: You must repeat this step after an upgrade.

3. Restart the ESM services for the changes to take effect.
4. On the ArcSight Console, add the following entries to `<installation_directory>/current/config/console.properties`:
 - `ssl.protocols=<TLS versions to support>`
 - `ssl.protocols.nonfips=SSLv2Hello,<TLS versions to support>`
5. Restart the console for the changes to take effect.
6. If a Forwarding Connector is installed on the Manager server, complete the procedure for configuring a SmartConnector to use TLS 1.1 or 1.0.

Generating a New Key Pair When Changing a Manager Hostname for FIPS Mode

You perform this set of tasks only if you have changed the Manager hostname.

1. Delete the existing Manager key pair:
If you are generating a key pair on the Manager, first delete the one that is there by default:

```
bin/arcsight keytool -store managerkeys -delete -alias mykey
```
2. Generate a new key pair for the Manager:
For FIPS 140-2:

```
bin/arcsight keytool -store managerkeys -genkeypair -dname "CN=<hostname>" -alias mykey -keyalg rsa -keysize 2048 -validity <in_days>
```

For FIPS Suite B 192:

```
bin/arcsight keytool -store managerkeys -genkeypair -dname "CN=<hostname>"  
-alias mykey -keyalg ec -keysize 384 -validity <in_days>
```

3. Stop the Manager and start all services so the Manager can start using the self-signed certificate. Run the following commands to do so:

```
/etc/init.d/arcsight_services stop manager
```

```
/etc/init.d/arcsight_services start all
```

Always perform these steps after generating a key pair.

4. Stop each connector.
5. Use `agentsetup` to import the new certificates into the Console.
6. Restart each connector
7. Restart the Console.
8. Import the new certificate into the client truststore on the manager. This is necessary so that manager utilities will continue to work. Delete the existing manager certificate from the manager's client truststore with the following command:

```
bin/arcsight keytool -store clientcerts -delete -alias <hostname>
```

then add the new certificate with these commands:

On the Manager:

```
bin/arcsight keytool -store managerkeys -exportcert -alias mykey -file  
mykey.cer
```

On the client (connector or console):

```
bin/arcsight keytool -store clientcerts -importcert -alias <hostname> -  
file mykey.cer
```

Changing a Default Mode Installation to FIPS 140-2

Note that before migrating from default mode to FIPS mode, keep in mind that pre-v4.0 Loggers cannot communicate with a FIPS-enabled Manager. Also note:

- Default to Suite B conversion is not supported in compact mode or in distributed correlation mode.
- FIPS mode conversion (of a non-FIPS system to FIPS 140-2) is supported in compact mode **only**. You cannot convert your system's FIPS mode if your system is in distributed correlation mode.
- Non-FIPS systems in distributed correlation mode cannot be converted to FIPS 140-2.

To convert an existing default mode installation to FIPS mode, on each component, migrate the existing certificates and key pairs from the component's cacerts and keystore to the component's FIPS keystore. The following sub-sections provide you step-by-step instructions on how to do so for each component.

Manager

The tasks below require that you use `keytool`; `keytoolgui` is not supported in FIPS mode.

To convert an existing Manager from default mode to FIPS mode you will export the certificate and import the key pair. Then you will run commands from the Manager's home directory to verify the key pair import and import the certificate.

To convert the Manager from default mode to FIPS 140-2:

1. Log in as user *arcsight*.
2. Stop the Manager if it is running.

```
/etc/init.d/arcsight_services stop manager
```

3. Run `bin/arcsight managersetup`.
 - a. Select **Run Manager in FIPS Mode**.
 - b. Select **FIPS 140-2**.
 - c. Complete `managersetup`.
4. If you have installed SSL client certificates on the manager, this command automatically copies them to the FIPS keystore:

```
bin/arcsight keytool -importkeystore -store managerkeys -srckeystore  
config/jetty/truststore -srcstoretype JKS -srcstorepass <old managercerts  
password> -deststorepass <managerkeys password>
```

Note: The `-srcstorepass` and `-deststorepass` options are not necessary if the `<old managercerts password>` matches the `managerkeys password`. If you have not changed these passwords, both will be *changeit*.

5. Copy the current manager key to the FIPS keystore:

```
bin/arcsight keytool -importkeystore -store managerkeys -srckeystore  
config/jetty/keystore -srcstoretype JKS -alias mykey -srckeypass <old  
managerkeys password> -destkeypass <managerkeys password>
```

If you have not changed the old password the `<old managerkeys password>` will be *password*. When asked if should overwrite existing keystore, select **Yes**.

6. Restart services:

```
/etc/init.d/arcsight_services start
```

ArcSight Console

The tasks below require that you use `keytool`; `keytoolgui` is not supported in FIPS mode.

Follow these steps to convert an existing ArcSight Console from default mode to FIPS mode. Follow these steps after you have converted the Manager to FIPS, as detailed in section ["Manager" on the previous page](#).

To convert the Console from default mode to FIPS 140-2:

1. Stop the ArcSight Console if it is running.
2. Export the certificate and copy it to the Console current directory:

```
bin/arcsight keytool -exportcert -store managerkeys -alias mykey -file manager.cert
```
3. Run the Console setup program by running `bin/arcsight consolesetup`.
 - a. Select **No, I do not want to transfer the settings**.
 - b. Select **Run Console in FIPS Mode**.
 - c. Select **FIPS 140-2**.
 - d. Follow the prompts in the next few screens until the wizard informs you that you have successfully configured the Console.

Note: In the unlikely event you see the message:

Warning: Custom SSL keystore properties for client are detected, manual configuration may be necessary. when running Console setup, check the values in `console/client.properties`. Make sure the value of `ssl.keystore.password` matches that of `ssl.truststore.password`, and that the value of `ssl.keystore.path` matches that of `ssl.truststore.path`. If the paths do not match, change them so they do. If the passwords do not match follow the steps in ["Changing Keystore/Truststore Passwords" on page 107](#) to change passwords. It is much simpler to change the password of the truststore, since the truststore contains no keys.

4. If `config/keystore.client` exists, this indicates that SSL client certificates are in use. Run the following command to migrate them to the FIPS keystore:

```
bin/arcsight keytool -importkeystore -store clientkeys -srckeystore config/keystore.client -srcstoretype JKS
```
5. Remove the old manager certificate if it exists:

```
bin/arcsight keytool -delete -store clientcerts -alias <hostname of manager>
```
6. Import the manager certificate into the Console truststore:

```
bin/arcsight keytool -importcert -store clientcerts -alias <hostname of manager> -file manager.cert
```

Select **Yes** when asked if this certificate should be trusted.

7. Start the console with `bin/arcsight console`. When you start the Console, you should see a message in the `logs/console.log` file telling you that the Console has started in FIPS mode.

Connectors

For information on configuring Connectors for FIPS, refer to *SmartConnector Configuration Guide* for each SmartConnector.

Changing Keystore/Truststore Passwords in FIPS Mode

It is a good security practice to change the keystore and truststore passwords after installing ESM or ESM console. In addition to changing the keystore password, you need to separately change the value that ESM uses for this password, so that ESM can continue to access the keystore. FIPS has a single shared keystore/truststore, so the keystore and truststore passwords must be the same. Changing passwords using `bin/arcsight changepassword` is recommended since this program will encrypt the passwords in the configuration file.

Note: Key pairs also have passwords. ESM expects that these passwords will be the same as the keystore passwords, so both must be changed.

Below is an example of how to change the passwords on the Manager keystore.

Note: These steps must be performed in the order given.

1. `/etc/init.d/arcsight_services stop manager`
2. `bin/arcsight keytool -store managerkeys -keypasswd -alias mykey`
The command `keytool` will prompt for the new password.
3. `bin/arcsight keytool -store managerkeys -storepasswd`
The command `keytool` will prompt for the new password. Enter the same password as for step 2.
4. `bin/arcsight changepassword -f config/esm.properties -p server.privatekey.password`
The command `changepassword` will prompt for the new password. Enter the same

password as for step 2.

5. `/etc/init.d/arcsight_services start all`

Here is an example of how to change the password on a Console truststore to match that of the console keystore. This can be needed to convert a default mode installation (with separate keystore/truststore) to FIPS mode with a single keystore/truststore. The console should not be running. Note that no `keytool -keypasswd` command is needed, as there are no keys in the truststore.

1. `bin/arcsight keytool -store clientcerts -storepasswd`

The command `keytool` will prompt for the new password. Enter the password for the `clientcerts` keystore.

2. `bin/arcsight changepassword -f config/client.properties -p
ssl.truststore.password`

The command `changepassword` will prompt for the new password. Enter the password for the `clientcerts` keystore.

Configure Your Browser for FIPS

To connect a browser to a FIPS web server, the browser must be configured to support FIPS. Review the documentation for your browser and follow the instructions to make it FIPS compliant before using it for ArcSight Console online help or to connect to the ArcSight Command Center.

Make sure that all SSL protocols are turned off. For example, on Microsoft Internet Explorer (IE):

1. Select **Tools > Internet Options**.
2. Select the **Advanced** tab.
3. Scroll down to the **Security** section.
4. Uncheck **Use SSL 2.0** and **Use SSL 3.0**.
5. Check the TLS options. For details on TLS support, see the topic *TLS Support in the ESM Installation Guide*.

Other browsers (and other versions of IE) may have different menu items or options for doing this, so refer to your browser documentation.

When using a browser with Suite B, it matters how you generate your key pair. For information about the encryption to use with browsers, see ["Key Pair Types Used in FIPS Mode" on page 231](#).

Send Documentation Feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this computer, click the link above and an email window opens with the following information in the subject line:

Feedback on Administrator's Guide (ESM 7.2 Service Pack 1)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to arcsight_doc@microfocus.com.

We appreciate your feedback!