
Micro Focus Security ArcSight ESM

Software Version: 7.2

ESM High Availability Module User's Guide

Document Release Date: November 2019

Software Release Date: November 2019



Legal Notices

Copyright Notice

© Copyright 2001-2019 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: Introduction	6
Chapter 2: Configuring the High Availability System - All Scenarios	8
Placing Systems in a High Availability Environment	8
Understanding Network Requirements	9
Using the Service IP Address to Identify the Cluster	10
Using the High Availability Module and ArcSight Transformation Hub	11
Licensing the High Availability Module	11
Chapter 3: Installing the High Availability Module and ESM	12
New Installation: Understanding Hardware Requirements	12
New Installation: Understanding Software Requirements	14
New Installation: Configuring the High Availability System	15
New Installation: Running the High Availability Module Installation Wizard	17
New Installation: Running the First Boot Wizard	18
Installing ESM	20
New Installation: Completing Module Post-Installation Tasks	21
Chapter 4: Installing the High Availability Module on an Appliance	23
Configuring the Appliance	23
Appliance: Running the High Availability Module Installation Wizard	25
Appliance: Running the First Boot Wizard	26
Appliance: Completing Module Post-Installation Tasks	28
Chapter 5: Installing the High Availability Module with an Existing ESM Installation	30
Existing ESM Installation: Understanding Hardware Requirements	30
Planning for the Initial Disk Synchronization	32
Existing ESM Installation: Understanding Software Requirements	32
Existing ESM Installation: Configuring the High Availability System	34
Existing ESM Installation: Running the High Availability Module Installation Wizard	37

Existing ESM Installation: Running the First Boot Wizard	38
Existing ESM Installation: Completing Module Post-Installation Tasks	40
Chapter 6: Upgrading ESM and the High Availability Module	43
Verifying the High Availability Module and ESM Upgrade	45
Chapter 7: Uninstalling Software Components	47
Chapter 8: An Example HA Implementation	49
Server Configuration	49
Initial Setup and Installation	50
Hardware	50
DNS Setup	50
Operating System Installation	50
Disk Partition Setup	51
Interconnect Cable Setup	52
Set Up Connected Hosts	52
Install ArcSight Software	53
Increase Disk Space	53
Chapter 9: Maintain and Monitor the Cluster System	55
The arcsight_cluster Script	55
Command Syntax	55
clusterParameters	56
diagnose	57
increaseDisk	57
license	58
offline	59
online	59
status	59
Status Output Example	59
Status Output Explanation	60
tuneDiskSync	63
Log Output	63
Changing Hostname, IP Address, or Service IP	64
Changing the Cluster's Service IP Address	64
Changing the Secondary Hostname or IP Address only	66
Changing the Primary Hostname or IP Address Only	66

Changing Both Server Hostnames or IP Addresses	67
Changing the Interconnect IP Address	68
Replacing a Server	68
Changing Mount Options	69
Chapter 10: Troubleshooting the Systems	71
Installation Issues and Solutions	71
General Problems	75
Changing ESM to IPv6	75
Audit Events	75
highavailability:100	76
highavailability:200	76
highavailability:300	76
highavailability:500	77
Failover Triggers	77
Processes Killed During Failover	77
System does not Failover	78
Network Interface Commands Stall Disk Mirroring	78
No ESM Uninstall Links on the Primary	78
Stopping the Network on the Secondary Kills ESM	79
Disks on Cluster System Fail to Connect	79
Appendix A: The highavail.properties File	81
Appendix B: Upgrading the Operating System on Software ESM	82
Appendix C: Upgrading the Operating System on a High Availability Appliance	83
Upgrading to RHEL 7.5	83
Upgrading from RHEL 7.5 to RHEL 7.7	86
Send Documentation Feedback	88


Chapter 1: Introduction

The ESM High Availability module (HA module) provides for a backup ESM system with automatic failover capability should the primary ESM system experience communication or operational problems. The HA module is only supported with ESM, and not with other ArcSight products. There are no separate configuration requirements to run the HA module with ESM in FIPS mode rather than in default mode.

You install the HA module on the primary of two adjacent systems that are connected by an Ethernet crossover cable. The HA module replicates the installation and all data by mirroring the hard disk partition to the secondary system.

The two systems each have an individual host IP address that is static. When you install the HA module, you specify a separate service IP address that identifies the cluster. During a failover, the HA module dynamically reassigns the service IP address to the new primary system.

Ordinarily, one ESM instance runs on the primary system and the HA module mirrors selected hard-disk writes to the secondary system. The HA module monitors the health of the primary system. When a failover is triggered, the HA module starts the secondary ESM instance, which takes over. During the failover process, events are cached at the connectors so that data is not lost.



You must complete configuration set up tasks on both the primary and secondary systems before you install the HA module. Some tasks are the same regardless of your installation scenario. For more information about these tasks, see [Configuring the High Availability System - All Scenarios](#). Other tasks are specific to your installation scenario. This guide covers the following scenarios:

- Both systems are new and do not have ESM installed.
For more information, see [Installing the High Availability Module and ESM](#).
- You are installing the HA module on an appliance.
For more information, see [Installing the High Availability Module on an Appliance](#).
- One of the systems has ESM installed.
For more information, see [Installing the High Availability Module with an Existing ESM Installation](#).
- You are upgrading both ESM and the HA module.
For more information, see [Upgrading ESM and the High Availability Module](#).

The configuration steps ensure that both systems are configured properly and that the configuration is aligned across the two systems.

You install or upgrade ESM and the HA module on the primary system only. After installation is complete, the HA module requires time to synchronize the secondary system with the primary system. In general, new ESM installations require less time than upgrading existing ESM systems because of the amount of data to be synchronized.

Chapter 2: Configuring the High Availability System - All Scenarios

This chapter describes networking and other HA system requirements. This information applies to all installation scenarios and will help you prepare for setting up the cluster systems and installing or upgrading the HA module. For information about requirements for installing or upgrading ESM, see the [ESM Installation Guide](#) or the [ESM Upgrade Guide](#).

Important: If you already have ESM and a license for a High Availability solution that was implemented before the HA Module 1.0 release, you will need a new ESM license that supports this product. The new High Availability module uses software to manage failovers and requires a different hardware configuration.

When planning your installation or upgrade, keep the following points in mind:

- The HA module is not supported with SELinux in enforcing mode. If SELinux is installed, it should be in disabled or permissive mode. Micro Focus recommends disabled mode.
- If you are planning to install ESM in distributed correlation mode, note that only the persistor node in the distributed correlation cluster supports the HA module. Non-persistor nodes do not support the HA module.

For an example HA implementation, see [An Example HA Implementation](#)

Placing Systems in a High Availability Environment

The HA module helps to ensure continued availability of ESM at the application level. However, a complete solution requires that high availability be designed at multiple points in a network architecture. This document does not discuss designing a high availability network architecture. However, you can take the following actions to help ensure continued availability of ESM:

- For the primary and secondary systems, provide redundant power supplies for each system from different sources.
- Use application management software to notify you of any issues with the primary or secondary systems.

Understanding Network Requirements

The HA module requires that the network meets the following requirements:

- Set up at least one host on the network that is separate from the cluster systems (called a **connected host**). The HA module pings this host to check for network connectivity. You will specify the host name or IP address of the connected host when you run the First Boot Wizard. Connected hosts can be IPv4, IPv6, or a combination of the two.
- The two HA systems must be part of the same IPv4 or IPv6 subnet. If you change the ESM subnet after you install the HA module (for example, by changing from IPv4 to IPv6), you must uninstall and then reinstall the HA module. This will require approximately 30 minutes of downtime and a resync of the data.
- The primary and secondary machines must be close enough that the cable connection between them requires no intervening routers or switches.
- Ensure that Port 1 (bottom left port) is connected to the network and complete the OS configuration. Then use the second port at the bottom to connect the crossover link.
- Obtain at least five IP addresses for the two systems:
 - Two IP addresses (one per system) are the static host IP addresses used to receive network communication.
 - Two IP addresses (one per system) are used for direct communication between the two systems in the cluster using crossover cables. These can be IPv4 or IPv6 addresses.

Note: You can use private IP addresses if you are certain that ESM will not route communication to these addresses.

- One IP address is the service IP address that is assigned to the ESM cluster. You will specify the host IP addresses and the service IP address when you run the First Boot Wizard. The service IP address is dynamically reassigned when a failover occurs and when the primary is brought back online. The service IP address must be on the same subnet as the host IP addresses. For more information, see [Using the Service IP Address to Identify the Cluster](#).

Note: Micro Focus recommends that all IP addresses be either IPv4 or IPv6. The crossover cable IP addresses can differ in protocol from the other IP addresses, but if they do the cluster management communication between the two hosts will only use network communication ports (not the crossover cable ports) and will not be redundant.

- If you are converting from a single system deployment to a cluster deployment using the HA module, you can save time by using the original ESM IP address as the new

service IP address, and then giving the original ESM system a new IP address. This enables you to reuse the ArcSight Manager SSL certificate, rather than having to generate a new certificate and import it to all connectors and clients.

- Micro Focus recommends using DNS to manage IP addresses and host names for all components in the cluster. Using DNS enables you to manage the service IP address in relation to the numerous connectors, consoles, and command centers associated with a specific ArcSight Manager. Also, using DNS enables you to keep the IP addresses or host names consistent for the primary and secondary systems in your cluster. However, you would not want to use DNS to track the IP addresses for the primary and secondary cables; there is no benefit from using DNS in this case.
- The HA module uses ports 5404, 5405, and 7789 on each IP address in the cluster environment. These ports must be dedicated to HA module communication. Do not configure other applications to use these ports.
- Both systems use the ports and protocols that are listed below. Ensure that `firewalld` and `iptables` do not block these ports. Set up your network firewalls to allow access to the connected hosts. A connected host is any other device on the network that the HA module can ping to verify that it is still on the network.

Protocol	Outgoing communication from...	On Port	Incoming communication to...	On Port
ICMP	<ul style="list-style-type: none">◦ Primary IP address◦ Secondary IP address	N/A	<ul style="list-style-type: none">◦ Primary IP address◦ Secondary IP address◦ Service IP address◦ Connected host	N/A
TCP	<ul style="list-style-type: none">◦ Primary crossover cable◦ Secondary crossover cable	Any	<ul style="list-style-type: none">◦ Primary system cable◦ Secondary system cable	7789
UDP	<ul style="list-style-type: none">◦ Primary IP address◦ Primary crossover cable◦ Secondary IP address◦ Secondary crossover cable	Any	<ul style="list-style-type: none">◦ Primary IP address◦ Primary crossover cable◦ Secondary IP address◦ Secondary crossover cable	Any

Using the Service IP Address to Identify the Cluster

The service IP address is an important element of the cluster systems. The HA module uses the service IP address for communication across the network. When you configure the ArcSight Manager IP address or host name during ESM installation, you will specify the service IP address and not an individual host IP address. The ArcSight Manager host name should resolve to the service IP address.

The ArcSight Console uses the service IP address to connect to the ArcSight Manager. Also, the ArcSight Command Center URL will specify the service IP address. When a failover occurs, the service IP address is dynamically assigned to the new primary system. Other than specifying the service IP address when installing the HA module and ESM and ensuring that no other hosts use this IP address, you will not need to configure it further. The HA module automatically configures it on the same interface that the host IP addresses use.

Note: Micro Focus recommends that you configure a host name during ESM installation. Host name changes are easier to manage using DNS and are required for IPv6 systems.

Using the High Availability Module and ArcSight Transformation Hub

In an HA ESM environment, if there is a mismatch between the actual ESM host name or IP address and the host name or IP address that is listed under the ESM consumer group on the Transformation Hub Manager, the cause might be that the underlying third party library in the Transformation Hub prefers the ESM primary or secondary host name or IP address instead of the service IP address.

This does not impact ESM and Transformation Hub functionality. For more information, see the [Transformation Hub Administrator's Guide](#).

Licensing the High Availability Module

The license file for the HA module is an ESM license file with the HA module included.

If you have ESM installed without the HA module, obtain a new ESM license that includes the module. After you upgrade ESM, install the new ESM/HA module license as described in the [ESM Administrator's Guide](#).

If ESM is not already installed, specify the ESM/HA module license file when you install the HA module and then again when you install ESM. For more information about installing ESM, see the [ESM Installation Guide](#).

If you are upgrading from ESM 7.0 Patch 1 and HA module 7.0 Patch 1, you do not need a new license file. If you are running a version of ESM and the HA module that is earlier than 7.0 Patch 1, you must upgrade to 7.0 Patch 1 before you can upgrade to version 7.2.

Chapter 3: Installing the High Availability Module and ESM

This chapter describes how to configure your systems and then run the High Availability module installation wizard and First Boot Wizard. This chapter applies to the scenario where you are installing ESM and the High Availability module for the first time.

Before you begin the installation, review [Configuring the High Availability System - All Scenarios](#).

Note: Do not install ESM until the HA disk synchronization is complete. Attempting to install ESM while HA disk synchronization is in process can cause the ESM installation to fail.

New Installation: Understanding Hardware Requirements

The HA module requires two identical systems that conform to the latest ESM version hardware and software requirements, except where described in this document. The HA module is not supported on virtual machines.

Ensure that your environment meets the following hardware requirements:

- You must use identical server class systems that support running either RHEL, CentOS, or SUSE Linux Enterprise Server (SLES).
- Running ESM with the HA module requires significant disk space. Because of the synchronization process, the cluster systems must meet minimum storage requirements. The ESM and archival storage must be on the same shared disk.

For information about hard disk requirements for running ESM, see the [ESM Installation Guide](#). In addition to the ESM requirements, the HA module has additional storage requirements:

Purpose	Minimum storage	Note
ESM and HA module installation binaries	3 GB	Ensure that there is enough space for the downloaded installation binaries.

Temporary installation files	6 GB	Space required to run the installation wizard and the First Boot Wizard
Shared disk partition	Varies	The HA module mirrors this partition between the two systems. The volume size depends on the specific implementation needs. ESM requires approximately 10 TB (mid-range) - 12 TB (high performance) disk space for event storage, plus at least one 1 GB (with no upper limit) of event archive space.
HA synchronization metadata	Varies	The volume size depends on the size of the ESM online storage.

- Set up the following disk partitions on the primary and secondary systems.

Because the installation erases the contents of the shared disk on the secondary system, ensure that it does not contain data of value.

Ensure that processes on the primary and secondary systems do not use the shared disk file system.

The shared disk partition does not support bind mounts. The installation wizard flags them as errors. Use symbolic links instead.

Partitions	Space required	Location	Notes
Shared disk partition		Either /opt or /opt/arcsight	This is a recommendation. Alternatively, you can create an /opt or /opt/arcsight symbolic link to the physical location. The HA module mirrors this partition between the two systems.
Metadata partition	Determined by the size of the shared disk partition	/dev/<sub_path>	Contains disk synchronization metadata This volume location must start with /dev. The metadata partition size is calculated as: size (in mebibytes)= (P/32) + 1 P = shared disk partition size in gibibytes
/ (root) partition	20 GiB (generous)		An operating system recommended partition
swap	8 GiB (minimum)		An operating system recommended partition.
temp	10 GiB (or more)	/tmp	An operating system recommended partition

- If the shared disks have write caches enabled, the write caches must be battery backed write caches (BBWC). If the shared disks do not have battery backup, there is a chance that the disks will be out-of-sync if a power failure occurs.
- The network interface cards should be at 1 Gigabit (Gb) or higher and use a cable that supports this bandwidth.

- The network interface that is used for the interconnection of the two servers should run at 1 or 10 Gigabits (Gb)/sec. The benefit of the higher bandwidth is seen during the initial synchronization between the primary and secondary systems. This is useful when you upgrade ESM on the primary system and there is a significant amount of data to synchronize.
- If your servers have high-speed disk subsystems, you might see improved performance with a 10 Gb network interface. The mirrored disk performance is limited by the slower of either the disk write throughput or the throughput on the crossover link.

New Installation: Understanding Software Requirements

The HA module has the following software requirements:

- The HA module version, ESM version, and operating system version must be compatible. For version compatibility, see the [ESM Support Matrix](#).
- The cluster systems must run either RHEL, CentOS, or SUSE Linux Enterprise Server (SLES). Both systems must have the same operating system and version installed.

Caution: The HA module incorporates components that are operating system version-specific. If you upgrade to a version of the operating system that is not supported, the module might not work properly. Do not upgrade to a newer version of your operating system until there is a version of the HA module that supports it.

If avahi software is running on your system, the HA module might not work properly. If you are running RHEL or CentOS 7.6 or SLES, run the following command as user root to detect whether avahi is present:

```
systemctl status avahi-daemon
```

If you are running RHEL or CentOS 6.10, run the following command as user root:

```
service avahi-daemon status
```

If the response indicates that the avahi-daemon is enabled or running, see the documentation for your operating system for instructions to stop it and to disable it upon reboot.

- The file system for the mirrored disk partitions can be EXT4 or XFS. You cannot change the file system type while installing the HA module or during an ESM upgrade. Both systems must use the same file system type.
- Configure both systems to access a yum (for RHEL or CentOS) or yast (for SLES) repository in order to install dependencies that the HA module requires. The

repository can be a remote yum repository that the operating system vendor provides, a repository that you create from the operating system ISO or CD, or a directory location on the local system. See the vendor-specific documentation for information about configuring yum or yast and connecting to yum or yast repositories.

- Micro Focus recommends that you use the operating system's Logical Volume Management (LVM) tools to manage volumes and partitions on the HA cluster systems. These tools make the process of configuring and managing disk space much simpler than using physical disk management.

An LVM partition must be a multiple of the LVM chunk size. If you use 32 MiB for the chunk size and need a 33 MiB partition, create a 64 MiB partition (because you would need two chunks). For an example, see [Disk Partition Setup](#).

New Installation: Configuring the High Availability System

You must configure the primary and secondary systems so that they are identical. Complete the following steps as directed on the primary system, the secondary system, or both systems. The HA module installation scripts verify the configuration and generate error messages if dependencies are not met.

To configure the HA system:

1. From the [Licensing and Downloads](#) site, download ArcSightESMSuite-7.2.0.xxxx.tar and ArcSight-ActivePassiveHighAvailability-7.2.0.xxxx.tar to the primary system.
2. Extract both .tar files. Do NOT extract them into what will later be the shared directory (generally /opt), because files in that directory are deleted during installation.
3. Ensure that both systems are using the correct version of the operating system timezone package. This is a requirement for ESM. For more information, see the [ESM Installation Guide](#).
4. To synchronize the time between the primary and secondary systems, configure them to run the Network Time Protocol (NTP).
5. Connect the systems with crossover cables and configure the interfaces with the appropriate IPv4 or IPv6 addresses. Both systems must use the same IP version. Ping from one system to the other over the configured interfaces to ensure proper configuration.
6. On the primary and secondary systems, select the partitions to be mirrored.

Typically, this is the partition mounted as `/opt` for your ESM installation. This partition must exist on the primary and secondary systems and must have the same device name, mounted location, and size. If the partition is not mounted as `/opt` or `/opt/arcsight`, create a symbolic link to `/opt` or `/opt/arcsight` on both systems.

Note: After installation, this partition is only mounted on the primary system. Only the primary system can make changes to it.

7. If the mirrored disks are SSD drives, such as Fusion, configure TRIM support on the primary and secondary systems.
8. Ensure that all file system options are set up as desired on the primary system. The HA module mounts the file system on the secondary system exactly as you mounted it on the primary system.
9. On the primary and secondary systems, create a metadata partition. This is a small partition on each system that is used for disk-synchronization metadata. The size to allocate for each partition is calculated in mebibytes:

$$\text{size (in mebibytes)} = (P/32) + 1$$

where P is the size of the shared disk partition in gibibytes. For example, if the shared disk partition size is 1 TiB (that is, 1,024 GiB), the metadata partition size would be 33 MiB.

For an example, see [Disk Partition Setup](#). If you increase the size of the shared disk partition, also increase the size of the metadata partition accordingly. Decreasing the size of the mounted partition is not supported.

If the metadata partition will be a physical volume (for example, `/dev/sda8`), create it now. If the metadata partition will be a logical volume (for example, `/dev/mapper/vg00-meta`), then you only need to ensure that enough free disk space is available in a volume group. The `prepareHA.sh` script will create the metadata volume.

10. Ensure that the password for the `root` user is the same on both systems. You can change the password after installation.
11. As user `root`, run the following commands:

```
cp -r Tools /tmp
cd /tmp/Tools/highavail
cp template.properties highavail.properties
chmod 644 highavail.properties
```
12. Specify the following information in `/tmp/Tools/highavail/highavail.properties`:
 - `service_hostname`= host name of ESM in the HA cluster
 - `shared_disk`= mount point of the disk to be mirrored across both systems

- `metadata_volume`= volume name for the metadata volume (for example, `/dev/mapper/vg00-meta`)
 - `primary_cable_ip`= IP address of the interface to the cable that is connected to the secondary system
 - `primary_hostname`= host name of the primary system
 - `secondary_cable_ip`= IP address of the interface to the cable that is connected to the primary system
 - `secondary_hostname`= host name of the secondary system
13. As user `root`, run the following command on the primary system:
- ```
/tmp/Tools/highavail/prepareHA.sh
```
- a. Confirm the names of the primary and secondary systems.
  - b. If the metadata partition does not exist and it will be a logical volume, allow the script to create it.
  - c. Provide the password for user `arcsight`.
- If there are errors, correct them and run `prepareHA.sh` again.
14. As user `root`, run the following command:
- ```
scp -r /tmp/Tools <secondary hostname>:/tmp
```
15. Reboot the primary system.
16. On the secondary system, as user `root`, run the following command:
- ```
/tmp/Tools/highavail/prepareHA.sh
```
- If there are errors, correct them and run `prepareHA.sh` again.
17. Reboot the secondary system.
- When the configuration tasks for the primary and secondary systems are complete, continue to [New Installation: Running the High Availability Module Installation Wizard](#).

## New Installation: Running the High Availability Module Installation Wizard

After you complete the tasks in [New Installation: Configuring the High Availability System](#), you can run the HA installation wizard in either console mode (via the command line) or GUI mode (using X Windows).

If you install the HA module in GUI mode, the First Boot Wizard starts automatically when the HA installation wizard is complete. You can also run the First Boot Wizard independently at any time to change the HA module configuration.

**Note:** Run the HA installation wizard on the primary system only.

### To run the installation wizard:

1. As user `arcsight`, change directory to the location where you extracted the `.tar` files and run the installation wizard:

```
./ArcSight-ActivePassiveHighAvailability-7.2.0.xxxx.bin -i console for
console mode
```

or

```
./ArcSight-ActivePassiveHighAvailability-7.2.0.xxxx.bin for GUI mode
```

2. Accept the license agreement.

The installation wizard generates a pre-installation summary and reports the installation progress.

If you ran the installation wizard in console mode, the wizard prompts you to start the First Boot Wizard when the installation is complete.

If you ran the installation wizard in GUI mode, the First Boot Wizard starts automatically when the installation is complete.

For information about running the First Boot Wizard, see [New Installation: Running the First Boot Wizard](#).

## New Installation: Running the First Boot Wizard

It is important that the primary and secondary systems match with respect to hardware, software, and configuration. The First Boot Wizard examines relevant characteristics in detail. Messages about inconsistencies are relatively common, especially the first time, and the messages should supply enough information to correct the inconsistency, run the First Boot Wizard again, and complete the installation.

**Note:** Run the First Boot Wizard on the primary system only.

### To run the First Boot Wizard:

1. If the First Boot Wizard did not start automatically after you completed the HA installation wizard, run the following command as user `arcsight`:

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

2. Specify the path to the ArcSight license file.
3. When prompted, load the `highavail.properties` file that defines the cluster

configuration (the file that you created at  
/tmp/Tools/highavail/highavail.properties).

4. Specify host names and the following information:

| Field              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Shared Disk        | <p>Mount point of the disk that is shared between the primary and secondary systems</p> <p>The options provided include all relevant mount points.</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>shared_disk</code>.</p> <p>The installation does not support bind mounts and flags them as errors. Use symbolic links instead.</p> <p>Because the installation completely erases the contents of the shared disk on the secondary system, ensure that it does not contain data of value.</p> <p>Ensure that no processes on the primary or secondary systems are using this file system. Otherwise, the installation will exit with errors.</p> <p>You cannot change this value on subsequent runs of the First Boot Wizard.</p> |
| Metadata Volume    | <p>Volume that contains disk-synchronization metadata</p> <p>The volume is expected to start with <code>/dev</code>.</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>metadata_volume</code>.</p> <p>The contents of the metadata volume on both the primary and the secondary systems will be removed.</p>                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Service Hostname   | <p>Service host name that is assigned to the service IP address</p> <p>This is a virtual host name that is used to connect to the cluster regardless of which physical computer is acting as the primary system. You can also use the service IP address, but Micro Focus recommends using the service host name. You can use the <code>hosts</code> file or DNS to resolve the host names.</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>service_hostname</code>.</p>                                                                                                                                                                                                                                                            |
| Secondary Hostname | <p>Host name of the secondary system</p> <p>This is the host name that is assigned specifically to the secondary system.</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>secondary_hostname</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Primary Cable IP   | <p>IP address of the interface that is connected to the interconnect cable on the primary system</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>primary_cable_ip</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Secondary Cable IP | <p>IP address of the interface that is connected to the interconnect cable on the secondary system</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>secondary_cable_ip</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

5. At the **Parameter Configuration** prompt, specify the following information:

| Field           | Description                                                                                                                                                                                                                                                                                                                                    |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Connected Hosts | <p>These hosts are other machines in the network that the HA module can ping to verify that it is connected to the network.</p> <p>Enter a space-separated list of host names or IP addresses that the module can ping. Do not specify a host name or IP address for the primary and secondary systems.</p> <p>This field is not required.</p> |
| Ping Timeout    | <p>Seconds to wait before reporting that a ping failed</p> <p>The default is 2 seconds.</p>                                                                                                                                                                                                                                                    |
| Ping Attempts   | <p>Number of pings to attempt before reporting that the pings failed</p> <p>The default is 2 pings.</p>                                                                                                                                                                                                                                        |

The wizard generates a summary of the host names and other configuration parameters. The wizard resolves IP addresses to host names and resolves host names to IP addresses. The wizard determines whether to use IPv4 or IPv6 for the service IP address and provides an explanation. If you do not agree with the choice, you might be able to force the wizard to choose IPv4 or IPv6 by specifying an IPv4 or IPv6 address instead of a host name.

6. Provide the password for user `root`.

The password enables the HA configuration script to complete actions that must be performed as the `root` user. The password must be the same on the primary and secondary systems. The wizard does not permanently store the password. You can change the password after the installation is complete.

After you select to continue, the wizard displays the status of each operation. It might take approximately one hour to complete.

7. When the installation is complete, check the log files on both servers. For information about resolving errors, see [Installation Issues and Solutions](#).

After you resolve errors, as user `arcsight`, run the First Boot Wizard again:

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

When the First Boot Wizard is complete, continue to [Installing ESM](#).

## Installing ESM

After you complete the First Boot Wizard, you can install ESM.

**Note:** Before you install ESM, ensure that the HA module is running.

## To install ESM:

1. As user root, create the folder /opt/arcsight and set ownership to user arcsight:

```
chown arcsight:arcsight /opt/arcsight
```

If the mount point for mirroring is /opt or /opt/arcsight, the HA module mirrors the change to the secondary system.

2. On the primary system, install ESM. For more information, see the [ESM Installation Guide](#).

**Note:** When the ESM Configuration Wizard prompts you for the ArcSight Manager host name or IP address, specify the cluster service host name or service IP address and **not** the host name of a single server.

Because you already installed the HA module, you do not need to run prepare\_system.sh during ESM installation.

3. During installation, ensure that you include the ArcSight ESM HA Monitoring Foundation Package.

The package is required in order to receive up-to-date HA module status information.

When the ESM installation is complete, perform post-installation tasks as described in the [ESM Installation Guide](#), and then continue to [New Installation: Completing Module Post-Installation Tasks](#).

## New Installation: Completing Module Post-Installation Tasks

After you complete post-installation tasks as described in the [ESM Installation Guide](#), verify the HA module and ESM installation.

### To complete post-installation tasks:

1. Ensure that the primary and secondary systems are using the correct version of the operating system timezone package.

For more information, see the [ESM Installation Guide](#). If you need to install the timezone package, you must install it on the primary and secondary systems because it is not installed in the shared directory.

2. On the primary system, ensure that all ArcSight services are running:

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

3. Check the cluster status:

```
./arcsight_cluster status
```

The `arcsight_cluster` script is located in `/usr/lib/arcsight/highavail/bin`. For more information about the script, see [The `arcsight\_cluster` Script](#).

If you installed the HA module on RHEL 7.6 and later want to upgrade the operating system to RHEL 7.7, see [Upgrading the Operating System on Software ESM](#).

# Chapter 4: Installing the High Availability Module on an Appliance

This chapter describes how to configure the appliance and then run the ESM High Availability module installation wizard and First Boot Wizard.

Before you begin the installation, review [Configuring the High Availability System - All Scenarios](#).

## Configuring the Appliance

The HA module requires two appliances that are configured so that they are identical. Complete the following steps as directed on the primary system, the secondary system, or both systems. The HA module installation scripts check the configuration and generate error messages if the appliances do not meet dependencies.

1. From the [Licensing and Downloads](#) site, download `ArcSight-ActivePassiveHighAvailability-7.2.0.xxxx.tar` to the primary system.
2. If you are migrating an appliance from a single installation to a high availability cluster, also download `esm_apha_support_rpms_rhel77_G10.tar.gz` to the primary system.
3. Extract the `.tar` files. Do NOT extract them into what will later be the shared directory (generally `/opt`), because files in that directory are deleted during installation.
4. New appliances include setup scripts that run at first boot to install the operating system and ESM. Do not install ESM on the secondary system. The scripts can run on the primary appliance, but must not run on the secondary appliance. Complete the following steps to stop the scripts after the operating system is installed and prevent ESM from being installed in the future:
  - a. Allow the script to install ESM on the primary appliance.
  - b. When you boot the secondary appliance, allow the operating system installation script to run.
  - c. At the prompt "Install Anywhere will guide you through the installation of ArcSight ESM 7.2 Suite," enter `quit`.
  - d. To prevent the script from installing ESM on the secondary appliance after you reboot it, edit the `.bash_profile` in the root user's home directory on the primary and the secondary appliances. Remove the line `# run OS configuration and ESM installer if not yet done` and all of the lines that follow it.

- e. To ensure that the ESM installation script does not run, reboot the primary and secondary appliances.
5. Ensure that both systems are using the correct version of the operating system timezone package. This is a requirement for ESM. For more information, see the [ESM Installation Guide](#).
6. To synchronize the time between the primary and secondary systems, configure them to run the Network Time Protocol (NTP).
7. Connect the systems with crossover cables and configure the interfaces with the appropriate IPv4 or IPv6 addresses. Both systems must use the same IP version. Ping from one system to the other over the configured interfaces to ensure proper configuration.
8. As user root, run the following command on the primary system to stop ESM:  
`/opt/arcsight/manager/bin/remove_services.sh`
9. On the primary and secondary systems, install support packages that are required to run the HA module:
  - a. Login as user root.
  - b. Copy `esm_apha_support_rpms_rhel77_G10.tar.gz` file to the `/tmp` partition.
  - c. Run the following commands:  

```
cd /tmp
tar xzf esm_apha_support_rpms_rhel77_G10.tar.gz
cd install
./install_ha_support_pkgs.sh
```
10. Ensure that the password for the root user is the same on both systems. You can change the password after installation.
11. As user root, run the following commands on the secondary system:  

```
mv /opt/amsd /usr/local
mv /opt/arcsight /usr/local
mv /opt/smartstorageadmin /usr/local
ln -s /usr/local/amsd /opt
ln -s /usr/local/arcsight /opt
ln -s /usr/local/smartstorageadmin /opt
umount /opt
ln -s /usr/local/amsd /opt
ln -s /usr/local/arcsight /opt
ln -s /usr/local/smartstorageadmin /opt
mount /opt
```
12. If you want to re-use the ArcSight Manager SSL certificate rather than generate a new certificate, complete the following steps:



- a. Add a new IP address to the interface that has the current host IP address. The new IP address will become the new the host IP address. The original IP address will become the service IP address that identifies the cluster.
- b. Set up `/etc/hosts` or DNS to resolve the new host IP address to the new host name.
- c. Configure the host to use the new host name.

Now that you have removed the original IP address from the network interface, you can re-use it as the cluster service IP Address when you run the First Boot Wizard.

**Note:** If you change the system host name during installation, ensure that the change persists across reboots. Reboot the system, and then use the `hostname` command to verify the system host name.

13. On the primary and secondary systems, run the following commands as user root:

```
cd /usr/lib
mkdir arcsight
chown arcsight:arcsight /usr/lib/arcsight
```

When the configuration tasks for the primary and secondary systems are complete, continue to [Appliance: Running the High Availability Module Installation Wizard](#).

## Appliance: Running the High Availability Module Installation Wizard

After you complete the tasks in [Configuring the Appliance](#), you can run the HA installation wizard in either console mode (via the command line) or GUI mode (using X Windows).

If you install the HA module in GUI mode, the First Boot Wizard starts automatically when the HA installation wizard is complete. You can also run the First Boot Wizard independently at any time to change the HA module configuration.

**Note:** Run the HA installation wizard on the primary system only.

### To run the installation wizard:

1. As user `arcsight`, change directory to the location where you extracted the `.tar` files and run the installation wizard:

```
./ArcSight-ActivePassiveHighAvailability-7.2.0.xxxx.bin -i console for
console mode
```

or

```
./ArcSight-ActivePassiveHighAvailability-7.2.0.xxxx.bin for GUI mode
```

2. Accept the license agreement.

The installation wizard generates a pre-installation summary and reports the installation progress.

If you ran the installation wizard in console mode, the wizard prompts you to start the First Boot Wizard when the installation is complete.

If you ran the installation wizard in GUI mode, the First Boot Wizard starts automatically when the installation is complete.

For information about running the First Boot Wizard, see [Appliance: Running the First Boot Wizard](#).

## Appliance: Running the First Boot Wizard

It is important that the primary and secondary systems match with respect to hardware, software, and configuration. The First Boot Wizard examines relevant characteristics in detail. Messages about inconsistencies are relatively common, especially the first time, and the messages should supply enough information to correct the inconsistency, run the First Boot Wizard again, and complete the installation.

**Note:** Run the First Boot Wizard on the primary system only.

### To run the First Boot Wizard:

1. If the First Boot Wizard did not start automatically after you completed the HA installation wizard, run the following command as user `arcsight`:  
`/usr/lib/arcsight/highavail/bin/arcsight firstBootWizard --console`
2. Specify the path to the ArcSight license file.
3. When prompted, load the `highavail.properties` file that defines the cluster configuration (the file that you created at `/tmp/Tools/highavail/highavail.properties`).

4. Specify host names and the following information:

| Field                                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Shared Disk                                      | <p>Mount point of the disk that is shared between the primary and secondary systems</p> <p>The options provided include all relevant mount points.</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>shared_disk</code>.</p> <p>Because the installation completely erases the contents of the shared disk on the secondary system, ensure that it does not contain data of value.</p> <p>Ensure that no processes on the primary or secondary systems are using this file system. Otherwise, the installation will exit with errors.</p> <p>You cannot change this value on subsequent runs of the First Boot Wizard.</p> |
| Metadata Volume for ESM Express or ESM Appliance | <p>Volume that contains disk-synchronization metadata</p> <p>The metadata partition for ESM Express or ESM Appliance is <code>/dev/sda6</code>.</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>metadata_volume</code>.</p> <p>The contents of the metadata volume on both the primary and the secondary systems will be removed.</p>                                                                                                                                                                                                                                                                                    |
| Service Hostname                                 | <p>Service host name that is assigned to the service IP address</p> <p>This is a virtual host name that is used to connect to the cluster regardless of which physical computer is acting as the primary system. You can also use the service IP address, but Micro Focus recommends using the service host name. You can use the <code>hosts</code> file or DNS to resolve the host names.</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>service_hostname</code>.</p>                                                                                                                                                 |
| Secondary Hostname                               | <p>Host name of the secondary system</p> <p>This is the host name that is assigned specifically to the secondary system.</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>secondary_hostname</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Primary Cable IP                                 | <p>IP address of the interface that is connected to the interconnect cable on the primary system</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>primary_cable_ip</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Secondary Cable IP                               | <p>IP address of the interface that is connected to the interconnect cable on the secondary system</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>secondary_cable_ip</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                        |

5. At the **Parameter Configuration** prompt, specify the following information:

| Field           | Description                                                                                                                                                                                                                                                                                                                                    |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Connected Hosts | <p>These hosts are other machines in the network that the HA module can ping to verify that it is connected to the network.</p> <p>Enter a space-separated list of host names or IP addresses that the module can ping. Do not specify a host name or IP address for the primary and secondary systems.</p> <p>This field is not required.</p> |
| Ping Timeout    | <p>Seconds to wait before reporting that a ping failed</p> <p>The default is 2 seconds.</p>                                                                                                                                                                                                                                                    |
| Ping Attempts   | <p>Number of pings to attempt before reporting that the pings failed</p> <p>The default is 2 pings.</p>                                                                                                                                                                                                                                        |

The wizard generates a summary of the host names and other configuration parameters. The wizard resolves IP addresses to host names and resolves host names to IP addresses. The wizard determines whether to use IPv4 or IPv6 for the service IP address and provides an explanation. If you do not agree with the choice, you might be able to force the wizard to choose IPv4 or IPv6 by specifying an IPv4 or IPv6 address instead of a host name.

6. Provide the password for user `root`.

The password enables the HA configuration script to complete actions that must be performed as the `root` user. The password must be the same on the primary and secondary systems. The wizard does not permanently store the password. You can change the password after the installation is complete.

After you select to continue, the wizard displays the status of each operation. It might take approximately one hour to complete.

7. When the installation is complete, check the log files on both servers. For information about resolving errors, see [Installation Issues and Solutions](#).

After you resolve errors, as user `arcsight`, run the First Boot Wizard again:

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

When the installation is complete, perform post-installation tasks as described in the [ESM Installation Guide](#), and then continue to [Appliance: Completing Module Post-Installation Tasks](#).

## Appliance: Completing Module Post-Installation Tasks

After you complete post-installation tasks as described in the [ESM Installation Guide](#), verify the HA module and ESM installation on the appliance.

### To complete post-installation tasks:

1. Ensure that the primary and secondary systems are using the correct version of the operating system timezone package.

For more information, see the [ESM Installation Guide](#). If you need to install the timezone package, you must install it on the primary and secondary systems because it is not installed in the shared directory.

2. On the primary system, ensure that all ArcSight services are running:

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

3. Check the cluster status:

```
./arcsight_cluster status
```

The `arcsight_cluster` script is located in `/usr/lib/arcsight/highavail/bin`. For more information about the script, see [The arcsight\\_cluster Script](#).

# Chapter 5: Installing the High Availability Module with an Existing ESM Installation

This chapter describes how to configure your systems and then run the ESM High Availability module installation wizard and First Boot Wizard. This chapter applies to the scenario where you are installing the HA module with an existing ESM installation.

Before you begin the installation, review [Configuring the High Availability System - All Scenarios](#).

## Existing ESM Installation: Understanding Hardware Requirements

The HA module requires two identical systems that conform to the latest ESM version hardware and software requirements, except where described in this document. The HA module is not supported on virtual machines.

Ensure that your environment meets the following hardware requirements:

- You must use identical server class systems that support running either RHEL, CentOS, or SUSE Linux Enterprise Server (SLES).
- Running ESM with the HA module requires significant disk space. Because of the synchronization process, the cluster systems must meet minimum storage requirements. The ESM and archival storage must be on the same shared disk.

For information about hard disk requirements for running ESM, see the [ESM Installation Guide](#). In addition to the ESM requirements, the HA module has additional storage requirements:

| Purpose                                 | Minimum storage | Note                                                                        |
|-----------------------------------------|-----------------|-----------------------------------------------------------------------------|
| ESM and HA module installation binaries | 3 GB            | Ensure that there is enough space for the downloaded installation binaries. |

|                              |        |                                                                                                                                                                                                                                                                                                        |
|------------------------------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Temporary installation files | 6 GB   | Space required to run the installation wizard and the First Boot Wizard                                                                                                                                                                                                                                |
| Shared disk partition        | Varies | The HA module mirrors this partition between the two systems. The volume size depends on the specific implementation needs. ESM requires approximately 10 TB (mid-range) - 12 TB (high performance) disk space for event storage, plus at least one 1 GB (with no upper limit) of event archive space. |
| HA synchronization metadata  | Varies | The volume size depends on the size of the ESM online storage.                                                                                                                                                                                                                                         |

- Set up the following disk partitions on the primary and secondary systems.

Because the installation erases the contents of the shared disk on the secondary system, ensure that it does not contain data of value.

Ensure that processes on the primary and secondary systems do not use the shared disk file system.

The shared disk partition does not support bind mounts. The installation wizard flags them as errors. Use symbolic links instead.

| Partitions            | Space required                                      | Location                              | Notes                                                                                                                                                                                                                   |
|-----------------------|-----------------------------------------------------|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Shared disk partition |                                                     | Either<br>/opt<br>or<br>/opt/arcsight | This is a recommendation. Alternatively, you can create an /opt or /opt/arcsight symbolic link to the physical location. The HA module mirrors this partition between the two systems.                                  |
| Metadata partition    | Determined by the size of the shared disk partition | /dev/<sub_path>                       | Contains disk synchronization metadata<br>This volume location must start with /dev.<br>The metadata partition size is calculated as:<br>size (in mebibytes)= (P/32) + 1<br>P = shared disk partition size in gibibytes |
| / (root) partition    | 20 GiB (generous)                                   |                                       | An operating system recommended partition                                                                                                                                                                               |
| swap                  | 8 GiB (minimum)                                     |                                       | An operating system recommended partition.                                                                                                                                                                              |
| temp                  | 10 GiB (or more)                                    | /tmp                                  | An operating system recommended partition                                                                                                                                                                               |

- If the shared disks have write caches enabled, the write caches must be battery backed write caches (BBWC). If the shared disks do not have battery backup, there is a chance that the disks will be out-of-sync if a power failure occurs.
- The network interface cards should be at 1 Gigabit (Gb) or higher and use a cable that supports this bandwidth.

- The network interface that is used for the interconnection of the two servers should run at 1 or 10 Gigabits (Gb)/sec. The benefit of the higher bandwidth is seen during the initial synchronization between the primary and secondary systems. This is useful when you upgrade ESM on the primary system and there is a significant amount of data to synchronize. For more information, see [Planning for the Initial Disk Synchronization](#).
- If your servers have high-speed disk subsystems, you might see improved performance with a 10 Gb network interface. The mirrored disk performance is limited by the slower of either the disk write throughput or the throughput on the crossover link.

## Planning for the Initial Disk Synchronization

After you install the HA module on an existing ESM system, the module must synchronize the entire shared disk partition on the existing ESM primary system to the secondary system. Depending on the amount of data to be synchronized, the speed of the network interface card, and the disk I/O rates, it could take two or more days to complete the synchronization.

The synchronization speed is determined by the slower of the disk I/O rate and the data transfer rate across the cable. You can run ESM on the primary system during this time, but the secondary system will not be ready to take over until the synchronization is complete. Typical ESM installations use very fast server class disks, which can be much faster than a 1 Gb cable. In such cases, providing a 10 Gb interface might noticeably reduce the time required for the initial synchronization.

SSD drives (for example, Fusion) contribute to improving the synchronization speed. SSD drives require and support TRIM to manage free space. The disk synchronization process can use TRIM to identify free blocks on the drive and skip them during synchronization. For example, if you have 12 TB of SSD storage, 4 TB of which are used, and if you run the Linux `fstrim` command immediately after installing the HA module, then the disk synchronization process passes TRIM information to the SSD drives. The disk synchronization process uses this information to detect which blocks are free and skips these blocks. In this example, only 4 TB of data would need to be synchronized, instead of 12.

## Existing ESM Installation: Understanding Software Requirements

The following software requirements apply when adding the HA module to an existing ESM installation:



- The HA module version, ESM version, and operating system version must be compatible. For version compatibility, see the [ESM Support Matrix](#).
- Replace the original ESM license with a new license that enables both ESM and the HA module.
- During the module installation, ensure that you install the ArcSight ESM HA Monitoring Foundation Package.
- The cluster systems must run either RHEL, CentOS, or SUSE Linux Enterprise Server (SLES). Both systems must have the same operating system and version installed.

**Caution:** The HA module incorporates components that are operating system version-specific. If you upgrade to a version of the operating system that is not supported, the module might not work properly. Do not upgrade to a newer version of your operating system until there is a version of the HA module that supports it.

If avahi software is running on your system, the HA module might not work properly. If you are running RHEL or CentOS 7.6 or SLES, run the following command as user root to detect whether avahi is present:

```
systemctl status avahi-daemon
```

If you are running RHEL or CentOS 6.10, run the following command as user root:

```
service avahi-daemon status
```

If the response indicates that the avahi-daemon is enabled or running, see the documentation for your operating system for instructions to stop it and to disable it upon reboot.

- If you are adding the HA module to an existing ESM installation that is earlier than ESM 7.0, upgrade to ESM 7.0 *before* you install the module.
- If you plan to convert the system from IPv4 to IPv6, convert *after* you upgrade to ESM 7.0 and *before* you install the module.
- The mirrored disk mount point (for example, /opt) must be the same on the primary and secondary systems. The mounted volume name (for example, /dev/sda5 or /dev/mapper/vg00-opt) must also be the same on both systems. If the primary system uses a physical volume for the mirrored disk, then the secondary system must also use a physical volume.
- Configure both systems to access a yum repository in order to install dependencies that the HA module requires. The repository can be a remote yum repository that the operating system vendor provides, a repository that you create from the operating system ISO or CD, or a directory location on the local system. See the vendor-specific documentation for information about configuring yum and connecting to yum repositories.

- Create a secondary system that has the same volumes as the primary system. If the primary system uses physical volume management, the secondary system must also use physical volume management.

## Existing ESM Installation: Configuring the High Availability System

You must configure the primary and secondary systems so that they are identical. Complete the following steps as directed on the primary system, the secondary system, or both systems. The HA module installation scripts verify the configuration and generate error messages if dependencies are not met.

### To configure the HA system:

1. From the [Licensing and Downloads](#) site, download `ArcSightESMSuite-7.2.0.xxxx.tar` and `ArcSight-ActivePassiveHighAvailability-7.2.0.xxxx.tar` to the primary system.
2. Extract both `.tar` files. Do NOT extract them into what will later be the shared directory (generally `/opt`), because files in that directory are deleted during installation.
3. On the primary system, run the following command as user `root` to stop ESM:  
`/opt/arcsight/manager/bin/remove_services.sh`

**Note:** If ESM was running in distributed correlation mode, you must run `sshSetup` again. When you run `remove_services`, it removes the configuration of `sshSetup` after you install the HA module and run `setupServices`. For more information, see the [ESM Administrator's Guide](#).

4. If you want to re-use the ArcSight Manager SSL certificate rather than generate a new certificate, complete the following steps:
  - a. Add a new IP address to the interface that has the current host IP address. The new IP address will become the new host IP address. The original IP address will become the service IP address that identifies the cluster.
  - b. Set up `/etc/hosts` or DNS to resolve the new host IP address to the new host name.
  - c. Configure the host to use the new host name.

Now that you have removed the original IP address from the network interface, you can re-use it as the cluster service IP Address when you run the First Boot Wizard.

**Note:** If you change the system host name during installation, ensure that the change persists across reboots. Reboot the system, and then use the `hostname` command to verify the system host name.

5. Ensure that both systems are using the correct version of the operating system timezone package. This is a requirement for ESM. For more information, see the [ESM Installation Guide](#).
6. To synchronize the time between the primary and secondary systems, configure them to run the Network Time Protocol (NTP).
7. Connect the systems with crossover cables and configure the interfaces with the appropriate IPv4 or IPv6 addresses. Both systems must use the same IP version. Ping from one system to the other over the configured interfaces to ensure proper configuration.
8. On the primary and secondary systems, select the partitions to be mirrored:
  - a. On the primary system, run `df /opt/arcsight`.  
The mount point for `/opt/arcsight` is displayed in the **Mounted on** column.
  - b. On the secondary system, create and mount a volume with the same volume name, size, and file system type.  
If `/opt/arcsight` is a symbolic link on the primary system, the installation wizard will create the same symbolic link on the secondary system.
9. If the mirrored disks are SSD drives, such as Fusion, configure TRIM support on the primary and secondary systems.
10. Ensure that all file system options are set up as desired on the primary system. The HA module mounts the file system on the secondary system exactly as you mounted it on the primary system.
11. On the primary and secondary systems, create a metadata partition. This is a small partition on each system that is used for disk-synchronization metadata. The size to allocate for each partition is calculated in mebibytes:

$$\text{size (in mebibytes)} = (P/32)+1$$

where P is the size of the shared disk partition in gibibytes. For example, if the shared disk partition size is 1 TiB (that is, 1,024 GiB), the metadata partition size would be 33 MiB.

For an example, see [Disk Partition Setup](#). If you increase the size of the shared disk partition, also increase the size of the metadata partition accordingly. Decreasing the size of the mounted partition is not supported.

If the metadata partition will be a physical volume (for example, `/dev/sda8`), create it now. If the metadata partition will be a logical volume (for example, `/dev/mapper/vg00-meta`), then you only need to ensure that enough free disk space

is available in a volume group. The `prepareHA.sh` script will create the metadata volume.

12. Ensure that the password for the `root` user is the same on both systems. You can change the password after installation.
13. As user `root`, run the following commands:

```
cp -r Tools /tmp
cd /tmp/Tools/highavail
cp template.properties highavail.properties
chmod 644 highavail.properties
```
14. Specify the following information in `/tmp/Tools/highavail/highavail.properties`:
  - `service_hostname` = host name of ESM in the HA cluster
  - `shared_disk` = mount point of the disk to be mirrored across both systems
  - `metadata_volume` = volume name for the metadata volume (for example, `/dev/mapper/vg00-meta`)
  - `primary_cable_ip` = IP address of the interface to the cable that is connected to the secondary system
  - `primary_hostname` = host name of the primary system
  - `secondary_cable_ip` = IP address of the interface to the cable that is connected to the primary system
  - `secondary_hostname` = host name of the secondary system
15. As user `root`, run the following command on the primary system:

```
/tmp/Tools/highavail/prepareHA.sh
```

  - a. Confirm the names of the primary and secondary systems.
  - b. If the metadata partition does not exist and it will be a logical volume, allow the script to create it.
  - c. Provide the password for user `arcsight`.

If there are errors, correct them and run `prepareHA.sh` again.
16. As user `root`, run the following command:

```
scp -r /tmp/Tools <secondary hostname>:/tmp
```
17. Reboot the primary system.
18. On the secondary system, as user `root`, run the following command:

```
/tmp/Tools/highavail/prepareHA.sh
```

If there are errors, correct them and run `prepareHA.sh` again.
19. Reboot the secondary system.

When the configuration tasks for the primary and secondary systems are complete, continue to [Existing ESM Installation: Running the High Availability Module Installation Wizard](#).

## Existing ESM Installation: Running the High Availability Module Installation Wizard

After you complete the tasks in [Existing ESM Installation: Configuring the High Availability System](#), you can run the HA installation wizard in either console mode (via the command line) or GUI mode (using X Windows).

If you install the HA module in GUI mode, the First Boot Wizard starts automatically when the HA installation wizard is complete. You can also run the First Boot Wizard independently at any time to change the HA module configuration.

**Note:** Run the HA installation wizard on the primary system only.

### To run the installation wizard:

1. As user `arcsight`, change directory to the location where you extracted the `.tar` files and run the installation wizard:

```
./ArcSight-ActivePassiveHighAvailability-7.2.0.xxxx.bin -i console for
console mode
```

or

```
./ArcSight-ActivePassiveHighAvailability-7.2.0.xxxx.bin for GUI mode
```

2. Accept the license agreement.

The installation wizard generates a pre-installation summary and reports the installation progress.

If you ran the installation wizard in console mode, the wizard prompts you to start the First Boot Wizard when the installation is complete.

If you ran the installation wizard in GUI mode, the First Boot Wizard starts automatically when the installation is complete.

For information about running the First Boot Wizard, see [Existing ESM Installation: Running the First Boot Wizard](#).

## Existing ESM Installation: Running the First Boot Wizard

It is important that the primary and secondary systems match with respect to hardware, software, and configuration. The First Boot Wizard examines relevant characteristics in detail. Messages about inconsistencies are relatively common, especially the first time, and the messages should supply enough information to correct the inconsistency, run the First Boot Wizard again, and complete the installation.

**Note:** Run the First Boot Wizard on the primary system only.

### To run the First Boot Wizard:

1. If the First Boot Wizard did not start automatically after you completed the HA installation wizard, run the following command as user `arcsight`:  
`/usr/lib/arcsight/highavail/bin/arcsight firstBootWizard --console`
2. Specify the path to the ArcSight license file.
3. When prompted, load the `highavail.properties` file that defines the cluster configuration (the file that you created at `/tmp/Tools/highavail/highavail.properties`).
4. Specify host names and the following information:

| Field           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Shared Disk     | <p>Mount point of the disk that is shared between the primary and secondary systems</p> <p>The options provided include all relevant mount points.</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>shared_disk</code>.</p> <p>The installation does not support bind mounts and flags them as errors. Use symbolic links instead.</p> <p>Because the installation completely erases the contents of the shared disk on the secondary system, ensure that it does not contain data of value.</p> <p>Ensure that no processes on the primary or secondary systems are using this file system. Otherwise, the installation will exit with errors.</p> <p>You cannot change this value on subsequent runs of the First Boot Wizard.</p> |
| Metadata Volume | <p>Volume that contains disk-synchronization metadata</p> <p>The volume is expected to start with <code>/dev</code>.</p> <p>In the <code>highavail.properties</code> file, this value is identified as <code>metadata_volume</code>.</p> <p>The contents of the metadata volume on both the primary and the secondary systems will be removed.</p>                                                                                                                                                                                                                                                                                                                                                                                                                          |

| Field              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Service Hostname   | Service host name that is assigned to the service IP address<br>This is a virtual host name that is used to connect to the cluster regardless of which physical computer is acting as the primary system. You can also use the service IP address, but Micro Focus recommends using the service host name. You can use the <code>hosts</code> file or DNS to resolve the host names.<br>In the <code>highavail.properties</code> file, this value is identified as <code>service_hostname</code> . |
| Secondary Hostname | Host name of the secondary system<br>This is the host name that is assigned specifically to the secondary system.<br>In the <code>highavail.properties</code> file, this value is identified as <code>secondary_hostname</code> .                                                                                                                                                                                                                                                                  |
| Primary Cable IP   | IP address of the interface that is connected to the interconnect cable on the primary system<br>In the <code>highavail.properties</code> file, this value is identified as <code>primary_cable_ip</code> .                                                                                                                                                                                                                                                                                        |
| Secondary Cable IP | IP address of the interface that is connected to the interconnect cable on the secondary system<br>In the <code>highavail.properties</code> file, this value is identified as <code>secondary_cable_ip</code> .                                                                                                                                                                                                                                                                                    |

5. At the **Parameter Configuration** prompt, specify the following information:

| Field           | Description                                                                                                                                                                                                                                                                                                                     |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Connected Hosts | These hosts are other machines in the network that the HA module can ping to verify that it is connected to the network.<br>Enter a space-separated list of host names or IP addresses that the module can ping. Do not specify a host name or IP address for the primary and secondary systems.<br>This field is not required. |
| Ping Timeout    | Seconds to wait before reporting that a ping failed<br>The default is 2 seconds.                                                                                                                                                                                                                                                |
| Ping Attempts   | Number of pings to attempt before reporting that the pings failed<br>The default is 2 pings.                                                                                                                                                                                                                                    |

The wizard generates a summary of the host names and other configuration parameters. The wizard resolves IP addresses to host names and resolves host names to IP addresses. The wizard determines whether to use IPv4 or IPv6 for the service IP address and provides an explanation. If you do not agree with the choice, you might be able to force the wizard to choose IPv4 or IPv6 by specifying an IPv4 or IPv6 address instead of a host name.

6. Provide the password for user `root`.

The password enables the HA configuration script to complete actions that must be performed as the `root` user. The password must be the same on the primary and secondary systems. The wizard does not permanently store the password. You can change the password after the installation is complete.

After you select to continue, the wizard displays the status of each operation. It might take approximately one hour to complete.

7. When the installation is complete, check the log files on both servers. For information about resolving errors, see [Installation Issues and Solutions](#).

After you resolve errors, as user `arcsight`, run the First Boot Wizard again:

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

When the installation is complete, perform post-installation tasks as described in the [ESM Installation Guide](#), and then continue to [Existing ESM Installation: Completing Module Post-Installation Tasks](#).

## Existing ESM Installation: Completing Module Post-Installation Tasks

In this scenario, the ESM instance is running on a single system and you converted the installation to an HA module cluster. You must now switch to the new service host name or service IP address.

### To complete post-installation tasks:

1. On the primary system, as user `root`, run the following command to set up the ESM services:

```
/opt/arcsight/manager/bin/setup_services.sh
```

The script automatically detects the HA module and makes appropriate changes to the primary and the secondary systems.

2. If the shared disk is a solid state drive (SSD), run the following command:

```
fstrim <shared disk>
```

On the primary system, if the drive has a large amount of free disk space, the command shortens the time required to synchronize the secondary disk.

**Note:** You can skip steps 3-10 if you changed the original single system hostname and are now using the original IP as the Service IP for the cluster. You can also skip steps 3-10 if your ESM installation uses the hostname for the SSL certificate.

3. As user `arcsight`, run the following command to stop the ArcSight Manager:

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

4. From `/opt/arcsight/manager/bin`, run the following command as user `arcsight` to start the setup program for the ArcSight Manager:



```
./arcsight managersetup
```

5. When prompted for the ArcSight Manager host name, and in every field where the previous host name or IP address is displayed, specify the cluster service host name or cluster service IP address (specify the same value that you set in the First Boot Wizard).
6. When prompted, select the self-signed keypair option and enter the required information to generate the self-signed certificate with the cluster service IP address.

**Note:** If ESM is configured for FIPS mode, you must complete this step manually from the command line. For more information, see the [ESM Administrator's Guide](#).

7. As user arcsight, run the following command to start the ArcSight Manager:

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

8. As user arcsight, ensure that the ArcSight Manager is running:

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

Run this command about once per minute until you receive a notification that the Manager is available.

9. Start the ArcSight Command Center:

```
https://<Service Hostname>:8443/
```

where <Service Hostname> is the host name that is defined for the cluster

If you are running Internet Explorer, host names with underscores do not work, so use the service IP address.

If you are not using DNS to resolve host names, use the service IP address instead.

10. Change the ArcSight Manager IP address to the cluster service IP address for every connector and console that connects to this Manager.
11. Update any URLs (for example, bookmarks) to ArcSight Command Center.
12. Import the newly-generated certificate for the ArcSight Manager to all clients, consoles, and connectors that access the Manager.

You can use `keytoolgui` to import the certificate. For more information, see the [ESM Administrator's Guide](#).

If ESM is configured to use FIPS, use the `arcsight keytool` utility. For more information, see the [ESM Administrator's Guide](#).

13. Ensure that clients can connect to the ArcSight Manager using the service IP address or service host name, and ensure that peer configuration works as expected.

The ESM installation is only mounted and visible on the primary system. To run ESM utilities and commands (for example, `/opt/arcsight/manager/bin/arcsight`), do so from the server that is currently the primary system.

14. If you have not already activated the ArcSight ESM HA Monitoring Foundation Package, activate it from the ArcSight Console.

For more information about activating standard content, see the [ESM Standard Content Guide](#).

15. Ensure that the primary and secondary systems are using the correct version of the operating system timezone package.

For more information, see the [ESM Installation Guide](#). If you need to install the timezone package, you must install it on the primary and secondary systems because it is not installed in the shared directory.

If you installed the HA module on RHEL 7.6 and later want to upgrade the operating system to RHEL 7.7, see [Upgrading the Operating System on Software ESM](#).

# Chapter 6: Upgrading ESM and the High Availability Module

This information guides you through the process of upgrading both ESM and the HA module running on a two-system cluster. The HA module version, ESM version, and operating system version must be compatible. For version compatibility, see the [ESM Support Matrix](#).

The following upgrade paths are supported:

- If you are running version 7.0 of ESM, you must upgrade ESM to version 7.0 Patch 2 before you can upgrade to version 7.2. If you are running version 7.0 of the HA module, you can upgrade the module directly to version 7.2.
- If you are running version 7.0 Patch 1 or Patch 2 of ESM, you can upgrade ESM directly to version 7.2. If you are running version 7.0 Patch 1 of the HA module, you can upgrade the module directly to version 7.2.

**Note:** The upgrade will cause ESM downtime for the duration of the upgrade.

Starting with version 7.2, the HA module no longer uses port 694. Instead, it uses ports 5404 and 5405. Ensure that these ports are available.

Starting with version 7.2, the network cables and the interconnect cables must be running either IPv4 or IPv6 in order to use them as redundant communication channels. To ensure communication redundancy, Micro Focus recommends moving the interconnect cable to the same communication protocol as the network ports. You can do this before or after the upgrade.

Upgrade ESM and the HA module on the primary system only. After the upgrade is complete, the HA module synchronizes the secondary system with the primary system.

Before you begin the upgrade, review [Configuring the High Availability System - All Scenarios](#).

## To upgrade ESM and the HA module:

1. From the [Licensing and Downloads](#) site, download ArcSight-ActivePassiveHighAvailability-7.2.0.xxxx.tar and ArcSightESMSuite-7.2.0.xxxx.tar to the primary system (where xxxx is the build number).

Do NOT place the installation binary or unpacked content on the shared disk partition (usually /opt/arcsight). The upgrade process deletes the shared disk partition.

2. As user `arcsight`, untar `ArcSight-ActivePassiveHighAvailability-7.2.0.xxxx.tar`.
3. Copy the `preUpgrade.sh` file to the secondary system.
4. As user `root`, run `preUpgrade.sh` on the secondary system.
5. If an operating system upgrade is required, upgrade the operating system version on the secondary system.

If this is a software installation, see the operating system vendor documentation for upgrade instructions.

For an appliance installation, see [Upgrading the Operating System on a High Availability Appliance](#). If you upgrade the operating system, download the HA support packages for that operating system and install them.

6. If you upgraded the operating system on software ESM, reboot the secondary system.

**Note:** This is not necessary on an appliance. The appliance will automatically reboot.

7. As user `arcsight`, untar `ArcSightESMSuite-7.2.0.xxxx.tar`.
8. On the primary system, as user `root`, run `Tools/stop_services.sh` to shut down ESM.
9. As user `root`, run `preUpgrade.sh` on the primary system.
10. If an operating system upgrade is required, upgrade the operating system version on the primary system.

If this is a software installation, see the operating system vendor documentation for upgrade instructions.

For an appliance installation, see [Upgrading the Operating System on a High Availability Appliance](#). If you upgrade the operating system, download the HA support packages for that operating system and install them.

11. If you upgraded the operating system on software ESM, reboot the primary system.

**Note:** This is not necessary on an appliance. The appliance will automatically reboot.

12. If you have not already done so, disable SELinux and then reboot the primary and secondary systems.
13. On the primary system, as user `arcsight`, run `ArcSight-ActivePassiveHighAvailability-7.2.0.xxxx.tar` to start the HA Module Installation Wizard.
14. On the primary system, as user `root`, run the following command:  
`/usr/lib/arcsight/highavail/install/upgrade.sh`

The log file for the HA module upgrade is located at: `/usr/lib/arcsight/highavail/logs/upgrade.log`.

15. On the primary system, upgrade to the supported ESM version.

For detailed instructions, see the [ESM Upgrade Guide](#). Because you have already stopped the ArcSight services, you do not need to run `Tools/stop_services.sh`.

**IMPORTANT:** The HA module must be running before you begin upgrading ESM.

After the ESM upgrade is complete, the HA module synchronizes the primary system and the secondary system.

16. As user root, start the ArcSight services:

```
/opt/arcsight/manager/bin/setup_services.sh
```

17. Ensure that the ArcSight services are running:

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

18. If you have not already done so, use the ArcSight Console to activate the ArcSight ESM HA Monitoring Foundation Package.

For more information, see the [ESM Standard Content Guide](#).

When the upgrade is complete, perform post-upgrade tasks as described in the [ESM Upgrade Guide](#), and then continue to [Verifying the High Availability Module and ESM Upgrade](#).

## Verifying the High Availability Module and ESM Upgrade

After you complete post-upgrade tasks as described in the [ESM Upgrade Guide](#), verify the HA module and ESM upgrade.

### To verify the upgrade:

1. Ensure that the primary and secondary systems are using the correct version of the operating system timezone package.

For more information, see the [ESM Installation Guide](#). If you need to install the timezone package, you must install it on the primary and secondary systems because it is not installed in the shared directory.

2. On the primary system, ensure that all ArcSight services are running:

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

3. Check the cluster status:

```
./arcsight_cluster status
```

The `arcsight_cluster` script is located in `/usr/lib/arcsight/highavail/bin`. For more information about the script, see [The `arcsight\_cluster` Script](#).

# Chapter 7: Uninstalling Software Components

You can uninstall both ESM and the HA module or you can uninstall only the HA module.

When you uninstall only the HA module, the systems are no longer part of a cluster installation. After you uninstall the module, you will convert one of the systems to a single ESM installation. When you reconfigure the server, you can choose from the following options:

- Use the server's individual IP address and host name.
- Use the service IP address and host name.

If you use the server's individual IP address or host name to identify the ArcSight Manager instance, you must also change the Manager host name or IP address that is defined for every connector and console instance that connects to this ArcSight Manager. You must also update all bookmarks or URL references to ArcSight Command Center. If you reuse the service IP address and host name, you do not have to update the connected clients.

**Note:** For greater flexibility in configuration, Micro Focus recommends using a host name (rather than an IP Address).

## To uninstall ESM and the HA module:

1. On the primary system, uninstall ESM. For more information, see the [ESM Installation Guide](#).
2. On the primary system, run the HA module uninstallation script as user root:

```
/usr/lib/arcsight/highavail/install/uninstall.sh
```

After you confirm that you want to uninstall the HA module, the script uninstalls it on both systems.

## To uninstall only the HA module:

1. On the primary system, as user root, run the following command:  

```
/opt/arcsight/manager/bin/remove_services.sh
```
2. On the primary system, as user root, run the following command:

```
/usr/lib/arcsight/highavail/install/uninstall.sh
```

After the module uninstallation is complete, all of the files required to run ESM remain on both systems. Choose which server you want to convert to a single ESM installation.

3. If you are not reusing the service IP address, change the IP address. For information about changing the IP address of an ESM server, see the [ESM Installation Guide](#).
4. If you are reusing the service IP address, complete the following steps:

- a. As user root, run the following command to update the IP address configuration on the selected server:

```
ip addr add <service_ip> dev <primary interface>
```

Where <service\_ip> is the IP address and <primary interface> is the interface on which the IP address of the host name is configured (for example, eth0).

- b. Update the ARP cache:

```
arping -U -I <primary interface> -s <service_ip> <default_gateway_ip>
```

- c. Complete the uninstallation on the secondary system that you removed from the HA cluster.

- d. If you uninstalled the HA module in a distributed correlation environment where the persistor was part of the HA cluster, run the following command:

```
/etc/init.d/arcsight_services sshSetup
```

**Note:** This is only necessary in a distributed correlation environment.

- e. On the primary system, run the following command as user root:

```
/opt/arcsight/manager/bin/setup_services.sh
```

At this point ESM is running on the server. However, if you reboot this server, the service IP address will not be available on the primary interface and ESM will not be accessible.

- f. To ensure that the service IP address is available after reboot, modify the appropriate scripts in /etc/sysconfig/network-scripts/.



# Chapter 8: An Example HA Implementation

This chapter describes an example implementation of HA, giving some details which are not provided in the main document. These examples should clarify and make specific the general statements in the main document.

- [Server Configuration](#) how the systems in this example are configured.
- [Initial Setup and Installation](#) goes through the steps required to set up this system.
- [Increase Disk Space](#) shows how to increase the disk space available to ESM in a HA configuration.

## Server Configuration

Each server in this example cluster meets the recommended hardware requirements specified in the *ESM Installation Guide*.

- 2 TiB of RAID 10 storage is provided via 15K RPM disks.
- The network interface runs at 1 GB.
- One 1 GB interface on each server will be interconnected by a cable.
- RedHat 7.6 is used with ESM 7.2 software with the HA Module.
- The company's internal DNS server is used for name-to-address translation for the cluster. This is generally the best choice, because there can be thousands of connectors, and dozens of ESM clients. Changing the ESM hostnames on this many machines would be difficult.
- Linux configuration files are used to define the hostname, the IP addresses for each interface, DNS server addresses, and the default route. In a corporate environment, a more common choice would be to set these values via DHCP. For the purposes of this example it is convenient to configure these on the machine directly, so what is going on can be seen. In any case, it is likely that the interconnect ports would be statically defined, since they connect to each other, and do not have access to a DHCP server.
- The shared disk partition and the metadata partition are allocated space via the Logical Volume Manager (LVM). This is strongly recommended that you use Logical Volume Manager (LVM) tools to manage disk space. It will be much easier for you to increase the disk space later using LVM tools.

# Initial Setup and Installation

## Hardware

A new rack was placed in a server room, and wired for two independent power sources. Two servers with the following characteristics were placed in the rack:

- Two CPUs (16 cores)
- 64G RAM
- One NIC card supporting 4 1Gb Ethernet interfaces
- Eight 600GB 15000 RPM hard drives
- Redundant power supplies

On each server, eth1 (port 2) is connected to the other server by a 1G cable. On each server, eth0 is connected to the network switch (and the internet).

## DNS Setup

We will assume that the company puts its intranet on Net 10 - in the private IP space. Many companies would use public IPs for their intranet - this is a company decision. Here are some example values that we will use:

| Type      | Hostname                       | IP          |
|-----------|--------------------------------|-------------|
| Primary   | ha1.internal.<yourcompany>.com | 10.10.10.2  |
| Secondary | ha2.internal.<yourcompany>.com | 10.10.10.3  |
| Service   | esm.internal.<yourcompany>.com | 10.10.10.10 |

Clients of ESM will connect to esm.internal.<yourcompany>.com. The primary and secondary hostname are required for configuration of those servers, and are convenient for accessing them.

## Operating System Installation

The RedHat installation supports formatting of hard drives, including formatting multiple hard drives to a RAID partition. So first format all the drives into a single RAID 10 disk array. After accounting for redundant storage support this leaves the system with 2.4TB = 2.2TiB.

The root (/), swap, and boot partitions should be physical partitions allocated during installation. Allocate 20 GiB (generous) for root, 8 GiB (minimum) for swap, and 2 GiB

for boot. The remaining disk space can be put into a single LVM volume group (vg00) for later allocation to support ESM.

Give the primary and secondary machines the hostnames specified in the previous section, and configure the IP address of the primary and secondary on the eth0 interface of the respective servers.

## Disk Partition Setup

It is a good idea to configure a separate /tmp partition - in this case a 6GiB partition in ext4 format. You can easily create such a partition from the existing volume group by running the following commands as user *root*:

```
lvcreate -L 6G -n tmp vg00
mkfs -t ext4 /dev/mapper/vg00-tmp
```

Then add the following line to /etc/fstab to make the mount survive across reboots:

```
/dev/mapper/vg00-tmp /tmp ext4 defaults 1 2
```

To mount the /tmp partition, run:

```
mount /tmp
```

Next, set up a partition for /opt that is as large as possible. However, it is necessary to save a little space for the metadata partition required for HA installation. Assuming that the disk will be 2.2 TiB (2,306,867 MiB), then the metadata partition must be at least 72 MiB, where:

$$\text{size} = (2,306,867 \text{ MiB} / 32768) + 1$$

Assuming the chunk size of the volume group is 32 MiB, we need to allocate 96 MiB.

Create this partition with the following command:

```
lvcreate -L 96M -n metadata vg00
```

There is no need to make a file system or mount in this case.

You can make a partition big enough to fill the volume group by running these commands as user *root*:

```
lvcreate -l 100%FREE -n opt vg00
mkfs -t xfs /dev/mapper/vg00-opt
```

Then, as with /tmp, you add an entry to /etc/fstab and mount /opt with the command `mount /opt`. The fstab entry is:

```
/dev/mapper/vg00-lv_opt /opt xfs defaults,inode64 1 2
```

Note that we use the `inode64` option here. That is a good idea for very large file systems - but probably this filesystem is large enough to benefit. In any case, if you have any

special mount options you want, mount your filesystem with them if you want them to be used after the HA installation.

## Interconnect Cable Setup

This section shows how to configure the interconnected interfaces. The eth1 interface on each machine will be connected with a crossover cable. Pick IP addresses for the interconnect interfaces. A private subnet that is not routed to other nodes is a good choice. In this example, we will use subnet 192.168.10.0/24. Address 192.168.10.2 will be the primary IP and 192.168.10.3 will be the secondary IP.

To set this up, first modify the interface scripts ifcfg-eth1 on both machines. This file is in /etc/sysconfig/network-scripts. An example of an ifcfg-eth1 script after the configuration changes:

```
DEVICE=eth1
HWADDR=12:34:56:78:90:AB
UUID=3835e99d-2ef2-422b-9455-75697e092689
IPADDR=192.168.10.2
NETMASK=255.255.255.0
TYPE=Ethernet
ONBOOT=yes
BOOTPROTO=none
USERCTL=no
NM_CONTROLLED=no
IPV6ADDR=fdfd::1:2/120
```

The first three lines come from the original file that was created when the operating system was installed. Delete any other lines from the original file. The next line, defining the IP address, is unique to each machine. On the secondary, we will use the IP Address 192.168.10.3. The remaining lines are the same for all such files - you may copy them in.

To bring up the connection, run ifup eth1 as *root* on both the primary and the secondary. At this point pings to 192.168.10.3 on the primary and pings to 192.168.10.2 on the secondary should succeed.

## Set Up Connected Hosts

In this case, we will set up the network to allow pings to hosts on three different subnets of the intranet - 10.10.11.5, 10.10.12.5, and 10.10.13.5 .

## Install ArcSight Software

This is a new installation, so it is faster to install the HA Module before ESM. After the installations described below are complete, then ESM will be running in HA mode.

### Install HA Module

HA Module is installed on `ha1.internal.acme.com` . Here are the parameters to use to install HA:

| Parameter          | Value                                           |
|--------------------|-------------------------------------------------|
| Shared Disk        | <code>/opt</code>                               |
| Metadata volume    | <code>/dev/mapper/vg00-metadata</code>          |
| Service hostname   | <code>esm.internal.&lt;mycompany&gt;.com</code> |
| Secondary hostname | <code>ha2.internal.&lt;mycompany&gt;.com</code> |
| Primary cable IP   | <code>192.168.10.2</code>                       |
| Secondary cable IP | <code>192.168.10.3</code>                       |
| Connected hosts    | <code>10.10.11.5 10.10.12.5 10.10.13.5</code>   |
| Ping timeout       | <code>2</code>                                  |
| Ping attempts      | <code>2</code>                                  |

### Install ESM

ESM is installed as described in either the ESM Installation Guide. The only special step is when you are prompted for Manager Information. One value will be entered differently than if you are setting up a single ESM system.

Manager host name (or IP): The correct value to enter for **Manager host name (or IP)** is `esm.internal.<mycompany>.com`.

Administrator user name: There is no change to this variable.

Administrator password: There is no change to this variable.

Password confirmation: There is no change to this variable.

## Increase Disk Space

Assume that this ESM system is experiencing heavier than expected event traffic on ESM, and as a result it is necessary to increase the size of the shared disk to 5TiB (5,242,880 MiB). This section describes how to do that. Note that this process can be accomplished without stopping ESM or unmounting the shared disk.

Purchase a new disk array for each server with the needed capacity. For this example, we assume that the system purchased was a 12x600GB (15K RPM) disk array. Using the Red Hat Facilities to format this as a single RAID 10 partition yields 3.6TB of usable disk space, which is equivalent to 3.3TiB. Assume the name of this partition is `/dev/md11`. Add this partition to the volume group on each server by running (as *root*) the following command:

```
vgextend vg00 /dev/md11
```

This change requires an increase to the size of the metadata volume. The metadata volume on each server must be at least 177 MiB, using the equation:

$$\text{size} = (5767168 \text{ MiB} / 32768) + 1$$

Rounding up to the nearest multiple of 32 gives 192 MiB for the new metadata partition size. The following command is run as *root* on each server to increase the size of the metadata partition:

```
lvresize -L 192M vg00/metadata
```

Increase the size of the shared disk partition (not the filesystem) on both the primary and the secondary to its maximum size. Do that with the following command (as *root*):

```
lvresize -l +100%FREE vg00/opt
```

Inform the HA software that the partition has increased in size by running the following command as *root* on the primary:

```
./arcsight_cluster increaseDisk
```

Increase the size of the filesystem on the primary. As the command below uses `/dev/drbd1`, the filesystem increases will be mirrored on the secondary. `xfs_growfs` is used since this is an XFS filesystem. For an ext4 filesystem `resize2fs` would be used. Run the following command as *root* on the primary only:

```
xfs_growfs /dev/drbd1
```

After you run this command, the `/opt` filesystem will be about 5.5 TiB in size.

Finally, go to the ArcSight Command Center, navigate to **Administration > Storage and Archive**, to the **Storage** tab, and configure the **Default Storage Group** to take advantage of this additional disk space. See the *ArcSight Command Center Users Guide* for further details.

# Chapter 9: Maintain and Monitor the Cluster System

This section covers tasks related to maintaining the primary and secondary systems in the HA Module cluster and also provides guidelines for monitoring the health of the cluster.

## The arcsight\_cluster Script

The `arcsight_cluster` script supports maintenance functions such as retrieving status, and taking servers in and out of service. In this way it is analogous to the `arcsight_services` script that controls services in ESM, as described in the Administrator's Guide.

This script is installed at `/usr/lib/arcsight/highavail/bin/arcsight_cluster` on both the primary and the secondary. Except for specific actions noted below, and unlike ESM commands, `arcsight_cluster` can be run from either the primary or the secondary. To run it you must be logged in as user *root*. The help provides a description of its usage, and the functions it performs.

## Command Syntax

The `arcsight_cluster` command syntax and options are described below. The actions (except help) have more detailed explanations in the topics that follow.

|             |                                                                                        |
|-------------|----------------------------------------------------------------------------------------|
| Description | A tool for managing the HA Module. Run this as user <i>root</i> .                      |
| Applies to  | HA Module on either the primary or secondary machine.                                  |
| Syntax      | <code>/usr/lib/arcsight/highavail/bin/arcsight_cluster &lt;action&gt; [options]</code> |

|          |                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------|---------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Actions  | <code>clusterParameters [--console]</code>                                                        | Update the cluster parameters using the Cluster Parameters Wizard. Only run this on the primary. The <code>--console</code> option displays in console mode. GUI mode is the default.                                                                                                                                                                                                                                   |
|          | <code>diagnose</code>                                                                             | Checks the system health. If any problems are found it corrects them or suggests how the user can correct them. After correcting a problem, run it again to see if there are any other problems.                                                                                                                                                                                                                        |
|          | <code>help (or -h)</code>                                                                         | Provides command usage and HA version.                                                                                                                                                                                                                                                                                                                                                                                  |
|          | <code>increaseDisk</code>                                                                         | Increase the size of the shared partition to fill the volume that backs it. Only run this on the primary. There is no option; it increases the size to the maximum possible size.                                                                                                                                                                                                                                       |
|          | <code>offline [hostname]</code>                                                                   | Makes hostname ineligible to be the primary. If hostname is not specified, the secondary is taken offline. Once off line, a server stays in that state, even if it is or becomes operational, until the online action is issued.                                                                                                                                                                                        |
|          | <code>online [hostname]</code>                                                                    | This action makes the server [hostname] a candidate to be the primary.<br><br>If there is already a primary, the other server is brought online as the secondary and specifying [hostname] is optional.<br><br>If both servers are offline (but ready to be brought on line) you must specify the server to bring online.<br><br>If online is not successful, it will suggest how the user may bring the server online. |
|          | <code>status</code>                                                                               | Print the status of the cluster.                                                                                                                                                                                                                                                                                                                                                                                        |
|          | <code>tuneDiskSync</code>                                                                         | Update the configuration to improve disk sync speed. Do this whenever the speed of the interconnect cable is changed.                                                                                                                                                                                                                                                                                                   |
| Examples | <pre>./arcsight_cluster status</pre> <pre>./arcsight_cluster online myfirstesm.mydomain.com</pre> |                                                                                                                                                                                                                                                                                                                                                                                                                         |

## clusterParameters

This command option starts the Cluster Parameters Wizard. Whether you run it in console or GUI mode, it asks you to provide the following parameters:

- connected hosts
- ping attempts
- ping timeout



## diagnose

The command `arcsight_cluster diagnose` runs a set of tests on your cluster, finds problems, and recommends actions to clear them. The diagnose action deals with the following problems:

- Checks for communication problems between the nodes.
- Suggests ways to bring nodes that are offline to online mode.
  - a. Detects if `arcsight_cluster offline` has been used to take a node offline, and if so, recommends using `arcsight_cluster online`.
  - b. Suggests that you run `crm cluster start`, if appropriate.
  - c. Recovers from `ifdown/ifup`.
- If the disk state is Diskless, it recommends ways to get out of that state.
- Any failures associated with resources are cleared.

If the command returns `2015-11-30 15:07:10 Reconnect attempt failed.`, this may indicate a split-brain condition. See ["Disks on Cluster System Fail to Connect" on page 79](#) for additional steps to evaluate whether that is the case.

## increaseDisk

The `increaseDisk` action provides a way to increase the size of the shared disk. This cannot be done directly because this partition contains disk-synchronization metadata, which must be modified as well. Therefore use this command action as part of the following procedure. You can increase the size of the shared disk without taking the disk or ESM off line.

To increase the size of disk:

1. Determine if the metadata volume needs to be increased in size using the following formula:

The size in mebibytes (MiB, 1,048,576 bytes) can be calculated as

$$\text{size}=(P/32)+1$$

where P is the size of the shared disk partition in gibibytes. For example, if the shared disk partition size is 1 TiB, then  $P=1,048,576$  MiB, and the metadata partition size would be 33 MiB.

If you ever need to increase the size of the shared disk partition, increase the size of the metadata partition accordingly. Decreasing the size of the shared disk partition is not supported.

Use the operating system's Logical Volume Management (LVM) tools to simplify changes. An LVM partition must be a multiple of the LVM chunk size. If you use 32 MiB for the chunk size, for example, then to get a 33 MiB partition, you would take a 64 MiB partition, because you would need two chunks.

Make sure to increase the size of the metadata on both the primary and secondary. They must be the same size. If you are using LVM, the command `lvresize` provides a simple way to do online resizing.

2. Increase the size of the backing device on both the primary and the secondary. Do not increase the size of the file system at this point. This will be done later. The backing device is listed in the file `/etc/drbd.d/opt.res`, on either the primary or the secondary. The line looks like this:

```
disk /dev/mapper/vg00-lv_opt;
```

Increase the size so that the backing devices on the primary and secondary have identical sizes. Again, if you are using LVM, the command `lvresize` provides a simple way to do online resizing.

3. On the primary system run:

```
./arcsight_cluster increaseDisk
```

It will only allow you to proceed if both disks have been increased by the same amount and the metadata volumes are big enough to accommodate this larger size.

4. Increase the size of the `/dev/drbd1` filesystem on the primary. This filesystem is the one mounted at `/opt` or `/opt/arcsight`. The type of the `/dev/drbd1` filesystem is the same as the type of the backing device. If the filesystem is of type `ext4`, use the `resize2fs` command to change the size. If the filesystem is of type `xfs`, use the command `xfs_growfs`.
5. Verify that the command succeeded by running `df -h /opt` on the primary, and noting that the available disk space has increased.

To take advantage of this increased disk space, you may also need to increase the size of the ESM Default Storage Group. You can do this from the ArcSight Command Center (, navigate to **Administration > Storage and Archive**, under the Storage tab). See the *ArcSight Command Center Users Guide* for further details.

## license

The `license` action starts a wizard that allows you to update the license file. For example, after upgrading to ESM version 7.2, you might want to change from an ArcSight license to an Autopass license. You can run the wizard in GUI mode or console mode.

## offline

The offline action lets you take any server out of service for the purpose of performing maintenance on it. Taking the primary offline forces a failover to the secondary. You get a “Do you want to continue?” prompt in that case.

A server won't become “offline” automatically unless all communications with it are lost. Typically, a server is only off line because someone issued the offline action. A server can be in the “offline” state and be operating normally, for example, after the maintenance is completed. An server cannot act as secondary while it is off line. This means that even if it is operating normally, it cannot take over as primary in a failover.

To bring it back on line use the online action.

## online

The online command brings the specified server back online, if it is in the offline state. If that server is already online, no action is taken. Changing a server state to online does not make it the primary; it is merely *eligible* to be the primary.

If there is already a primary server online, then [hostname] is optional; the action brings the server that is not the primary online as the secondary. If both servers are off line, you must specify [hostname].

If you specify online [hostname] for an offline server that is not fully operational, the server's state is changed to online. In that state, it automatically becomes the secondary when it becomes fully operational.

Sometimes the HA Module hesitates to start a resource that has recently and frequently failed. You can clear memory of all failures with the diagnose action. This may help to start resources.

## status

The status action provides you with the current status of the cluster.

### Status Output Example

```
Tue Sep 30 14:39:34 PDT 2014 FAIL Disk: UpToDate/Inconsistent, 0 Nodes
offline, 0 Resources Stopped
```

```
prod01.test.acme.com: online
prod02.test.acme.com: online primary
```

```
Disk: SyncSource UpToDate/Inconsistent
[=====>.....] sync'ed: 38.1% (319920/512200)K
finish: 0:00:08 speed: 38,456 (38,456) K/sec

OK Network-prod01.test.acme.com
OK Network-prod02.test.acme.com

Started ESM
Started Audit-Event-prod01.test.acme.com
Started Audit-Event-prod02.test.acme.com
Started Filesystem
Started Ping-prod02.test.acme.com
Started Ping-prod02.test.acme.com
Started STONITH-SSH-prod01.test.acme.com
Started STONITH-SSH-prod02.test.acme.com
Started Service-IP
```

## Status Output Explanation

The following topics describe different sections of the status output example, above.

### Summary

```
Tue Sep 30 14:39:34 PDT 2014 FAIL Disk: UpToDate/Inconsistent, 0 Nodes
offline, 0 Resources Stopped
```

This line gives the current date and time followed by OK, when the overall status of the HA cluster is OK. In the case above, FAIL indicates that the HA cluster is not OK. In the example provided, the secondary disk is out-of-date (primary status/secondary status). FAIL appears if one or more of the following cases apply:

- The cluster service is down.
- One of the servers is not online.
- The disk communication state is other than Connected.
- One or more of the pacemaker resources is stopped.
- Network communication has failed to one or more servers.

This action (including all options) returns an exit code of zero when it's OK, and non-zero if there is a failure.

The following example indicates that the cluster function has failed:

```
Tue Sep 30 14:48:32 PDT 2014 FAIL Disk: Unconfigured
Cluster is stopped. Run "crm cluster start" to restart it.
Disk: Unconfigured
```

It is possible that even though the server on which you ran this command is reporting this issue, the other server is running as primary without any problems.

## Server Status

The next lines give the status of the servers in the network. Each is either online or offline:

```
prod01.test.acme.com: online
prod02.test.acme.com: online primary
```

Offline may mean that it was put in offline mode by the administrator, or that there has been a failure causing it to go offline. Primary indicates that this server is the primary.

If the secondary was offline or it's cluster function stopped, these lines would look like this:

```
prod01.test.acme.com: offline
prod02.test.acme.com: online Primary
```

## Disk Status

There is only one line if the synchronization is up to date. If the disks are inconsistent, the next line shows a simple progress bar with the percent synchronized and the bytes synchronized out of the total.

```
Disk: SyncSource UpToDate/Inconsistent
 [=====>.....] sync'ed: 38.1% (319920/512200)K
 finish: 0:00:08 speed: 38,456 (38,456) K/sec
```

The first line shows the disk connection state, followed by the disk state of /opt on this server followed by the disk state of /opt on the other server. The next two lines appear if the disk state is SyncSource or SyncTarget. The first means sync is underway from this machine to the other. The second means it is underway from the other machine to this one. These lines contain information about how much space requires sync, how much remains, an estimate of how long the sync will take, and how fast the sync is running.

If the secondary was offline or its cluster function stopped, these lines would be like:

```
Disk: WfConnection UpToDate/Outdated
```

The first word after Disk: indicates the Communication state. The shared disk may have one of the following communication states:

| Connection State | Description                      |
|------------------|----------------------------------|
| Connected        | Data is being mirrored normally. |
| StandAlone       | There is no network connection.  |

| Connection State | Description                                                                                                        |
|------------------|--------------------------------------------------------------------------------------------------------------------|
| SyncSource       | Disk synchronization is underway from the local machine to the other machine. That is, this machine is the primary |
| SyncTarget       | Disk synchronization is underway from the other machine to this machine. That is, this machine is the secondary.   |
| WFConnection     | This machine is waiting for the other machine to connect to it.                                                    |
| Unconfigured     | The server where this command was executed is offline.                                                             |

The second word gives the disk state of this server, followed by a /, followed by the disk state of the other server. The table below shows common disk states.:

| Disk State   | Description                                                                                                            |
|--------------|------------------------------------------------------------------------------------------------------------------------|
| UpToDate     | The data on the disk is current and correct.                                                                           |
| Outdated     | The data on the disk is out of date. No sync is currently going on.                                                    |
| Inconsistent | The data on the disk is out of date, and a sync is going on to correct this.                                           |
| Diskless     | No data can be accessed on the disk. May indicate disk failure.                                                        |
| DUnknown     | The D is for Disk. The other server disk state is not known because there is no communication between the servers.     |
| Consistent   | This server's disk state is correct, but until communication is re-established, it will not be known if it is current. |

If a server is offline, it will say Disk: Unconfigured.

## Connectivity

These lines indicate the connectivity of each server to the network.

OK Network-prod01.test.acme.com

OK Network-prod02.test.acme.com

OK means the server can ping one or more of the hosts specified as a cluster parameter. FAIL means all pings to all hosts on the list failed. When a server is offline, it's network connectivity shows as FAIL.

## Resource Status

The remaining lines report on certain internal resources that the HA Module is managing. In parentheses after each item is the string you can use to search the logs for these entries.

- **ESM** is the ESM instance on the primary (ESM services). The Started status begins when the startup process begins. ESM takes several minutes to complete the startup process and become accessible. During this interval, ESM is not available, even

though the status is Started. Wait a few minutes and try again.

- **Audit-Event-`<hostname>`** is an agent that provides status information that ESM uses to generate audit events.
- **Filesystem** refers to the shared disk filesystem mounted on the ESM machine. (Filesystem)
- **STONITH-SSH-`<hostname>`** is an agent that will reboot the other machine in the cluster when this is necessary.
- **Service-IP** is the service IP of for the ESM machine. (IPAddr2)
- **Ping-`<hostname>`** is a program that checks this machine's connectivity to the network using a ping command. An instance runs on each machine. (ping)

An F after started means that this resource has a positive failure count. You can reset the counter using the ["diagnose" on page 57](#) action. This action will restart the resource.

## tuneDiskSync

The tuneDiskSync action adjusts the disk sync parameters to match the speed of the interconnect cable. It only needs to be run when the speed of these cables is changed. Doing so results in no interruption of service. This is done automatically at installation. If it is not done when the interconnect cable configuration changes, then background sync performance (sync after the systems have been disconnected) may suffer. In particular, if the speed of the interconnect cable is increased, the increase is not translated to an improvement in sync performance until this command is run.

## Log Output

The HA Module produces the following log output:

The HA Module produces log files in `/usr/lib/arcsight/highavail/logs`. These logs are concerned with user-initiated operations. The HA Module configures the operating system to rotate these log files.

This folder contains the following log files:

- `upgrade.log` contains information about the upgrade process.
- `arcsight_cluster.log` contains descriptions of `arcsight_cluster` requests and responses to the user.
- `install-console.log` contains console output for installations run on this machine.
- `install.log` is the installation file for installations run on this machine and contains more detail than `install-console.log`.

- `secondaryHelper.log` contains detailed installation output for installation operations run on this machine, which were actually initiated when the other machine was the primary.

Log rotation occurs at most weekly. Logs are rotated when their size exceeds 1Mbyte. Rotated logs are named `<log-name>-YYYYMMDD`, for example, `install.log-20140501`. The original log plus five rotated logs are kept. The oldest log is removed each time a new log is created.

Cluster log messages from resources other than DRBD are placed in `/var/log/cluster/corosync.log`. Some of these messages plus DRBD messages are sent to the syslog facility `local5`. The storage location of that file depends on the configuration in `rsyslogd.conf`. By default, this output goes to `/var/log/messages`.

In the subtopic ["Resource Status" on page 62](#), each resource description is followed by a string you can use to search `/var/log/cluster/corosync.conf` and `/var/log/messages` to find messages from each of the resources.

## Changing Hostname, IP Address, or Service IP

Choose from the following procedures:

["Changing the Cluster's Service IP Address " below](#)

["Changing the Secondary Hostname or IP Address only" on page 66](#)

["Changing the Primary Hostname or IP Address Only" on page 66](#)

["Changing Both Server Hostnames or IP Addresses" on page 67](#)

["Changing the Interconnect IP Address" on page 68](#)

### Changing the Cluster's Service IP Address

In case you want to change the service IP address of your machines after running the First Boot Wizard successfully, follow these steps. Wherever you see just "hostname," it means "service hostname or service IP address."

To complete these steps, you will need to generate a new key pair (and self-signed certificate) using the new Service IP address.

1. Change the service IP of the cluster using the First Boot Wizard. On the primary, as user *arcsight*, run:

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



There is a field for the Service hostname on the Parameter Configuration panel.  
Finish the First Boot Wizard.

2. Stop the Manager by running (as user *arcsight*):

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

3. While logged in as user *arcsight*, run the following to start the setup program for the Manager from `/opt/arcsight/manager/bin` directory:

```
./arcsight managersetup
```

This opens the Manager's setup wizard.

- a. Enter the new service hostname or service IP address (that you set in the First Boot Wizard) in the Manager Hostname field when prompted by the Manager setup wizard and in every other field where the old hostname is displayed.
  - b. Select the self-signed keypair option when prompted and enter the required information to generate the self-signed certificate with the new service IP address. If ESM is configured for FIPS mode, you will not get this option. The key-pair must be generated manually using the `runcertutil` utility.
4. Start the Manager and all other processes by running (as user *arcsight*):  

```
/etc/init.d/arcsight_services start
```
  5. 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. Go on to the next step when you see the line "manager service is available".
  6. Make sure you can start the ArcSight Command Center by browsing to the following URL:  

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

Where `<hostname>` is the new hostname (note that hostnames with underscores do not work on IE, so use the IP address.)
  7. Import the Manager's newly-generated certificate on all clients (ArcSight Console and connectors) that access the Manager. Use `keytoolgui`. See the "SSL Authentication" section of the ESM Administrator's Guide for details about this tool. Use `runcertutil` if you are running ESM using FIPS mode. See "Tools Used to Configure Components in FIPS" in the ESM Administrator's Guide for details about the `runcertutil` tool.
  8. Test to make sure that:
    - The clients can connect to the Manager.
    - Peer configuration works as expected. If not, redo the peer configuration.

## Changing the Secondary Hostname or IP Address only

Use the following procedure to change the hostname or IP address of the secondary server only. During this procedure, ESM remains running on the primary; there is no interruption.

1. Run the following command on the secondary as user *root*:  
`crm cluster stop`
2. Change the hostname and/or IP address of the secondary as required.
3. If you changed the system hostname:
  - a. Run the following commands on the secondary system as user *root*:  
`systemctl disable corosync OR chkconfig --del corosync`  
`systemctl disable pacemaker OR chkconfig --del pacemaker`
  - b. Reboot the secondary system.
  - c. Test that the change persists across reboots. Use the `hostname` command to show the system hostname.
4. On the primary, as user *arcsight*, run:  
`/usr/lib/arcsight/highavail/bin/arcsight firstBootWizard`  
In the First Boot Wizard, specify the new hostname for the secondary system.

When the First Boot Wizard completes, the cluster restarts and you are done.

## Changing the Primary Hostname or IP Address Only

Use the following procedure to change the hostname or IP address of the primary server only. Basically, you force the primary to fail over then, when it has become the secondary, you use the procedure for changing the secondary.

1. Run the following command on the primary system as user *root*:  
`crm cluster stop`
2. Wait until the failover to the other ESM is complete.
3. On the same machine, which is now the secondary, change the hostname and/or IP address of the (new) secondary (formerly the primary) as required.
4. If you changed the system hostname:
  - a. Run the following commands on the secondary system as user *root*:  
`systemctl disable corosync OR chkconfig --del corosync`  
`systemctl disable pacemaker OR chkconfig --del pacemaker`

- b. Reboot the secondary system.
  - c. Test that the change persists across reboots. Use the `hostname` command to show the system hostname.
5. On the primary, as user *arcsight*, run:  
`/usr/lib/arcsight/highavail/bin/arcsight firstBootWizard`
6. In the First Boot Wizard, specify the new hostname or IP address for the secondary. When the First Boot Wizard completes, the cluster restarts.

## Changing Both Server Hostnames or IP Addresses

**IMPORTANT:**The following procedure can be used only if both of the new IP Addresses are in the same subnet as the old ones. If the new IP Addresses are in a different subnet (for example, if you are converting from IPv4 to IPv6), you must uninstall and then re-install the HA Module.

1. Run the following command on the secondary (System B) as user *root*:  
`crm cluster stop`
2. Change the hostname and/or IP address of the secondary (System B) as required.
3. If you changed the system hostname:
  - a. Run the following commands on the secondary system as user *root*:  
`systemctl disable corosync OR chkconfig --del corosync`  
`systemctl disable pacemaker OR chkconfig --del pacemaker`
  - b. Reboot the secondary system.
  - c. Test that the change persists across reboots. Use the `hostname` command to show the system hostname.
4. On the primary system (System A), as user *arcsight*, run:  
`/usr/lib/arcsight/highavail/bin/arcsight firstBootWizard`  
In the First Boot Wizard, specify the new hostname for the secondary (System B) system. When the First Boot Wizard completes, the cluster restarts and you are done with the secondary (System B). Wait for the shared disk to complete its sync. When run this command as user *root*:  
`/usr/lib/arcsight/highavail/bin/arcsight_cluster status`  
the `Disk` line in the status information should read:  
`Disk: Connected UpToDate/UpToDate`  
Note that it may take some time for the sync to complete.
5. Run the following command on the primary (System A) as user *root*:

```
crm cluster stop
```

The primary (System A) will failover to the secondary (System B).

6. On the same machine as the previous step (System A), change the hostname and/or IP address as required.
7. If you changed the system hostname:
  - a. Run the following commands on the secondary system as user *root*:

```
systemctl disable corosync OR chkconfig --del corosync
systemctl disable pacemaker OR chkconfig --del pacemaker
```
  - b. Reboot the secondary system.
  - c. Test that the change persists across reboots. Use the `hostname` command to show the system hostname.
8. On the new primary system (System B), as user *arcsight*, run:

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

9. In the First Boot Wizard, specify the new hostname or IP address for the new secondary (System A). When the First Boot Wizard completes, the cluster restarts.

## Changing the Interconnect IP Address

Use the following procedure to change the interconnect IP address on either the primary or the secondary system:

1. As user *root* on the secondary system, run `crm cluster stop`.
2. Change to the `/etc/sysconfig/network-scripts` directory.
3. Select and edit the file for the network interface that you want to change by changing the `IPADDR` value. For example the file might be `ifcfg-eth1`.
4. Run the `ifdown` and `ifup` commands (for example, `ifdown eth1; ifup eth1`).
5. Run the First Boot Wizard on the primary system and specify the new interconnect cable IP address(es).

## Replacing a Server

This topic describes how to use the First Boot Wizard to replace a server (for example, if it has hardware issues).

### To replace a server:

1. Power down the server to be replaced.  
The remaining server becomes the primary server.
2. Prepare the new server as described in [Existing ESM Installation: Configuring the High Availability System](#).  
The new server can have a different IP address and host name than the one it is replacing.
3. As user root, stop ESM services on the primary system:

```
/opt/arcsight/manager/bin/remove_services.sh
```

4. As user arcsight, run the First Boot Wizard on the primary system and specify the host name or IP address for the new secondary system if it is different from the original.
5. If you are replacing a server in a distributed correlation environment where the persister is part of the HA cluster, run the following command:

```
/etc/init.d/arcsight_services sshSetup
```

**Note:** This is only necessary in a distributed correlation environment.

6. As user root, restart the ESM services on the primary system:

```
/opt/arcsight/manager/bin/setup_services.sh
```

At this point, ESM should start on the primary system and the new server will become the secondary system. The synchronization process between the primary system and the new secondary system might take some time. For more information, see [Planning for the Initial Disk Synchronization](#).

## Changing Mount Options

Changing the -o options on a mount command is the same as without the HA Module, except that one extra command is required. To change the options, log into the primary as root and run the following command:

```
mount -t <file system type> -o remount,<new mount options> /dev/drbd1 <shared disk>
```

Where:

- <file system type> must be ext4 or xfs, and *cannot be changed*.
- <new mount options> are the new options you want.

- <shared disk> is where the shared disk is mounted, which *cannot be changed* (typically /opt or /opt/arcsight).
- /dev/drbd1 is the name of the mirrored volume.

Then run the following command as user *root* on the primary. This command makes the changes permanent across failovers:

```
./arcsight_cluster tuneDiskSync
```

# Chapter 10: Troubleshooting the Systems

The following information may help solve problems that occur while operating the HA system. In some cases, the solution can be found here or in specific ArcSight documentation. This chapter includes the following topics:

- [Installation Issues and Solutions](#) .....71
- [General Problems](#) .....75
- [Audit Events](#) .....75
- [Failover Triggers](#) .....77
- [Processes Killed During Failover](#) .....77
- [System does not Failover](#) .....78
- [Network Interface Commands Stall Disk Mirroring](#) .....78
- [No ESM Uninstall Links on the Primary](#) .....78
- [Stopping the Network on the Secondary Kills ESM](#) .....79
- [Disks on Cluster System Fail to Connect](#) .....79

## Installation Issues and Solutions

The following table describes possible installation error messages and how to resolve them. Angle brackets (< >) enclose values such as names or IP addresses that are unique to your message. The messages are in the following format:

[Primary|Secondary]: [Timestamp] ERROR - <message>

| Installation Message                                                              | Description                                                                                               |
|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| <b>User and Access Issues</b>                                                     |                                                                                                           |
| Fatal error on <hostname>. See <log file>.                                        | An unexpected error caused SSH to fail to <hostname>. For more information, check the specified log file. |
| Timeout on SSH to <hostname>. SSH access to <hostname> failed to connect quickly. | Resolve the SSH communication problem.                                                                    |
| Incorrect root password for <hostname> - please enter correct one.                | The password that you specified is not correct. Provide the correct password.                             |
| Failed to set up SSH access. See <log file> for details.                          | SSH access was not successful. For more information, check the specified log file.                        |

| Installation Message                                                                                        | Description                                                                                                                                                                           |
|-------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No arcsight user on secondary. Please create one identical to that on primary.                              | On the secondary, create an arcsight user that is identical to the arcsight user on the primary.                                                                                      |
| arcsight users on primary and secondary must be set up identically.                                         | The user or group IDs of the arcsight users differ on the primary and the secondary. Ensure that they are identical.                                                                  |
| arcsight users on primary and secondary must have the same home directory.                                  | Ensure that the arcsight users on the primary and the secondary have the same home directory.                                                                                         |
| <b>Crossover Cable Issues</b>                                                                               |                                                                                                                                                                                       |
| Speed of primary end of crossover cable is <primaryCableSpeed>M - must be at least 1000M.                   | The primary interface for interconnect is slower than the Gigabit ethernet. Use a faster interface.                                                                                   |
| Speed of secondary end of crossover cable is <secondaryCableSpeed>M - must be at least 1000M.               | The secondary interface for interconnect is slower than the Gigabit ethernet. Use a faster interface.                                                                                 |
| Primary Cable IP <primaryCableCIDR> and Secondary Cable IP <secondary_cable_ip> must be in the same subnet. | Ensure that the primary cable IP address and the secondary cable IP address are in the same subnet.                                                                                   |
| No interface found for <primary_cable_ip> on Primary.                                                       | The primary cable IP address does not correspond to an interface. This was probably a data entry error in the First Boot Wizard.                                                      |
| No interface found for <secondary_cable_ip> on Secondary.                                                   | The secondary cable IP address does not correspond to an interface. This was probably a list selection error in the First Boot Wizard.                                                |
| <b>Shared Disk Issues</b>                                                                                   |                                                                                                                                                                                       |
| Unmount of <shared_disk> failed. Fix the problem, and re-run this script.                                   | Correct the problem, and then run the First Boot Wizard again.                                                                                                                        |
| Permanently unmount the following mounts on <shared_disk>, and then retry installation: <mount name>        | The listed mounts mount on top of /opt or /opt/arcsight. This is not supported. Unmount them and then remove them from /etc/fstab.                                                    |
| <metadata_vol> should not be mounted.                                                                       | The metadata volume is mounted. Unmount the volumes. The HA module will probably also generate the <metadata_vol> appears to be in use error. Follow the instructions for that error. |



| Installation Message                                                                                                                                                                                                                                                                                                                                    | Description                                                                                                                                                                                                                                                                            |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>&lt;metadata_vol&gt; appears to be in use. See the following output from<br/>blkid -o export &lt;metadata_vol&gt;</p> <p>--- blkid output here ---</p> <p>If this volume is not in use, run<br/>dd if=/dev/zero of=&lt;metadata_vol&gt;<br/>as user root on hostname &lt;hostname&gt; to clear this volume and then rerun the First Boot Wizard.</p> | <p>The metadata volume is in use. Ensure that this is not the case, run the command specified in the message, and then run the First Boot Wizard again.</p>                                                                                                                            |
| <p>Disk status must be Connected to reconfigure cluster.</p>                                                                                                                                                                                                                                                                                            | <p>The HA module is already installed on both machines, so this call to the First Boot Wizard is to reconfigure the installation. This can only be done if the disk status is Connected (normal).</p> <p>Run ./arcsight_cluster diagnose and then run the First Boot Wizard again.</p> |
| <p>Please mount &lt;shared_disk partition&gt;, and re-run installation.</p>                                                                                                                                                                                                                                                                             | <p>Mount the shared disk and then run the installation again.</p>                                                                                                                                                                                                                      |
| <p>Size of metadata volume &lt;metadata_vol&gt; is less than required minimum of &lt;megabytes&gt;M</p>                                                                                                                                                                                                                                                 | <p>The metadata volume is too small to support the shared disk. Increase the size of the metadata volume.</p>                                                                                                                                                                          |
| <p>The size of &lt;volume&gt; on the secondary is &lt;megabytes&gt;M. It must be the same as the primary - &lt;megabytes&gt;M.</p>                                                                                                                                                                                                                      | <p>This could refer either to the shared disk volume or the metadata volume. The size of each must be the same on each server (rounded to the nearest Mbyte). Change the sizes to be identical.</p>                                                                                    |
| <p>&lt;volume&gt; is not a valid disk volume.</p>                                                                                                                                                                                                                                                                                                       | <p>Either the shared disk or the metadata volume is not really a volume. Check to see if there is a typographical error in the name you specified.</p>                                                                                                                                 |
| <p>Found &lt;megabytes&gt;M disk space used on &lt;shared_disk&gt;.</p> <p>The installation will not proceed with these files in place. If these files are not important, run "rm -rf &lt;shared_disk&gt;/*" as root on &lt;hostname&gt; and re-run the First Boot Wizard.</p>                                                                          | <p>The installation found more than 10MB of files on &lt;shared_disk&gt; on the secondary. The installation is terminated. Remove the files, and then run the First Boot Wizard again.</p>                                                                                             |
| <p>&lt;shared disk volume&gt; mounted on &lt;shared_disk&gt; on the primary and on &lt;secondary_disk&gt; on the secondary. It must be mounted on the same mount point on both machines.</p>                                                                                                                                                            | <p>Ensure that the volume of the shared disk is mounted on the same mount point on both machines.</p>                                                                                                                                                                                  |
| <p>Cannot do a Reconfiguration when disks are in &lt;status&gt; status. Please correct the disk status before doing reconfiguration.</p>                                                                                                                                                                                                                | <p>The &lt;status&gt; value in the message is either StandAlone or WfConnection. The reconfiguration will not work unless disk mirroring is functioning. You can usually use the arcsight_cluster script, ./arcsight_cluster diagnose, to fix this problem.</p>                        |
| <p><b>Primary/Secondary Host Issues</b></p>                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                        |

| Installation Message                                                                                                                                                                                                         | Description                                                                                                                                                                                                                                                                        |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No interface found for <primary_ip> on Primary                                                                                                                                                                               | The primary IP address or host name must be the first IP address on an interface. Configure the primary host name to correspond to an interface.                                                                                                                                   |
| No interface found for <secondary_ip> on Secondary                                                                                                                                                                           | The secondary IP address or host name must be the first IP address on an interface. Configure the secondary host name to correspond to an interface.                                                                                                                               |
| unsupported kernel version <version>                                                                                                                                                                                         | The kernel version on this server does not correspond to an operating system that the HA module supports. Upgrade the operating system to a supported version.                                                                                                                     |
| ERROR installing RPMs - Please check the log file <logfile> on <hostname> for details about the error.                                                                                                                       | RPMs were not installed. For more information, check the specified log file.                                                                                                                                                                                                       |
| Primary IP <primary IP> and Secondary IP <secondary IP> must be in the same subnet.                                                                                                                                          | Change the host IP addresses so that they are in the same subnet.                                                                                                                                                                                                                  |
| <hostname> - the hostname of this host does not correspond to the hostname given for either the Primary or the Secondary.                                                                                                    | Correct the host name.                                                                                                                                                                                                                                                             |
| <host> does not resolve to <IP> /                                                                                                                                                                                            | Correct the DNS or <code>/etc/hosts</code> so that <host> resolves to the IP address on the server.                                                                                                                                                                                |
| OS version on primary and secondary are different.                                                                                                                                                                           | Ensure that the primary and the secondary are on the same operating system version.                                                                                                                                                                                                |
| Could not send and return test string using ssh. Expected "test", saw "\$returnedString"                                                                                                                                     | There is a problem with the SSH login. Ensure that the root user can SSH between systems in both directions (i.e., from System A to System B and from System B to System A).                                                                                                       |
| remove added message of the day or login string from root logins. Expected "test" saw <returnedString>                                                                                                                       | A message of the day string has been detected. This might cause problems with SSH communication. Disable the SSH banner by creating an empty hushlogin file in the root user's home directory: <code># touch /root/.hushlogin</code> .                                             |
| Cluster did not come up after installation. See the status output above this message.                                                                                                                                        | This happens rarely. Check the <code>install.log</code> file for details. This message might have been generated because of a temporary condition, and within a few minutes the system will work as expected. If the problem persists, contact <a href="#">Technical Support</a> . |
| Host IP addresses are all <IPv4/IPv6> and cable IP addresses are all <IPv6/IPv4>. Change both to either IPv4 or IPv6 to support redundant pacemaker communication. Otherwise, pacemaker communication will not be redundant. | This is a warning only. You can improve the system's redundancy by using exclusively IPv4 or IPv6 addresses. If you do so, the system will be more resistant to communications failures.                                                                                           |

| Installation Message                                                                          | Description                                                                                                                       |
|-----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| SELinux on <primary/secondary> is Enforcing - HA does not support SELinux. Please disable it. | The HA module does not support SELinux. For information about disabling it, see the documentation for your operating system.      |
| <b>Cluster Upgrade Issues</b>                                                                 |                                                                                                                                   |
| Cluster should not be running during upgrade. Run "crm cluster stop" as root to stop cluster. | The system should not be running during the upgrade process. As user root, run <code>crm cluster stop</code> to stop the cluster. |

## General Problems

Your first resort for troubleshooting cluster problems should be the command:

```
./arcsight_cluster diagnose
```

This command clears some common problems automatically and provides simple solutions for others.

## Changing ESM to IPv6

If you change ESM from IPv4 to IPv6 after the HA Module is installed, it means that you are changing the subnet. Changing the subnet requires that you uninstall and reinstall the HA Module.

## Audit Events

Audit events are events generated within the Manager to mark a wide variety of routine actions that can occur manually or automatically, such as adding an event to a case or when synchronization of the two systems begins. Audit events have many applications, which can include notifications, task validation, compliance tracking, automated housekeeping, and system administration.

This topic lists the High Availability Option audit events you can use in rules, filters, and other analytical or administrative resources. Observe the way these events are used in the standard system-related content for examples of how to apply them.

From the table below, use the Device Event Class (DEC) ID string in rules and filters. The **Audit Event Description** reflects the event name you see in active channel grids.

| Device Event Class ID | Audit Event Description |
|-----------------------|-------------------------|
| highavailability:100  | Primary Manager started |
| highavailability:200  | HA system failure       |
| highavailability:300  | Disk sync in progress   |
| highavailability:500  | HA system restored      |

## highavailability:100

This event occurs when there is a failover causing the secondary system to take over and become the primary machine. It also occurs every time ESM starts up, with or without a failover.

Severity: 3

Device event category: /Monitor/Manager/HighAvailability/Primary/Up

## highavailability:200

This is a system-failure event that occurs if the secondary system becomes unavailable and cannot assume the role of the primary system. This event is generated every five minutes until the secondary system is restored. The event includes a **reason** field that provides more detailed information. There are numerous possible causes:

- Failure of either network interface card (NIC)
- Cross-over cable failure or disconnect
- Secondary system failure or shutdown
- Secondary system hard drive failure.
- You reboot the secondary system for any reason

Severity: 7

Device event category: /Monitor/Manager/HighAvailability/Status/Failed

## highavailability:300

This event occurs when the Distributed Replicated Block Device (DRBD) storage system begins the process of synchronizing the primary and secondary hard drives and continues every five minutes (by default) until the synchronization is complete. Each event includes the amount of data between the two systems that has been synchronized as a percentage until it reaches 100 percent. You can change the interval using the `highavailability.notification.interval` property as described in [“Configurable Properties” on page 1](#).

Severity: 4

Device event category: /Monitor/Manager/HighAvailability/Sync/InProgress

## highavailability:500

The HA system is restored. This event occurs when the secondary system changes from a failed status (highavailability: 200 or 300) to OK. It may take 30 seconds for this event to generate after the secondary system and high-availability service is restored.

Severity: 3

Device event category: /Monitor/Manager/HighAvailability/Status/OK

## Failover Triggers

The following situations can trigger a failover:

- You place the primary in offline mode using the `arcsight_cluster` command.
- The primary operating system goes down. In the case of a routine system restart, the machine doing the restart might continue to be primary. This is true when the system starts again before the failover had time to trigger.
- The hard disk on the primary system fails.
- The primary system loses an internet connection.

The following situations do not trigger a failover:

- You manually stop the ArcSight Manager or any of its services. For example, changing a property in the `server.properties` file and starting the Manager again does not trigger a failover.
- The network switch fails, causing a communications failure to both primary and secondary systems. Users will immediately detect that the ArcSight Console or ArcSight Command Center has lost communication with the Manager. The primary continues to run and connectors cache events until communications are restored, at which time the primary ESM continues as usual.
- The primary system runs out of disk space and the secondary system also runs out of space because of mirroring.

## Processes Killed During Failover

As a part of failover, the HA Module shuts down ESM and all processes on the old primary that are accessing its shared disk. This includes, for example, ESM wizards or shell windows that have changed directory to the shared disk. Killing these processes is a necessary step prior to unmounting the shared disk.

## System does not Failover

Failovers might fail to trigger on a system where the shared disk is in XFS format and the inode64 mount option is not used. This happens in particular if the inode64 option was used at some previous time, and then is not used later.

To fix this problem, follow the procedure described in ["Changing Mount Options" on page 69](#), adding the inode64.

Your mount command might look something like this:

```
mount -t xfs -o remount,inode64 /dev/drbd1 <shared disk>
```

## Network Interface Commands Stall Disk Mirroring

If you use network interface commands such as:

- `ifdown <interface>` followed by `ifup <interface>`,
  - `ifconfig <interface> down` followed by `ifconfig <interface> up`, or
  - `ip set <interface> down`, followed by `ip set <interface> up`
- ... the disk mirroring component does not recover automatically.

To recover, run `./arcsight_cluster diagnose`. This command clears the condition and restores normal operations.

## No ESM Uninstall Links on the Primary

The mirrored disk containing the ESM installation is only mounted on the current primary server. This may be different from the server where ESM was installed. ESM must always be uninstalled from the current primary.

When the machine on which ESM was originally installed fails over to the other machine, that other machine (now the primary) does not have the uninstall link if it was saved to a location outside the scope of the disk mirroring. To uninstall ESM from that machine, use the procedure described in the *ESM Installation Guide* topic entitled "Uninstalling ESM."

## Stopping the Network on the Secondary Kills ESM

If you run the command `systemctl stop network` or `service network stop` on the secondary, it *sometimes* results in the ESM on the primary shutting down. If that happens, it triggers a failover that cannot complete if the network service is stopped. The command breaks the secondary's connection to both the primary/secondary interconnect cable and the internet. Running `systemctl start network` or `service network start` by itself does not restore ESM.

To recover from this situation, run `systemctl start network` or `service network start`. Then run `./arcsight_cluster diagnose` on both machines. This command repairs the condition and restarts ESM on the original primary.

You might expect that if you stop the network on the primary it triggers a failover, but stopping it on the secondary is actually worse. It creates a situation that wants to trigger a failover, the failover cannot complete because the network is stopped on the secondary and you end up with ESM not running on either machine.

Avoid using `systemctl stop network` or `service network stop` on either machine.

## Disks on Cluster System Fail to Connect

In this scenario, the disk status will be either `WFConnection` or `Standalone` on both systems. The command `./arcsight_cluster diagnose` will clear this condition in simple cases (see details about ["diagnose" on page 57](#)). If you see the following output, there may be a split brain condition:

```
2015-11-30 15:07:10 Reconnect attempt failed.
```

To check whether this is a split brain condition, run the following command as the root user:

```
grep Split-Brain /var/log/messages
```

If the 'Split-Brain' keyword appears in recent messages, this confirms that the split brain condition has occurred. You must choose which machine has the most up-to-date data, called System A in the following procedure. The machine with the older data is called System B in the following procedure.

Perform the following steps to correct the split brain condition. When these steps are complete, data from System A will be synced to System B.

1. On System B, as the root user run `crm cluster stop`. It may take up to 10 minutes for ESM to stop.
2. On System B, make sure that the shared disk (e.g. `/opt`) is unmounted before you perform the next steps.
3. On System B, run the following commands as the root user:  
`drbdadm up opt`  
`drbdadm disconnect opt`  
`drbdadm secondary opt`  
`drbdadm connect --discard-my-data opt`
4. On System B, as the root user run either `crm cluster start`.
5. On System A (the machine with up-to-date data), run the following command:

```
drbdadm connect opt
```

The cluster should come up normally within a few minutes. If you get the following error, you can ignore it.

```
opt: Failure: (102) Local address(port) already in use. Command
'drbdsetup-84 connect opt ipv4:10.0.0.89:7789 ipv4:10.0.0.87:7789 --
protocol=C --max-buffers=128K --max-epoch-size=16K --sndbuf-size=0 --
csums-alg=sha1 --after-sb-0pri=discard-least-changes' terminated with exit
code 10
```



# Appendix A: The highavail.properties File

The First Boot Wizard generates the highavail.properties file that defines certain cluster configuration properties. If the First Boot Wizard was run at least once, this file should exist at: /usr/lib/arcsight/highavail/highavail.properties. The highavail.properties can be loaded in the First Boot Wizard during the HA Module installation process to simplify the wizard steps. It is required to run the prepareHA.sh script.

If you are installing the HA Module for the first time, this file will not exist. If you want to use it with the First Boot Wizard or prepareHA.sh script, you must create it with a text editor. Copy and rename the template.properties file, located in the "Tools/highavail" directory where you unpacked the ESM 7.2 Installation Package. The following example provides guidance about how to define each property value. The actual values will be unique to your deployment environment.

```
service_hostname=esm.internal.acme.com
shared_disk=/opt
metadata_volume=/dev/mapper/vg00-metadata
primary_cable_ip=198.166.11.4
primary_hostname=ha1.internal.acme.com
secondary_cable_ip=198.166.11.3
secondary_hostname=ha2.internal.acme.com
```

## Appendix B: Upgrading the Operating System on Software ESM

This appendix describes upgrading the operating system from RHEL 7.6 to RHEL 7.7 on software ESM with a version 7.2 High Availability (HA) module.

### To upgrade the operating system:

1. On the secondary system, as user `root`, run the following command:  
`/usr/lib/arcsight/highavail/install/preUpgrade.sh`
2. On the secondary system, upgrade the operating system to RHEL 7.7.  
For more information, see the vendor documentation.
3. Reboot the secondary system.
4. On the primary system, stop ESM:  
`/etc/init.d/arcsight_services stop`
5. On the primary system, stop the cluster:  
`crm cluster stop`
6. On the primary system, if `/usr/lib/arcsight/highavail.backup` exists, move it to `/usr/lib/arcsight/highavail.save`.
7. On the primary system, as user `root`, run the following command:  
`/usr/lib/arcsight/highavail/install/preUpgrade.sh`
8. On the primary system, upgrade the operating system to RHEL 7.7.  
For more information, see the vendor documentation.
9. Reboot the primary system.
10. On the primary system, as user `root`, run the following command:  
`/usr/lib/arcsight/highavail/install/upgrade.sh`  
The response includes the line `This will be an OS upgrade only`, indicating that the HA module will not be upgraded.
11. On the primary system, start ESM:  
`/etc/init.d/arcsight_services start`

# Appendix C: Upgrading the Operating System on a High Availability Appliance

This appendix describes upgrading the High Availability (HA) module from RHEL 7.4 to RHEL 7.5 and then to RHEL 7.7 on an appliance.

**Note:** You cannot upgrade directly from RHEL 7.4 to RHEL 7.7. Before upgrading to RHEL 7.7, you must upgrade both the primary and the secondary systems to RHEL 7.5. If you previously completed the upgrade to 7.5, continue to [Upgrading from RHEL 7.5 to RHEL 7.7](#).

## Upgrading to RHEL 7.5

When upgrading to RHEL 7.5, you must prepare both systems for upgrade, and then upgrade the secondary system first.

### To prepare the systems for upgrade:

1. On both systems:
  - a. As user root, disable drbd.service:  

```
systemctl disable drbd.service
```
  - b. Verify that drbd.service is disabled:  

```
systemctl list-unit-files --type=service |grep drbd
```
2. As user root, place the secondary system in standby mode:  

```
crm_standby -v true
```
3. As user root, take the secondary system offline:  

```
systemctl stop heartbeat
systemctl disable heartbeat
```

Proceed to upgrade the secondary system.

### To upgrade the secondary system to RHEL 7.5:

1. Log in as user root.
2. Transfer the following files to the /tmp partition:
  - `esm_osupgrade_rhel75_20180828142739.tar.gz`
  - `esm_apha_support_rpms_rhel75.tar.gz`

- HA\_7.0.0\_Update\_For\_7.5OS
3. From the directory where you placed the archive, extract `esm_osupgrade_rhel75_20180828142739.tar.gz`.
  4. Change directory:  
`cd rhel75-upgrade`
  5. Make the script executable:  
`chmod 0700 osupgrade`
  6. Start the upgrade:  
`./osupgrade 2>&1 | tee osupgrade.log`
  7. After the script completes, ensure that the system is rebooted.
  8. Verify the operating system version:  
`cat /etc/redhat-release`  
You should receive the following result:  
Red Hat Enterprise Linux Server release 7.5
  9. From the directory where you placed the archive, extract and install the HA support RPMs:  
`tar -zxvf esm_osupgrade_rhel75_20180828142739.tar.gz`  
`cd esm_apha_support_rpms_rhel75.tar.gz`  
`./install_apha_support_pkgs.sh`
  10. From the directory where you placed the archive, extract and install the HA update:  
`cd HA_7.0.0_Update_For_7.5OS/`  
`cd dist/`  
`tar xf HA_7.0.0_Update_For_7.5OS.tgz`  
`cd HA_7.0.0_Update_For_7.5OS/`  
`./HAUpdate.sh`
  11. As user root, bring the secondary system online:  
`systemctl start heartbeat`  
`systemctl enable heartbeat`
- Proceed to upgrade the primary system.

### To upgrade the primary system to RHEL 7.5:

1. Log in as user root.
2. Transfer the following files to the `/tmp` partition:
  - `esm_osupgrade_rhel75_20180828142739.tar.gz`
  - `esm_apha_support_rpms_rhel75`

- HA\_7.0.0\_Update\_For\_7.50S
3. Stop all ArcSight services:  
`systemctl stop heartbeat`  
`systemctl disable heartbeat`
  4. Take the primary system offline:  
`service heartbeat stop`
  5. From the directory where you placed the archive, extract `esm_osupgrade_rhel75_20180828142739.tar.gz`.
  6. Change directory:  
`cd rhel75-upgrade`
  7. Make the script executable:  
`chmod 0700 osupgrade`
  8. Start the upgrade:  
`./osupgrade 2>&1 | tee osupgrade.log`
  9. After the script completes, ensure that the system is rebooted.
  10. Verify the operating system version:  
`cat /etc/redhat-release`  
You should receive the following result:  
Red Hat Enterprise Linux Server release 7.5
  11. From the directory where you placed the archive, extract and install the HA support RPMs:  
`tar -zxvf esm_apha_support_rpms_rhel75.tar.gz`  
`cd esm_apha_support_rpms_rhel75`  
`./install_apha_support_pkgs.sh`
  12. From the directory where you placed the archive, extract and install the HA update:  
`cd HA_7.0.0_Update_For_7.50S/`  
`cd dist/`  
`tar xf HA_7.0.0_Update_For_7.50S.tgz`  
`cd HA_7.0.0_Update_For_7.50S/`  
`./HAUpdate.sh`
  13. Bring the primary server online:  
`systemctl start heartbeat`  
`systemctl enable heartbeat`
  14. Correct a Linbit bug in the 7.4 operating system that causes a large delay in the transfer of services when one system in the cluster goes offline:  
`crm configure property cluster-recheck-interval=1M`

The following messages are displayed:

```
WARNING: unrecognized CIB element <built-in function Comment>
```

```
ERROR: cib-bootstrap-options: attribute expected-quorum-votes does not exist
```

15. Return to the secondary system and, as user root, bring it online:

```
crm_standby -D
```

16. Check the status of the cluster:

```
/usr/lib/arcsight/highavail/bin/arcsight_cluster status
```

17. If any ArcSight services are not restarted automatically, restart them on the primary server:

```
service arcsight_services start
```

18. Verify that all of the ArcSight services have started.

19. Start the ArcSight Console and verify that it is operating correctly.

**Note:** If, after the upgrade, the disks do not connect, run `arcsight_cluster diagnose` to clear the problem.

## Upgrading from RHEL 7.5 to RHEL 7.7

After the upgrade to RHEL 7.5 is complete, you can proceed to upgrade to RHEL 7.7.

If you were prompted to upgrade the operating system in [Upgrading ESM and the High Availability Module](#), complete the following procedure on the system that requires the upgrade.

**Note:** Do not perform this procedure on both systems at the same time. When you are prompted to upgrade the operating system on the secondary system, complete this procedure on that system and then return to the point where you left off in [Upgrading ESM and the High Availability Module](#). When you are prompted to upgrade the operating system on the primary system, complete this procedure on that system and then return to the point where you left off in [Upgrading ESM and the High Availability Module](#).

### To upgrade to RHEL 7.7:

1. Because the default appliance firewall rules prevent ESM HA 7.2 from working, as user root, stop and disable `firewalld`:

```
systemctl stop firewalld
```

```
systemctl disable firewalld
```

2. From the [Licensing and Downloads](#) site, download the upgrade script (`esm_`

osupgrade\_rhel177\_xxxx.tar.gz) and the HA support package (esm\_apha\_support\_rpms\_rhel177\_G9.tar.gz).

After you download the files, contact [Technical Support](#) to verify that the signed software is from Micro Focus and has not been manipulated by a third party.

3. As user root, create a temporary directory for the download. For example, /downloadtmp. Do not use the /opt directory.
4. Transfer the files that you downloaded to the temporary directory.
5. From the temporary directory, extract the files:

```
/bin/tar xzf esm_osupgrade_rhel177_xxxx.tar.gz
```

6. Change directory to esm-rhel177upgrade.
7. Make the script executable:

```
chmod 0700 osupgrade
```

8. Start the upgrade:

```
./osupgrade 2>&1 | tee osupgrade.log
```

After the script completes, reboot the system.

9. Verify the operating system version:

```
cat /etc/redhat-release
```

The output should be similar to the following:

```
Red Hat Enterprise Linux Server release 7.7
```

10. Untar esm\_apha\_support\_rpms\_rhel177\_G9.tar.gz:

```
tar xzf esm_apha_support_rpms_rhel177_G9.tar.gz
```

11. Run the following commands:

```
cd esm_apha_support_rpms_rhel177
./install_apha_support_pkgs.sh
```

You can now return to [Upgrading ESM and the High Availability Module](#), resuming at step 7 or step 13, depending on whether you upgraded the operating system on the primary system or the secondary system.

# 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 ESM High Availability Module User's Guide (ESM 7.2)**

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](mailto:arcsight_doc@microfocus.com).

We appreciate your feedback!