HP ArcSight ESM High Availability Module

Software Version: 6.9.1

ESM High Availability Module User's Guide

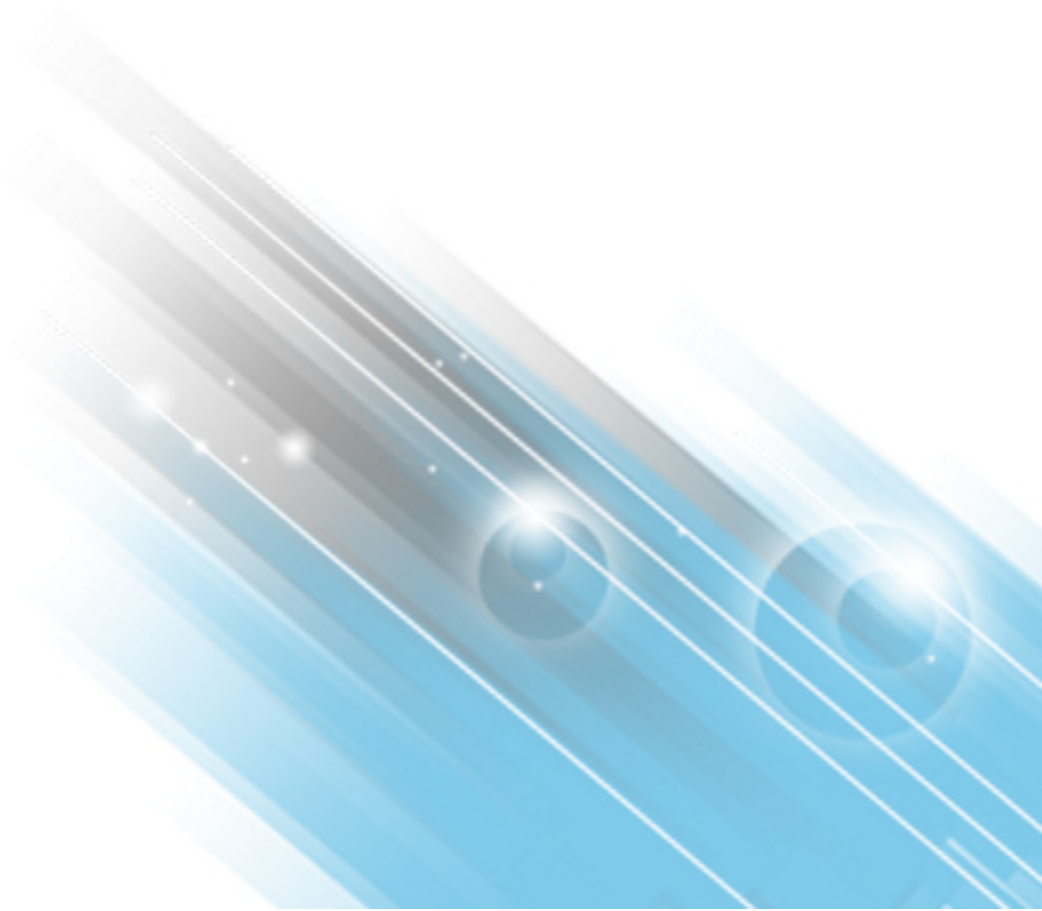

February 2016

Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

The network information used in the examples in this document (including IP addresses and hostnames) is for illustration purposes only.

HP ArcSight products are highly flexible and function as you configure them. The accessibility, integrity, and confidentiality of your data is your responsibility. Implement a comprehensive security strategy and follow good security practices.

This document is confidential.

Restricted Rights Legend

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

Copyright Notice

Follow this link to see a complete statement of copyrights and acknowledgements:
https://www.protect724.hpe.com/docs/DOC-13026

Support

Contact Information

Phone	A list of phone numbers is available on the HP ArcSight Technical Support Page: https://softwaresupport.hp.com/documents/10180/14684/esp-support-contact-list
Support Web Site	https://softwaresupport.hp.com
Protect 724 Community	https://www.protect724.hpe.com

Contents

Chapter 1: Introduction

The ESM High Availability Module (HA Module) provides for a backup ESM machine with automatic failover capability should the primary ESM machine experience any communication or operational problems. The HA Module is supported with ESM only and is not supported with other ArcSight software products. There are no separate configuration requirements to run the HA Module with ESM in FIPS mode vs Default mode.

The HA Module is installed on the primary of two adjacent machines connected by an Ethernet crossover cable. The HA Module replicates the installation and all data by mirroring the hard disk partition to the secondary machine.

The two systems each have an individual host IP address that is configured statically. In addition, you define a separate Service IP address that is used to identify the cluster. You will specify the Service IP during installation of the HA Module. During a failover, the HA Module reassigns the Service IP dynamically to the new primary system.

Ordinarily, one ESM instance runs on the primary machine and selected hard-disk writes are mirrored to the secondary machine. 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 no data is lost.

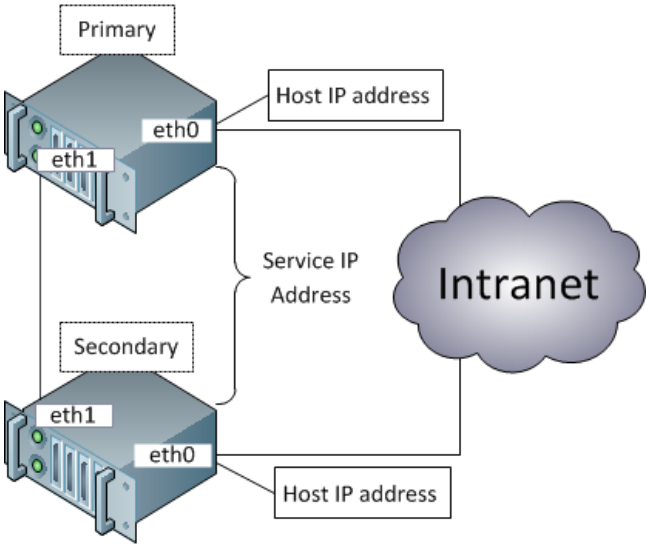

You will need to perform configuration set up tasks on both the primary and secondary systems before installing the HA Module. The order of steps that you perform will differ depending on whether both systems are new and without ESM installed, whether one of the systems has ESM installed, or whether you are upgrading both ESM and the HA Module. The goal of the configuration steps is to ensure that both systems are configured properly and that the configuration is aligned across these two systems.

You will install ESM and the HA Module on the primary system only. After installation is complete, a period of time will be required for the HA Module to sync the secondary system with the primary. In

general, new ESM installations take much less time than existing ESM systems because of the amount of data to be synced.

Chapter 2: Choosing Software, Hardware, and the Environment

This section describes the hardware, software, networking, and other requirements that are needed before the installation begins. This information will help you plan and prepare for the process of setting up the cluster systems and installing the HA Module. These are the requirements needed by the HA Module. Please see the *HP ArcSight ESM Installation and Configuration Guide* or the *HP ArcSight ESM Upgrade Guide* for the specific requirements to install or upgrade HP ArcSight ESM. You will use these documents together to plan your ESM and HA Module cluster installation.

The steps required to install the ESM High Availability Module are different for a new ESM installation than when upgrading an existing ESM to the latest version and installing the HA Module. After reading this section, read the section: "An Overview of the Installation Steps" on page 16 .

Refer to "An Example HA Implementation" on page 39 for a specific example an HA installation.

Important: If you already have ESM and are licensed for the a High Availability solution 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 a requires a different hardware configuration.

Place Systems in a High Availability Environment

The HA Module helps 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. The topic of designing a high availability network architecture is not the scope of this document. However, here are a few things that you can do independently of the HA Module to help ensure continued availability of ESM.

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

Hardware Requirements

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

- The systems can be either:
  two appliances (ESM Express or ESM Appliance).
  one appliance and one non-appliance system that supports the hardware requirements.
  two non-appliance systems that support the hardware requirements.

  Important: Appliances (ESM Express or ESM Appliance) have operating system and ESM installation scripts that run when the appliance is booted for the first time. If you're adding an appliance as the secondary system in the cluster, you must stop this process after the operating system is installed, but before ESM is installed. Do not install ESM on the secondary system. See the section "Configuring Systems before Installing the HA Module" on page 19for detailed information about how to perform this task.

- Running ESM with the HA Module requires significant disk space. There are minimum storage requirements of the cluster systems because of synchronization process. The ESM and archival storage must be on the same shared disk.

  Please see the *ESM Installation and Configuration Guide* for hard disk requirements required to run ESM. In addition to the ESM requirements, these additional storage requirements are needed to successfully install and run the HA Module.

Purpose	Minimum Storage	Note
ESM and HA Module installation binaries	3 GB	Ensure there is enough space for the downloaded installation binaries.
Temporary installation files	5 GB	Space required to run the Installation Wizard and the First Boot Wizard

Shared disk partition	Varies	This partition is mirrored 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, for Event Archive space.
HA sync metadata	Varies	The volume size depends on how large the ESM online storage is.

- You must use identical server class systems that support running either RHEL or CentOS.
- If the shared disks have write caches enabled, the write caches must be battery backed write caches (BBWC). If they do not have battery backup, there is a chance that the two disks will get out-of-sync when a power failure occurs.
- 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.
- The network interface cards should be at 1 Gigabit (Gb) or higher using a cable that supports this bandwidth.
- The network interface 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. This is useful when ESM is being upgraded on the primary system and has a significant amount of data that must be synchronized. See "Planning for the Initial Disk Synchronization" on page 15 for more detail about this process.
- If your servers have very high speed disk subsystems, you may 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.
- Set up the following disk partitions on both the primary and secondary systems.

Partitions	**Space required**	**Location**	**Notes**
Shared disk partition		Either /opt or /opt/arcsight	Recommendation only. You can alternatively, create an `/opt` or `/opt/arcsight` symbolic link to the physical location. This partition is mirrored between the primary and secondary.

Metadata partition	Determined by the size of the shared disk partition	/dev/<*sub_path*>	Contains disk synchronization metadata. This volume location should start with */dev*. The metadata partition size (in mebibytes) is calculated as: size = (P/32768) + 1 P = shared_disk_partition size in mebibytes
/ (root) partition	20 GiB (generous)		An operating system recommended partition.
swap	8 GiB (minimum)		An operating system recommended partition.
temp	6GiB (or more)	/tmp	An operating system recommended partition.

Notes about the shared disk partition:

- The contents of the shared disk on the secondary will be completely erased, so make sure it contains no data of value.
- Make sure that no process on the primary or secondary is using the shared disk file system.
- Bind mounts are not supported on the shared disk partition and are flagged as an error by the HA Module installation wizard. Use symbolic links instead.

Network Requirements

The following are the general requirements for the HA Module.

- You must set up at least one host on the network that is separate from the cluster systems (called a "Connected Host"). The HA Module will ping this host to check for network connectivity. You will specify the hostname or IP address of this connected host when running the First Boot Wizard during the HA Module installation.
- The two systems must be part of the same IPv4 subnet. The HA Module does not support IPv6.
- The primary and secondary machines must be close enough together that the network connection between them requires no intervening routers or switches.

- You will need to 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. 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 using the First Boot Wizard, which is run during installation of the HA Module. The Service IP address is dynamically reassigned to the system when a failover occurs and when the primary is brought back online.
- 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 reuse the ESM Manager SSL certificate, rather than regenerating a new certificate and importing it to all connectors and clients.
- We recommend you use DNS to manage IP addresses and host names for all the components in the cluster.
- The HA Module uses ports 694 and 7789 on each IP address in the cluster environment. These ports must be dedicated to HA Module communication only. Do not configure other applications to use these ports.
- The ports and protocols listed below are used by both systems and must not be blocked. Make sure that neither firewalld, nor iptables blocks the ports listed below. Set up your network firewalls to allow access to the Connected Hosts. A Connected Host is any other machine on the network that you have indicated can be pinged by the HA Module to verify that it is still on the network.

Protocol	Outgoing communication from...	On Port	Incoming communication to...	On Port
ICMP	■ the primary IP address ■ the secondary IP address	N/A	■ the primary IP address ■ the secondary IP address ■ the Service IP address ■ to the Connected Host	N/A
TCP	■ the primary crossover cable ■ the secondary crossover cable	Any	■ the primary system cable ■ the secondary system cable	7789

UDP	■ the primary IP address ■ the primary crossover cable ■ the secondary IP address ■ the secondary crossover cable	Any	■ the primary IP address ■ the primary crossover cable ■ the secondary IP address ■ the secondary crossover cable	Any

Using the Service IP 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 Manager IP address and Host Name during ESM installation, you will specify the Service IP address and not an individual host IP address. When the ArcSight Console connects to the Manager, it will use the Service IP address. The ArcSight Command Center URL will specify the Service IP address. When a fail over occurs, the Service IP address will be dynamically assigned to the new primary system. Other than specifying the Service IP address when installing the HA Module and ESM, and assuring 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 used by the host IP addresses.

Software Requirements

- The supported operating system is specified in the ArcSight ESM Support Matrix.
- The cluster systems must run either RHEL or CentOS. Both systems must have the same operating system and version installed. RHEL 7.1 and CentOS 7.1 are supported with new ESM installations only.

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

- Both systems must be configured to access a Yum repository which is needed to install dependencies required by the HA Module. This can be either a remote Yum repository provided by the operation system (OS) vendor, a repository created from the OS ISO or CD, or a directory location on the local system. Please see the vendor-specific documentation for information about configuring Yum and connecting to Yum repositories. If you're using ESM on an appliance, you will accomplish this using instructions in "Configuring Systems before Installing the HA Module" on page 19 to install the packages in `esm_ha_support_pkgs.tar.gz`.

- We strongly recommend 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 if you use physical disk management. Please note, an LVM partition must be a multiple of the LVM chunk size. If you use 32 MiB for the chunk size, then to get a 33 MiB partition, you would create a 64 MiB partition, because you would need two chunks. See "Disk Partition Setup" on page 41 for an example of how to do this.

- Download the compatible ESM and HA Module installation files from the HP SSO download site to the primary system and unpack the tar file. Do NOT place the installation binary or content in the shared directory (generally */opt/arcsight*), because they will be deleted during the installation process. You will install ESM and the HA Module on the primary system only. After installation is complete, the HA Module will synchronize the secondary system with the primary system. The installation files are:
  `ArcSightESMSuite-6.9.1.xxxx.tar` file
  `ArcSight-Highavail-6.9.1.xxxx.0.bin` file
  `esm_ha_support_pkgs.tar.gz` file: required only if you are migrating an appliance (ESM Express Appliance or ESM Appliance) from a single installation to an HA Module cluster

- The HA Module version, ESM version, and operating system version must be compatible. See the "ESM Support Matrix" for a summary of the ESM, HA Module, and Operating System version compatibility.

- ArcSight Risk Insight supports English only. If you are running ArcSight ESM with Risk Insight, the ArcSight ESM components must be set to support this requirement.

- If you plan to run ArcSight Risk Insight on the HA Module cluster, see the ArcSight Risk Insight *Technical Note: Setting Up Risk Insight for HA*. Certain installation and configuration requirements are required to support this deployment scenario. See the ArcSight Risk Insight Support Matrix for the details about supported platforms. These documents are available on the Protect 724 site.

Important: The order in which you install the software components will be different depending on the installation scenario:

- Installing the HA Module for the first time and upgrading the ESM version
- Installing the HA Module for the first time without upgrading ESM
- Installing ESM and the HA Module for the first time
- Upgrading both the ESM version and HA Module version

See "An Overview of the Installation Steps" on page 16 for more detail about these scenarios and the order in which you will install the software components.

Getting the License File

The license file for the HA Module is an ESM license file with the HA Module included. If you have ESM installed without HA, obtain a new ESM license that includes the HA Module. After upgrading ESM,

install the new ESM/HA Module license as described in the ESM Administrator's Guide, Chapter 2, "Configuration." The topic is "Installing New License Files Obtained from HP."

If ESM is not already installed, you will specify the same ESM/HA Module license file when you install the HA Module and then again when you install ESM. Refer to the *ESM Installation and Configuration Guide* for detailed information about installing ESM. For HA Module installation, see "Running the HA Module Installation Wizard" on page 25.

If you are upgrading from ESM 6.8 and HA Module 1.0, you do not need a new license file.

Planning for the Initial Disk Synchronization

NOTE: This information does not apply if you are installing the HA Module and ESM on a new system where ESM has not yet collected event data.

After HA Module is installed on a existing ESM system, the entire shared disk partition on the existing ESM primary system must be synchronized 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 during this time, but the secondary system is not 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 1G cable. In such cases, providing a 10G interface may lead to noticeable reductions in the time required for the initial synchronization.

SSD drives (Fusion, for example) contribute to improving the synchronization speed because they are fast. SSD drives require and support TRIM to manage free space. The HA Module disk synchronization process is TRIM-aware; it 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 TRIM information is passed to the SSD drives by way of the disk synchronization process. 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.

Chapter 3: An Overview of the Installation Steps

There are differences in the set up process for each system depending on whether you are either 1) installing the HA Module for the first time on an existing system and upgrading ESM to the correct version, 2) installing the HA Module for the first time on an existing system without upgrading ESM, 3) installing ESM and the HA Module for the first time on a clean system, or 4) upgrading both ESM and the HA Module.

If you are installing ESM for the first time on the primary system, you will also need the *HP ArcSight ESM Installation and Configuration Guide*. If you are upgrading ESM on the primary system, you will need the *HP ArcSight ESM Upgrade Guide.* Please download those documents from the Protect724 Customer Portal at https://protect724.hp.com/community/arcsight/productdocs/esm.

Choose one of the following installation scenarios for more detail:

- "Installing the HA Module on an existing ESM installation" below
- "Installing ESM and the HA Module for the first time" on the next page
- "Upgrading both the ESM version and HA Module version" on page 18

Installing the HA Module on an existing ESM installation

Follow this process if the ESM instance is running on a single system and you want to convert the installation to an HA Module cluster. Make sure that the ESM version is compatible with the HA Module. If the ESM version is not compatible, follow the *HP ArcSight ESM Upgrade Guide* to upgrade the system before continuing.

The general steps to follow are:

1. Set up a secondary system that has equivalent hardware to the existing primary system. Review the software, hardware, and configuration requirements to ensure that the HA Module will run successfully on the cluster systems. See "Choosing Software, Hardware, and the Environment" on page 8. Make sure that you also:
   - Replace the original ESM license with a new license that enables both ESM and the HA Module.
   - Install the "ArcSight ESM HA Monitoring" Foundation Package.
2. On the primary system, as the user *root* stop ESM by running:

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

3. (optional) This step will enable you to reuse the ESM Manager SSL certificate, rather than regenerate a new certificate. The general approach is to give the existing system a new IP address, and then re-use the original IP address as the cluster's new Service IP address. The detailed steps to perform on the existing system are:
   a. Add an new IP address to the interface that has the current host IP address. The new IP address will be the host's new IP address. The original IP address will become the Service IP address that identifies the cluster.

   b. Setup /etc/hosts or DNS to resolve the new host IP to the new hostname.

   c. Configure the host so it uses the new hostname.

      Now that the original IP Address has been removed from the network interface, you can re-use it as the cluster's Service IP Address when you run the First Boot Wizard.

      NOTE: If you change the system hostname during installation, test that the change persists across reboots. Reboot the system, and then use the `hostname` command to show the system hostname.

4. Complete the steps in "Configuring Systems before Installing the HA Module" on page 19 on both systems. Some of the pre-installation configuration steps will be similar to those described in the *HP ArcSight ESM Installation and Configuration Guide*. They are described in this document because of their importance to the HA Module installation process.

5. On the primary system, install the HA Module, as described in "Running the HA Module Installation Wizard" on page 25. After you install the HA Module, it will begin synchronizing the ESM instance on the primary system with the secondary system. This is called the initial synchronization step.

6. Complete the post-install steps in "Verifying the Systems After the Installation or Upgrade" on page 33.

Installing ESM and the HA Module for the first time

Follow this process if you are installing both ESM and the HA Module on new systems. Perform the installation steps in the following order:

1. On both the primary system and secondary systems, install the supported operating system.

2. Complete the steps in "Configuring Systems before Installing the HA Module" on page 19 on both systems. Some of the pre-installation configuration steps will be similar to those described in the *HP ArcSight ESM Installation and Configuration Guide*. They are described in this document because of their importance to the HA Module installation process.

3. On the primary system, install the HA Module as described in "Running the HA Module Installation Wizard" on page 25.

4. Login as *root* and create the folder `/opt/arcsight`. Set the ownership to user *arcsight*.

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

   This change is mirrored to the secondary system after the HA Module is installed, assuming your mount point for the mirroring is either `/opt` or `/opt/arcsight`.

5. On the primary system, install ESM. See the *HP ArcSight ESM Installation and Configuration Guide* for details.

   IMPORTANT: The HA Module must be running before you begin installing ESM.

   IMPORTANT: When the ESM Configuration Wizard asks you for the Manager Host Name or IP address, enter the cluster Service Host Name or Service IP Address and NOT the host name of a single machine.

   In the ESM Configuration Wizard, be sure to include the Foundation Package called “ArcSight ESM HA Monitoring,”. Installing this package with ESM is required if you want to acquire up-to-date HA Module status information from the outset. If you activate this package later from the ArcSight Console, there is no status information available until an HA event occurs, which could be a long time.

6. Complete the post-installation steps described in "Verifying the Systems After the Installation or Upgrade" on page 33.

Upgrading both the ESM version and HA Module version

Follow this process if ESM 6.8c and HA Module 1.0 are installed on the primary and secondary cluster systems and you are upgrading to ESM 6.9.1 and HA Module 6.9.1. You will upgrade ESM, the HA Module, and operating system in a specific order. See the "Upgrading ESM and the HA Module" on page 31 section for details about this process.

NOTE: If ArcSight Risk Insight is running on the cluster systems, see the ArcSight Risk Insight Technical Note: Upgrading Risk Insight 1.0 to 1.1 for platform and configuration details related to upgrading ESM, HA Module, and Risk Insight. This document is available on the Protect 724 site.

Chapter 4: Configuring Systems before Installing the HA Module

The primary and secondary machines must be set up so that they are nearly identical. The following steps must be performed on both primary and secondary systems to ensure that they are configured properly to run the HA Module. The configuration on both primary and secondary systems must be identical. The HA Module installation scripts will check for configuration dependencies and return an error message if dependencies are not met.

TIP: Some of the steps describe in this section have been automated by the prepareHA.sh script that comes packaged with the ESM 6.9.1 Installation Package. However, the prepareHA.sh script does not perform all of the configuration steps described below. The steps that are not automated must be performed manually. See the section "The prepareHA Script " on page 71 for detailed information about this script.

- If you are converting an appliance (ESM Express Appliance or ESM Appliance) as secondary system in the cluster, disable the certain first boot scripts. New systems that are configured to run ESM on an appliance include set up scripts that run at first boot to install the operating system and ESM. If you're using an appliance to be the secondary in the cluster, you must not install ESM on that system. The scripts can run normally on an appliance that is intended to be the primary. Perform the following steps on the appliance to both stop the scripts after the operating system is installed and prevent ESM from being installed in the future.
  a. When you boot the secondary machine (the appliance), let the operating system installation script run normally.

  b. At the prompt: "InstallAnywhere will guide you through the installation of ArcSight ESM 6.9.1c Suite", enter the text: `quit`

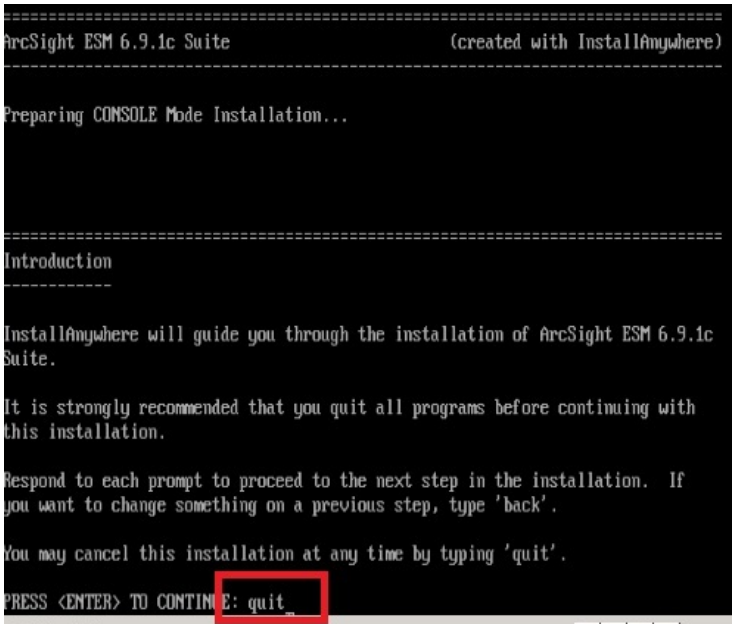

c. To prevent the appliance from installing ESM after it is rebooted, edit the .bash_profile in the root user's home directory. Remove the lines from "`# run OS configuration and ESM installer if not yet done`" to the end of the file.

d. Reboot the appliance to ensure that the ESM installation script does not run.

- When defining host names for each system, do not use capital letters in either name. You may encounter unexpected results if there are capital letters in either the primary hostname or the secondary hostname.
- Make sure that both systems have the latest operating system timezone package installed. This is a requirement for ESM. For instructions, refer to the *ESM Installation and Configuration Guide*, specifically the topic "Install Time Zone Package" in the "Installing ESM" chapter.
- Set up both primary and the secondary systems to run the Network Time Protocol (NTP) so that the system time is kept synchronized between them.
- If the mirrored disks are SSD drives, such as Fusion, make sure you have TRIM support configured on both the primary and secondary systems.
- If the operating system is RH/CentOS7.1, run the following commands on both the primary and secondary systems to create a symbolic link, called */usr/lib64/libpcre.so.0*, to the file */usr/lib64/libpcre16.so.0*.

  ```
  cd /usr/lib64
  ln -s libpcre16.so.0 libpcre.so.0
  ```

- Connect the two servers with crossover cables. Configure the interfaces with IPv4 addresses.

- If you are converting an appliance (ESM Express or ESM Appliance) single installation to an HA Module cluster installation:
  a. Run the following commands on both appliance systems.
  ```
  mkdir -p /usr/local
  mv -f /opt/hp /usr/local
  ln -s /usr/local/hp /opt
  ```
  b. And then, on the appliance that will be the secondary, delete all files under `/opt` except `/opt/hp`.

- On both the primary and secondary systems, select the partitions to be mirrored between the two servers.

  Typically, this is the partition mounted as `/opt` for your ESM installation. Use the command `df /opt/arcsight` to obtain the partition. This partition must exist on both the primary and secondary and must have the same device name, be mounted at the same location, and be the same size. If the partition is not mounted as `/opt` or `/opt/arcsight`, then create a symbolic link from `/opt` or `/opt/arcsight` on both the primary and secondary. Note that after installation, this partition is only mounted on the primary. Only that primary can make changes to it.

- Make sure all file system options are set up the way you want them on the primary system. The HA Module will mount the file system on the secondary exactly the way you mounted it on the primary system.

- On both the primary and secondary systems, create a metadata partition. This is a small partition on each server used for disk-synchronization metadata. The size to allocate for each partition is calculated in mebibytes (1 MiB=1,048,576 bytes):

  ```
  size=(P/32768)+1
  ```

  where P is the size of the shared disk partition in mebibytes (1 MiB=1,048,576 bytes). For example, if the shared disk partition size is 1 TiB (that is, 1,048,576 MiB), the metadata partition size would be 33 MiB.

  See "Disk Partition Setup" on page 41 for an example of how to do this. If you ever increase the size of the shared disk partition, be sure to increase the size of the metadata partition accordingly. Decreasing the size of the mounted partition is not supported.

  If the systems in the cluster are two appliances (ESM Express or ESM Appliance), then skip this step. The meta data partition already exists on each system.

  If the systems in the cluster are a non-appliance system and an appliance, you will set up the partition on the non-appliance system only. The appliance does not use LVM, so you cannot use LVM on the non-appliance system. Check the partition configuration on the appliance system and configure the non-appliance system with the exact same specifications.

  If the systems used in the cluster are both non-appliance systems, you must set up the metadata partitions on both systems as described above.

- Make sure the password for the *root* user is the same on both systems. This is required during the HA Module installation process. You may change the *root* passwords after installation.
- Make sure the *arcsight* user exists on both systems. This user must be configured to use the same home directory, UID, and GID on each system. If ESM is already installed on the primary, get the UID and GID for *arcsight* user, and then create the same configuration on the secondary system. Do this as user *root*:

  ```
  id arcsight
  ```

  Note: The user id, group id and home directory must be the same on both machines or the installation will fail.

  Use the following command to create the group and user on the second system, replacing <GID> and <UID> with the values returned from the previous command.

  ```
  groupadd -g <GID> arcsight

  useradd -c "arcsight_esm_owner" -g arcsight -d /home/arcsight -m -s /bin/bash -u <UID> arcsight
  ```

  If neither system has ESM installed, run these commands as user *root* to create a user and group called *arcsight*. Set the home directory, <GID>, and <UID> to be identical on both systems.

  ```
  groupadd –g 500 arcsight

  useradd -c "arcsight_esm_owner" -g arcsight -d /home/arcsight -m -s /bin/bash -u 500 arcsight
  ```

  Note: If 500 is not available for user and group ids, pick any free ID. The user id, group id, and home directory must be the same on both machines or the installation will fail.

  Change the password for user *arcsight*:

  ```
  passwd arcsight
  ```

  The passwords for the *arcsight* user can be different on each system.

- If you will install the HA Module for the first time, login as the user *root* on the primary system and create the folder `/usr/lib/arcsight`. Make the user *arcsight* the owner:

  ```
  mkdir /usr/lib/arcsight

  chown arcsight:arcsight /usr/lib/arcsight
  ```

- On both the primary and secondary systems, increase the user process limit as follows:
  a. In the `/etc/security/limits.d` folder, check whether you have this file specific to your operating system:

called either `20-nproc.conf` or `90-nproc.conf`

If you do not already have this file specific to your operating system, create it. Make sure to create the `limits.d` directory, if it does not exist either.

b. If the file specific to your operating system does exist, then delete all entries in the file and add these lines:

```
*    soft    nproc    10240
*    hard    nproc    10240
*    soft    nofile   65536
*    hard    nofile   65536
```

Note: Be sure to include the asterisk (*) in the new entries. It is important that you add all of the entries exactly as specified. Any omissions can cause system runtime errors.

c. Reboot the machine.

d. Log in as user *arcsight*.

e. Run the following command to verify the new settings:

```
ulimit -a
```

f. Verify that the output shows the following values for Open files and Max user processes:

```
open files 65536
max user processes 10240
```

- If the systems have been upgraded from 6.9.0 to 6.9.1, remove the reference to the system host name in the `/etc/hosts` file. In the following example, remove the text *host01.mycompany.com*. Leave the remaining text as is.

```
127.0.0.1 host01.mycompany.com localhost localhost.localdomain localhost4
localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain
```

- Both systems must be able connect to each other via SSH using the “root” user account with password authentication. If the SSH permissions are locked down to prevent connecting via SSH as user root, the installation will fail.

  Test that you can connect via SSH to localhost as user `root` on both systems, and then also that 1) the primary can connect to the secondary and 2) the secondary can connect to the primary via SSH as user root. Password authentication must be enabled and you should be prompted for a password. For example:

```
ssh root@localhost
ssh root@node1.domain.com
ssh root@node2.domain.com
```

The SSHD Config (/etc/ssh/sshd_config) on each system contain the lines:
```
PermitRootLogin yes
PasswordAuthentication yes
```

- If there's an SSH login banner enabled on the system for the "root" account, it must be disabled or the installation may fail. To disable the SSH login banner for the *root* user,

  Create an empty ".hushlogin" file in the *root* home directory.
  ```
  # touch /root/.hushlogin
  ```

  Open the /etc/ssh/sshd_config file, comment out the "Banner" line, and then restart the sshd service.

  ```
  # vi /etc/ssh/sshd_config
  # /etc/init.d/sshd restart
  ```

  Confirm that the login banner is disabled for the "*root*" account on both systems:

  ```
  # ssh root@localhost
  root@localhost's password:
  root@apslp0389:/root
  ```

  Note: If another process or script modifies the `root` user login banner, make sure that you disabled that process or code.

- If you are migrating an appliance (ESM Express Appliance or ESM Appliance) from a single installation to an HA Module cluster, perform the following steps on each appliance before you install the HA Module. These steps install supporting packages that are required to run the HA Module:
  a. Login as the `root` user on the appliance.

  b. Copy the `esm_ha_support_pkgs.tar.gz` file to the `/tmp` partition. You downloaded this file from the HP SSO download site in an earlier step.

  c. Run following commands to install the supporting packages:
  ```
  cd /tmp

  tar zxvf esm_ha_support_pkgs.tar.gz

   cd install

   ./install_ha_support_pkgs.sh
  ```

Chapter 5: Running the HA Module Installation Wizard

This section describes how to run the ESM High Availability Module installation wizard and First Boot Wizard.

It is assumed that you have already completed all the required tasks for the primary and secondary machines as described in "Configuring Systems before Installing the HA Module" on page 19.

You can run the installation wizard and First Boot Wizard in either console mode (via the command line) or GUI mode (using X Windows). The First-Boot Wizard enables you to configure the HA Module.

When installing the HA Module in GUI mode, the First Boot Wizard starts automatically when the installation wizard finishes, so it appears to be a seamless operation. You can also run the first boot wizard independently at any time to make changes to the HA Module configuration.

Upon completion of the First Boot Wizard prompts, a script is invoked to check that system configuration is complete and correct, and then reports inconsistencies and the location of logs to help you fix the issues. If there are no inconsistencies, the First Boot Wizard completes with the specified configuration.

It is important that the two systems match with respect to hardware, installed 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 that you can correct the inconsistency, re-run the First Boot Wizard, and finish the installation.

Running the HA Module Installation Script

These steps will be performed on the primary system only. To run the installation wizard:

1. Download the HA Module installation binary file, `ArcSight-Highavail-6.9.1.xxxx.0.bin` to the `/home/arcsight` directory. Do not download it to the shared partition.
2. Download the ESM license file that includes the HA Module.
3. Log in as the arcsight user and run the installer in either GUI or console mode. The installation prompts for each modes are comparable. Console mode provides a text-based interface. To run the installation file, execute either:

   `/home/arcsight/ArcSight-Highavail-6.9.1.<xxxx>.0.bin -i console` for console mode

   or

   `/home/arcsight/ArcSight-Highavail-6.9.1.<xxxx>.0.bin` for GUI mode

4. At the **Introduction** prompt, either click Next (GUI) or press Enter (console).
5. If HA Module 1.0 is installed on the system, you will be prompted to 'Press [YES] to upgrade this installation, Press [Quit] to cancel this installation.' If you are upgrading the HA Module, and running the script in Console mode, enter 1 and then press Enter.
6. At the **License Agreement** prompt,

   In GUI mode, scroll down and then select the **I accept the terms of the License Agreement** radio button to agree to the license agreement. The radio button is grayed out until you scroll to the bottom of the license agreement. Click Next.

   In Console mode, press Enter at each prompt to scroll to the end of the license agreement. Enter Y, and then press Enter accept the terms of the license agreement.
7. If you ran the installation wizard in console mode, a **Pre-installation Summary** screen will appear. Press Enter to continue.
8. The installation script will run and display progress in the window.
9. If you ran the installer in console mode, you will be prompted to perform next steps when the installer is complete. If you ran the installer in GUI mode, the First Boot Wizard will start automatically.

Running the First Boot Wizard

1. If you ran the installation wizard in GUI mode, the First Boot Wizard should appear automatically and you can skip to the next step.

   To run the First Boot Wizard, change user to *arcsight*, and then type:

   ```
   cd /usr/lib/arcsight/highavail/bin
   ./arcsight firstBootWizard [--console]
   ```

   If you specify `--console`, the First Boot Wizard runs in console mode. If not specified, it runs in GUI mode. Unless otherwise noted, all fields are required and omitting them will result in an error message.
2. At the **Welcome to the First Boot Wizard** prompt, click Next (GUI) or enter **yes** (console).
3. On the **License File** prompt, enter the full path to your ArcSight license file (either the zip file or the .lic file that is unpackaged from the zip)

   - In GUI mode, click the browse button (**...**) and navigate to the directory to which you downloaded the license file for the HA Module and select it.
   - In console mode, enter the full path to the file.

Click **Next** (GUI) or enter **yes** (console) to continue.

4. The **Properties File** prompt can be used if you are upgrading from HA Module 1.0 to HA Module 6.9.1. This prompt simplifies the First Boot Wizard process by enabling you to load the highavail.properties file that defines the cluster configuration.

   Enter the location of the HA Module highavail.properties. It is often found at:
   `/usr/lib/arcsight/highavail/highavail.properties`

5. In the **Parameter Configuration** prompt, enter the requested information, and then click Next (GUI) or enter **yes** (console) to continue.

Field	Description
Shared Disk	Enter the mount point of the disk shared between the primary and secondary (for example, '/opt'). The options provided include all relevant mount points. Bind mounts are not supported and are flagged as an error by the installation. Use symbolic links instead. The contents of the shared disk on the secondary will be ***completely erased***, so make sure it contains no data of any value. Make sure no process on the primary or secondary is using this filesystem or the installation will exit with errors. You cannot change this value on subsequent runs of the First Boot Wizard. This value is identified in the highavail.properties file as shared_disk.
Metadata Volume	Enter the volume containing disk-synchronization metadata. This volume is expected to start with "/dev". The contents of the metadata volume on both the primary and the secondary will be removed. This value is identified in the highavail.properties file as metadata_volume.
Service Hostname	Enter the service hostname assigned to the service IP. This is a virtual hostname that is used to connect to the cluster regardless of which physical computer is the acting as the primary system. The service IP address can also be used, but we recommend using the service hostname. You can use the `hosts` file or DNS to resolve the hostnames. This value is identified in the highavail.properties file as service_hostname.

Field	Description
Secondary Hostname	Enter the hostname of the secondary machine. This is the host name assigned specifically to the machine. This value is identified in the highavail.properties file as secondary_hostname.
Primary Cable IP	Select the IP address of the interface connected to the interconnect cable on the primary system. This value is identified in the highavail.properties file as primary_cable_ip.
Secondary Cable IP	Enter the IP address of the interface connected to the interconnect cable on the secondary system. This value is identified in the highavail.properties file as secondary_cable_ip.

Click **Next** (GUI) or press Enter to continue .

6. At the second **Parameter Configuration** prompt, enter the following information:

Field	Description
Preferred Primary	Enter the hostname of the machine you would prefer to be primary. This field is not required. Leave it blank if either machine may be the primary. If you supply a value, HA will always fail over from the non-preferred primary to the preferred primary when both machines are running properly. This may cause additional, unneeded failovers.
Connected Hosts	These hosts are other machines in the network that HA can ping to verify that it is connected to the network. Enter a space-separated list of hostnames or IP addresses that can be pinged. Do not enter any hostname or IP address for either the primary or the secondary machines. This field is not required. If you leave it blank there is no automatic failover if the primary loses contact with the network.
Connectivity-Down Timeout	Specify the time to wait, in seconds, before initiating a failover due to lack of internet connectivity on the primary. The default is 180.
Ping Timeout	Specify the seconds to wait before considering that a ping has failed. The default is 2 seconds.
Ping Attempts	Specify the number of pings to attempt before considering that the pings have failed. The default is 2 pings.

For more information in how these settings affect Failover, see "An overview of the Failover-Check Operation" on page 73.

7. At the "root password" prompt, enter the password for user *root*, and then either click **Next** (GUI) or enter **yes** (console).

   Supplying the password for the *root* user enables the HA configuration script to handle components and actions that have to be performed as the *root* user. The password must be the same on both machines. This password is not stored permanently. You may change this password after the installation completes.

8. If you are running on console mode,
   a. you are prompted about whether to hide the input for private parameters from the screen. Press Enter to hide these parameters.

   b. you are prompted to verify the root password.

9. If your shared disk is empty, you be prompted with additional information about the duration of the remaining processes. Click **Next** (GUI) or press **Enter** (console) to continue.

   In GUI mode, the dialog remains open while the installation script runs. This may take an hour or so depending on whether you are upgrading an existing ESM.

10. When the First Boot Wizard is finished, it displays the "First Boot Wizard is done" dialog (GUI) or "Installation Result" prompt (console) and shows any relevant messages. Click **Next** (GUI) or enter **yes** to complete.

11. In console mode, enter **yes** to return to the command prompt.

12. If there are errors, check both servers for log files. See "Installation Issues and Solutions" on page 61.

    Fix any errors noted in these logs and then re-run the First Boot Wizard by running the following command as user *arcsight*:

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

Verifying the HA Module Installation

Run the following commands as user *root* to check to see that the HA Module is running properly after installation.

1. On the primary system, check the shared disk partition:

   ```
   df -h /dev/drbd1
   ```

   Sample output is shown below:

   ```
   Filesystem Size Used Avail Use% Mounted on
   /dev/drbd1 1.6T 197M 1.6T 1% /opt
   ```

This shows that the shared disk partition (`/dev/drbd1`) is mounted on `/opt`. It should be mounted on the shared disk directory entered into the First Boot Wizard.

NOTE: You may notice that the shared disk partition on the secondary system is unmounted. This is normal. This configuration was changed by the HA Module installation scripts. Do not remount the shared disk partition on the secondary system after running the HA Module installation wizard.

2. From the directory `/usr/lib/arcsight/highavail/bin` run:

```
./arcsight_cluster status
```

Example output is shown below:

```
Tue Dec 16 14:04:04 PST 2014 OK
prod01.acme.com: online
prod02.acme.com: online Primary

Disk: Connected UpToDate/UpToDate

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

Started Failover-Check-prod01.acme.com
Started Failover-Check-prod02.acme.com
Started Filesystem
Started Ping-prod01.acme.com
Started Ping-prod02.acme.com
Started STONITH-SSH-prod01.acme.com
Started STONITH-SSH-prod02.acme.com

Started Service-IP
```

Check to be sure that the first line ends with `OK`. This indicates the cluster is running normally.

Chapter 6: Upgrading ESM and the HA Module

This information guides you through the process of upgrading both ESM and the HA Module in an environment where you have HA Module 1.0 and ESM 6.8c running on a two system cluster. The HA Module version, ESM version, and operating system version must be compatible. See the "ESM Support Matrix" for a summary of the ESM, HA Module, and operating system version compatibility. Upgrade process is supported on RHEL 6.7 and CentOS 6.7 only. RHEL 7.1 and CentOS 7.1 are supported with new installations only.

You will upgrade ESM and the HA Module on the primary system only. After upgrade is complete, the HA Module will synchronize the secondary system with the primary system.

1. Download the compatible ESM and HA Module installation files onto the primary system. Do NOT place the installation binary or unpacked content on the shared disk partition (generally */opt/arcsight*), because they will be deleted during the upgrade process.
   - Download the `ArcSightESMSuite-6.9.1.xxxx.tar` file and unpack it.
   - Download `ArcSight-Highavail-6.9.1.xxxx.0.bin` file to the `/home/arcsight` directory
2. Make sure you have access to the *HP ArcSight ESM Upgrade Guide*. This explains the detailed steps to upgrade ESM.
3. On the secondary system:
   a. As *root* user, run the command `/sbin/service heartbeat stop`.
   b. As the *root* user, run the command `/sbin/chkconfig --del heartbeat`.
   c. Setup the Yum repository files to point to the repository for the supported operating system version.
   d. Run the `yum update` command to upgrade the operating system.
   e. Reboot the secondary system.
4. On the primary system:
   a. As the `root` user, run the command `Tools/stop_services.sh` to remove the ArcSight Services. The Tools directory is in the location where you unpacked the ESM installation tar file. Note: ESM will go down at this point and will stay down until after you upgrade ESM and restart the services.
   b. As the `root` user, run the command `/sbin/service heartbeat stop` .
   c. As the `root` user, run the command `/sbin/chkconfig --del heartbeat`.

d. Setup the Yum repository configuration files to point to the repository for the supported operating system version

e. Run the `yum update` command to upgrade the operating system.

f. Reboot the primary system.

5. Make sure all configuration described in the section "Configuring Systems before Installing the HA Module" on page 19 is complete.

6. On the primary system, check that the cluster service is stopped by running the command `service heartbeat status` as the arcsight user.

7. On the primary system, execute the file `ArcSight-Highavail-6.9.1.xxxx.0.bin` to run the HA Module Installation Wizard. The file will ask if you want to upgrade. Either enter Yes, or select Yes, at the prompt. See the section "Running the HA Module Installation Wizard" on page 25 for detail about the running the HA Module Installation Wizard.

8. On the primary system, as the `root` user, run the command `/usr/lib/arcsight/highavail/install/upgrade.sh` to upgrade the HA Module. The upgrade script asks if you want to continue with the upgrade. Enter, or select, Yes at this prompt to complete the HA Module upgrade. The log file for the HA Module upgrade is located at: */usr/lib/arcsight/highavail/logs/upgrade.log*.

9. On the primary system, upgrade to the supported ESM version. See the *HP ArcSight ESM Upgrade Guide* for detailed instructions about upgrading ESM. Because you have already performed this step, you do not need to run the command `Tools/remove_services.sh` to remove the ArcSight services.

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

10. After the ESM upgrade is complete, the primary system should be running ESM. The HA Module will begin synchronizing the primary system and the secondary system.

11. As the root user, start the ArcSight services by executing the command:

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

12. Check that the ArcSight services are running by executing the command:

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

13. If you have not already done this, activate the "ArcSight ESM HA Monitoring" Foundation Package from within the ArcSight Console. See the *ArcSight Administration and ArcSight System Standard Content Guide* for instructions about activating standard content.

Chapter 7: Verifying the Systems After the Installation or Upgrade

The tasks that you perform after setting up the ESM and HA Module are determine by whether ESM and the HA Module were newly installed or were upgraded. Complete one of the following subsections. The steps in "Setting Configurable HA ModuleProperties" on page 36 are optional and enable you to tune the configuration if needed.

You installed the HA Module on an existing ESM installation

In this scenario, the ESM instance is running on a single system and you converted the installation to an HA Module cluster. Now is the time to switch to the new Service hostname or Service IP Address. Perform these steps on the primary system.

1. On the primary system, set up the ESM services by running this command as user *root*:

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

   It will automatically detect the HA Module and make appropriate changes to both the primary and the secondary.

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

   ```
   fstrim <shared disk>
   ```

   If the drive has a large amount of free disk space, this command dramatically shortens the time to synchronize the secondary disk.

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

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

   NOTE: You can skip steps 4-10 if you changed the original single system hostname and are now using the original IP as the Service IP for the cluster.

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

   ```
   ./arcsight managersetup
   ```

   This opens the Manager's setup wizard.

a. When prompted by the Manager setup wizard for the Manager Hostname, and in every field where the previous Hostname or IP address is displayed, enter the cluster Service Hostname or cluster Service IP Address (use the same value that you set in the First Boot Wizard).

b. 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. If ESM is configured for FIPS mode, this step has to be performed manually on the command line. Check the ESM Administrator guide for information about Generating a Key Pair.

5. Start the Manager by running the following command as user *arcsight*):

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

6. As the user *arcsight*, check that ArcSight Manager is running using the following 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”.

7. Make sure you can start the ArcSight Command Center by browsing to the following URL:

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

   Where `<Service Hostname>` is the hostname defined for the cluster. Hostnames with underscores do not work on IE, so use the Service IP address. If you are not using DNS to resolve host names, use the Service IP address instead.

8. Change the Manager IP (to the cluster Service IP) for every connector and Console that connects to this Manager. Change any URLs (for example bookmarks) to ArcSight Command Center.

9. Import the Manager’s newly-generated certificate on all clients, ArcSight Console and connectors, that access the Manager. Use keytoolgui. Keytoolgui is described in the *SSL Authentication* chapter of the ESM Administrator’s Guide for details. If you're using FIPS configuration, use the runcertutil utility, described in the ESM Administrator’s Guide.

10. Test to make sure that:

    - The clients can connect to the ArcSight Manager using the Service IP Address or Service host name.
    - Peer configuration works as expected. If not, redo the peer configuration.

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

11. If you have not already done this, activate the "ArcSight ESM HA Monitoring" Foundation Package from within the ArcSight Console. See the *ArcSight Administration and ArcSight System Standard Content Guide* for instructions about activating standard content.

You installed ESM and the HA Module for the first time

No additional configuration is required for the cluster set up. Make sure that you have performed the ESM-specific post-installation configuration, see the *ESM Installation and Configuration Guide*, specifically the chapter titled "Post-Installation Considerations". After the ESM post-installation configuration is complete:

1. On the primary system, check that all ArcSight services are running using the command:

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

   You should see a list of services and the status of each.

2. During the HA Module installation, the cluster is started automatically by starting heartbeat service. Check the cluster status using the arcsight_cluster script command:

   ```
   ./arcsight_cluster status
   ```

   The arcsight_cluster script was installed in the /usr/lib/arcsight/highavail/bin directory. See the section "The arcsight_cluster Script " on page 45 for details about the command arguments available.

You upgraded both ESM and the HA Module

No additional configuration is required for the cluster set up. For a list of ESM-specific post-upgrade configuration, see the*HP ArcSight ESM Upgrade Guide*. On the primary system, check that both ESM and the HA Module services are running.

1. On the primary system, check that all ArcSight services are running using the command:

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

   You should see a list of services and the status of each.

2. During the HA Module installation, the cluster is started automatically when starting heartbeat service. Check the cluster status using the arcsight_cluster script command:

   ```
   ./arcsight_cluster status
   ```

   The arcsight_cluster script was installed in the /usr/lib/arcsight/highavail/bin directory. See the section "The arcsight_cluster Script " on page 45 for details about the command arguments available.

Setting Configurable HA ModuleProperties

There are three ESM properties relevant to HA that you can configure. The properties are in `/opt/arcsight/manager/config/server.properties`.

`highavailability.monitor.on=true`
This property turns the HA Notification feature on or off. Use `false` to turn off notifications.

`highavailability.notification.interval=300`
This property sets the notification interval for failure conditions. It is configured in seconds and the default is 300 seconds (five minutes). Users get an email, audit event, and subsystem change console pop-up at the specified interval.

`whine.check.interval.HASubsystemChecker=30`
This property sets the polling interval of the tracker/checker that checks the `/usr/lib/arcsight/highavail/status.txt` file. It is configured in seconds and the default is 30 seconds.

If you change any of these properties, restart the ArcSight Manager for them to take effect. For more information about editing ESM properties files, refer to the “Configuration” chapter of the ESM Administrator’s Guide.

Chapter 8: Uninstalling Software Components

The HA Module uninstallation process can be done either with or without uninstalling ESM.

Uninstalling both ESM and HA Module

1. On the primary server, uninstall ESM using the ESM uninstallation instructions in the ESM Installation and Configuration Guide.
2. On the primary server, run the following HA Module uninstall script as user *root*:

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

   It will ask you if you really want to do the uninstall. if you say yes, the uninstall will be completed on both servers.

Uninstalling HA Module Only

When you uninstall the HA Module only, the systems are no longer part of a cluster installation. Use the following steps to uninstall HA Module and convert one of the systems to a single ESM installation. Options that you can choose from when reconfiguring the server:

- use the server's individual IP address and hostname.
- use the Service IP address and hostname.

If you use the server's individual host name or host IP address to identify the ESM Manager instance, you must also change the ESM Manager Host Name or IP address defined in every Connector and Console instance that connects to this ESM Manager. You must also update all bookmarks or URL references to the ArcSight Command Center. If you reuse the Service IP address and hostname, you will not have to make this change on clients that connect to that server.

NOTE: It's best practice to use a host name (rather than an IP Address) for greater flexibility in configuration.

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

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

2. On the same server, as user *root*, run:

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

After the HA uninstall is complete, all the files you need to run ESM are on both servers.

3. Choose which server will be the single ESM installation.
4. If you are not reusing the Service IP Address, use the procedure for changing the IP Address of an ESM Server described in the *ESM Installation and Configuration Guide*.
5. If you are reusing the Service IP address:
   a. Run the following command, as user *root,* 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 of the hostname is configured (for example, `eth0`).

   b. Update the ARP cache:

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

   c. Run the following command as user *root* on the server:

      ```
      /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 will not be brought up on the primary interface, and ESM will not be accessible.

   d. To make sure the ESM Service IP Address comes up at reboot on the selected server, change the appropriate scripts in `/etc/sysconfig/network-scripts/` on that server.

Chapter 9: An Example HA Implementation

This appendix 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 and Configuration Guide*.

- 2TiB of RAID 10 storage is provided via 15K RPM disks.
- The network interface runs at 1G.
- One 1G interface on each server will be interconnected by a cable.
- RedHat 7.1 is used with ESM 6.9.1c 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:

size = (2,306,867 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
```

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	/opt
Metadata volume	/dev/mapper/vg00-metadata
Service hostname	esm.internal.<mycompany>.com
Secondary hostname	ha2.internal.<mycompany>.com
Primary cable IP	192.168.10.2
Secondary cable IP	192.168.10.3
Preferred primary	
Connected hosts	10.10.11.5 10.10.12.5 10.10.13.5
Ping timeout	2
Ping attempts	2

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:

size = (5767168 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, 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 10: Maintaining and Monitoring 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 *root*.
Applies to	HA Module on either the primary or secondary machine.
Syntax	`/usr/lib/arcsight/highavail/bin/arcsight_cluster <action> [options]`

Actions	`clusterParameters [--console]`	Update the cluster parameters using the Cluster Parameters Wizard. Only run this on the primary. The `--console` option displays in console mode. GUI mode is the default.
	`diagnose`	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.
	`help (or -h)`	Provides command usage and HA version.
	`increaseDisk`	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.
	`offline [hostname]`	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.
	`online [hostname]`	This action makes the server [hostname] a candidate to be the primary. If there is already a primary, the other server is brought online as the secondary and specifying [hostname] is optional. If both servers are offline (but ready to be brought on line) you must specify the server to bring online. If online is not successful, it will suggest how the user may bring the server online.
	`prefer [hostname]`	System uses this server as the primary if both are eligible. If left blank the system only switches primary if the other server becomes ineligible (reduces fail-overs).
	`status`	Print the status of the cluster.
	`tuneDiskSync`	Update the configuration to improve disk sync speed. Do this whenever the speed of the interconnect cable is changed.

Examples	`./arcsight_cluster status`
	`./arcsight_cluster online myfirstesm.mydomain.com`
	`./arcsight_cluster prefer prod01`

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:

• preferred primary	• ping timeout
• connected hosts	• ping attempts
• connectivity down 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 `systemctl start heartbeat` or `service heartbeat 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 69 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

   size=(P/32768)+1

   where P is the size of the shared disk partition in mebibytes. 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 (under the Storage tab). See the *ArcSight Command Center Users Guide* for further details.

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.

prefer

The prefer action allows you to specify which server in the cluster should run ESM if both systems are configured properly and have that ability. When executed, the server defined by the hostname argument becomes the preferred primary. By default, neither system is preferred and the first available server will run ESM. You can use the prefer action to change the server running ESM. If a preferred primary has been defined, when it goes down the other system takes over. When the preferred primary system comes back up, another failover occurs so that this system again becomes the primary. This second failover would not happen if neither system was the preferred primary.

If there is currently a preferred server, and you run this action with no hostname, it changes the cluster to having no preference.

Run `arcsight_cluster prefer <hostname>` to set one system as the preferred primary.

Run `arcsight_cluster prefer` with no arguments to remove the preferred primary setting.

IMPORTANT: We recommend that you do not set a preferred primary server. In general, it is best to let the HA Module assign which system is the primary.

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 Preferred

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 Failover-Check-prod01.test.acme.com
Started Failover-Check-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 heartbeat 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 heartbeat function has failed:

```
Tue Sep 30 14:48:32 PDT 2014 FAIL Disk: Unconfigured
Cluster is stopped. Run "systemctl restart heartbeat" 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 Preferred
```

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. `Preferred` means this is the machine that the administrator wants to be the primary.

If the secondary was offline or it's heartbeat function stopped, and there was no preferred primary, 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 heartbeat 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.
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.
- **Failover-Check-<hostname>** is a program that checks if a failover is needed. An instance of it runs on each machine. For details see "An overview of the Failover-Check Operation" on page 73. (failover_check)
- **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 47 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 log output of three types, syslogs, HA logs, and upgrade logs

Upgrade Logs at */usr/lib/arcsight/highavail/logs/upgrade.log*. This contains information recorded about the upgrade process.

Syslogs, which generally get logged to `/var/log/messages`. These generally have to do with the status of the cluster, and any operations that are being performed. Linux automatically rotates these log files.

HA Log files in `/usr/lib/arcsight/highavail/logs`. These 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:

- arcsight_cluster.log Description of `arcsight_cluster` requests, and responses to the user.
- install-console.log Console output for installations run on this machine.
- install.log Installation file for installations run on this machine. Contains much more detail than install-console.log.
- secondaryHelper.log 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.

All syslog output from resources (plug-ins) goes 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 52, each resource description is followed by a string you can use to search `/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 56

"Changing the Primary Hostname or IP Address Only" on page 56

"Changing Both Server Hostnames or IP Addresses" on page 57

"Changing the Interconnect IP Address" on page 59

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 by running (as user *arcsight*):

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

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 (please note, 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 commands on the secondary as user *root*:

   ```
   systemctl stop heartbeat
   ```

   or

   ```
   service heartbeat 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 command on the secondary system as user *root*:
   ```
   systemctl disable heartbeat
   ```
   or
   ```
   chkconfig --del heartbeat
   ```

   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 heartbeat 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*:

   ```
   systemctl stop heartbeat
   ```

   or

   ```
   service heartbeat 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 command on the secondary system as user *root*:
   ```
   systemctl disable heartbeat
   ```
   or
   ```
   chkconfig --del heartbeat
   ```

   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 heartbeat restarts.

You may want to use `arcsight_cluster prefer <secondary hostname>` to fail back to the original server, but it is not necessary. In that case, run `arcsight_cluster prefer` with no arguments. This leaves ESM running on the new primary without an additional failover interruption.

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, you must uninstall and then re-install the HA Module.

1. Run the following command on the secondary (System B) as user *root*:

   ```
   systemctl stop heartbeat
   ```

or

```
service heartbeat 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 command on the secondary system as user *root*:
      ```
      systemctl disable heartbeat
      ```
      or
      ```
      chkconfig --del heartbeat
      ```
   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 heartbeat restarts and you are done with the secondary (System B).
5. Run the following command on the primary (System A) as user *root*:

   ```
   systemctl stop heartbeat
   ```

   or

   ```
   service heartbeat 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 command on the secondary system as user *root*:
      ```
      systemctl disable heartbeat
      ```
      or
      ```
      chkconfig --del heartbeat
      ```
   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 heartbeat 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. Change to the `/etc/sysconfig/network-scripts` directory.
2. 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`.
3. Run the ifdown and ifup commands (for example, `ifdown eth1; ifup eth1`).
4. 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 problem)s. Note that you need to bring down ESM during the installation on the new secondary. The procedure is given below:

1. Power down the server to be replaced. The other server will then become the primary.
2. Prepare the new server as described in "Configuring Systems before Installing the HA Module" on page 19. The new server may have different IP addresses and hostnames than the one it replaces and there are manual steps to perform on this machine as the secondary.
3. Stop ESM services on the primary by running the following command as user *root*:

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

4. Run the First Boot Wizard as user *arcsight* on the primary and specify the hostname or IP address for the new secondary system if it's different from the original.
5. Restart ESM services as user *root* on the primary:

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

At this point, ESM should come up again on the primary system. The new server will become the secondary system. The synchronization process between the primary system and this new secondary system may take some time. See the "Planning for the Initial Disk Synchronization" on page 15 section for more information.

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 11: 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

Each of the following messages would be prefixed with the following:

```
[Primary|Secondary]:  [Timestamp] ERROR – <message>
```

The following table lists the possible installation error messages, what they mean, and what to do if you get that message. Angle brackets (< >) enclose values such as names or IP addresses that are unique to your message.

Installation Message	Description
User and Access Issues	
Fatal error on <hostname>. See <log file>.	An unexpected error caused SSH to fail to <hostname> check the specified log file for suggestions.
Timeout on SSH to <hostname>. SSH access to <hostname> failed to connect quickly.	Fix the SSH communication problem.

Installation Message	Description
Incorrect root password for <hostname> - please enter correct one.	You entered an incorrect password. Enter the correct one.
Failed to set up SSH access. See <log file> for details.	SSH access didn't work. See the specified log for suggestions.
No arcsight user on secondary. Please create one identical to that on primary	Create a user *arcsight* on the secondary.
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 secondary. make them the same.
arcsight users on primary and secondary must have the same home directory.	Make them the same.
Crossover Cable Issues	
Speed of secondary end of crossover cable is <secondaryCableSpeed>M - must be at least 1000M.	Secondary interface for interconnect is slower than Gigabit ethernet. Use a faster interface.
Primary Cable IP <primaryCableCIDR> and Secondary Cable IP <secondary_cable_ip> must be in the same subnet.	Make the IP subnets consistent.
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.
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.
Speed of primary end of crossover cable is <primaryCableSpeed>M - must be at least 1000M.	Primary interface for interconnect is slower than Gigabit ethernet. Use a faster interface.
Shared Disk Issues	
Unmount of <shared_disk> failed. Fix the problem, and re-run this script.	Fix the problem and re-run the First Boot Wizard.
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 remove them from /etc/fstab.
<metadata_vol> should not be mounted.	The metadata volume is mounted - and it should not be. Unmount it. Most likely you will also get the *"<metadata_vol> appears to be in use."* error. Please follow the instructions for that error as well.

Installation Message	Description

<metadata_vol> appears to be in use. See the following output from
`blkid <metadata_vol>`
--- blkid output here ---
If this volume is not in use, run
`dd if=/dev/zero of=<metadata_vol>`
to clear this volume and then rerun the First Boot Wizard.
	It looks like someone is already using the metadata volume.
	Be certain this is not the case, then run the given `dd` command and re-run the First Boot Wizard.

Disk status must be Connected to reconfigure cluster.	The HA Module is already installed on both machines, so this call to the First Boot Wizard must be to reconfigure the installation. This can only be done if the disk status is Connected (normal). Run `arcsight_cluster diagnose` and then try re-running the First Boot Wizard.
Please mount <shared_disk partition>, and re-run installation.	Mount the shared disk.
Size of metadata volume <metadata_vol> is less than required minimum of <megabytes>M	The metadata volume is too small to support shared_disk. Increase the size of the metadata volume.
The size of <volume> on the secondary is <megabytes>M. It must be the same as the primary - <megabytes>M.	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 make them match.
<volume> not a valid disk volume.	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.
Found <megabytes>M disk space used on <shared_disk>. The installation will not proceed with these files in place. If these files are not important, run "rm -rf <shared_disk>/*" as root on <hostname> and re-run the First Boot Wizard.	The installation found more than 10MB of files on <shared_disk> on the secondary. The installation is terminated. Remove the files, and then re-run the First Boot Wizard.
<shared disk volume> mounted on <shared_ disk> on the primary and on <secondary_disk> on the secondary. It must be mounted on the same mount point on both machines.	Make sure the volume of the shared disk is mounted on the same mount point on both machines.

Installation Message	Description
<secondary host name> <time> ERROR - Found <size> disk space used on /opt. The installation will not proceed with these files in place. If these files are not important, run "rm -rf /opt/*" as root on <secondary host name> and re-run the installation.	The contents of the shared partition on the secondary machine must be empty before the installation wizard will run. Make sure it does not contain data of value and remove it.
Cannot do a Reconfiguration when disks are in <status> status. Please correct the disk status before doing reconfiguration.	The <status> 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.
Primary/Secondary Host Issues	
No interface found for <primary_ip> on Primary	The primary IP/hostname must be the first IP on an interface. Configure the primary hostname to correspond to an interface.
No interface found for <secondary_ip> on Secondary	The secondary IP/hostname must be the first IP on an interface. Configure the secondary hostname to correspond to an interface.
Primary IP <primary IP> and Secondary IP <secondary IP> must be in the same subnet.	Change host IP addresses so 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 incorrect hostname.
<hostname> is not a valid hostname or IP address.	Correct the incorrect hostname.
IP for <host> is <found IP> on this server and <other IP> on the other.	The hostname resolves to different IP addresses on the different servers. Make the server configurations consistent. A likely cause is inconsistencies in /etc/hosts.
OS version on primary and secondary are different.	Make them the same.
Could not send and return test string using ssh. Expected "test", saw "$returnedString"	There is a problem with the ssh login. Manually check that `root` user can ssh between systems in both directions (i.e. from System A to System B and from System B to System A).

Installation Message	Description
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 may cause problems with SSH communication. Please disable the SSH banner. See the section "Configuring Systems before Installing the HA Module" on page 19 for instructions about how to remove it.
Cluster did not come up after installation. See the status output above this message.	This happens rarely. Check the install.log file for details about the error condition. This message may appear because of a temporary condition, and within a few minutes the system will be working as expected. If the problem persists, contact Customer Support.
Cluster Upgrade Issues	
Cluster should not be running during upgrade. Run "systemctl stop heartbeat" as root to stop cluster.	The system should not be running during the upgrade process. Run "systemctl stop heartbeat" or "service heartbeat stop" as the `root` user 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.

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 date 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 "Setting Configurable HA ModuleProperties" on page 36.

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 occurrences can trigger a failover:

- You select the secondary as the preferred system using the `arcsight_cluster` command. This is the preferred way of forcing a failover.
- You put 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 may 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.
- Loss of an internet connection to the primary system. (it may take several minutes.)

The following occurrences do not trigger a failover:

- You can manually stop the ESM Manager or any of it's services without triggering a failover. For example, if you change a property in the server.properties file and have to restart the Manager, it does not trigger a failover.
- If the network switch fails causing a communications failure to both primary and secondary systems, there is no failover. Users would immediately detect that their ArcSight Console or ArcSight Command Center UIs have 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.
- If the primary system runs out of disk space, the secondary also runs out of space because of the mirroring. No failover is triggered.

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

The Failover-Check resource does not fail over if the Connected Hosts parameter is empty, or if none of the hosts respond to ping. For further information, see "An overview of the Failover-Check Operation" on page 73.

Failovers may 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 60, adding the inode64.

Your mount command might look something like this:

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

System Fails Over for no Reason

If the primary is preferred, and there is a failover to the secondary, and subsequently the primary comes back on line, a second failover to make the preferred server the primary occurs.

Make sure the Connectivity Down Timeout is more than 120 seconds. If Connectivity Down Timeout is less than 120 seconds, a single ping failure from the secondary to the primary causes a failover.

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 and Configuration 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`, if you haven't already. 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 47). 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 either `systemctl stop heartbeat` or `service heartbeat 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 `systemctl start heartbeat` or `service heartbeat 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 prepareHA Script

The prepareHA.sh script automates certain steps described in "Configuring Systems before Installing the HA Module" on page 19. It is provided as a convenience to help streamline some to the pre-configuration tasks.

The script must be run as the root user and requires a highavail.properties file that defines the cluster configuration. The script takes a single optional argument identifying the location of the highavail.properties file. If you do not specify the file location, the script looks for this file in the same directory as the script. See also "The highavail.properties File" on page 72 for more information. You can choose to either perform these steps manually or use this script.

- Turn off SSH login messages
- Create the metadata volume if it does not exist
- Create the arcsight user
- Create the /usr/lib/arcsight directory
- Create the /etc/security/limits.d/90-nproc.conf file
- Create the symbolic link to libpcre.so.0

The script also checks that:

- Both servers are running the same supported operation system version.
- Both servers use the same shared disk on the same volume with the same size.
- There are no mounts on the shared disk on either system.
- The metadata volume is large enough on both systems.
- Yum is configured on both systems
- The primary, secondary, and service hostnames all resolve to the same IP address on both the primary and secondary servers.
- It is possible to ping from the primary to the secondary and from the secondary to the primary through either the network or the crossover cable.
- The `arcsight` user UID and GID values match on both systems.

If you choose to use this script, copy the `Tools` directory (and all contents) to the `/tmp` directory which avoids complications that might arise if the directory or the contained files are copied elsewhere. Run the script from this new location. The `Tools` directory is in the location where you unpacked the ESM installation tar file.

```
cp -r Tools /tmp
```

Appendix B: 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 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 6.9.1 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 C: An overview of the Failover-Check Operation

This appendix describes how the Failover-Check resource determines that the cluster should failover to the secondary because of problems with access from the intranet to the primary. It is helpful background for understanding how to configure the Failover-Check resource, and for fixing problems when it doesn't fail over as expected.

The Failover-Check resource takes the following parameters:

- Connected Hosts – a list of hostnames or IPs to ping.
- Connectivity Down Timeout – The number of seconds to wait before considering that the primary internet connection is down and a failover should occur (Default 180).
- Ping Timeout – The number of seconds to wait before considering that a ping request has failed (Default 2).
- Ping Attempts – The number of times to try a ping before considering that it has failed (Default 2).

How Failover Check Works

A ping check uses the standard Linux `ping` command. This command sends one ping per second to the destination up to the number defined by the Ping Attempts parameter. A ping is considered to have failed if no response is within the number of seconds defined by the Ping Timeout parameter.

The Failover checking is done on the secondary system. Every two minutes, it goes through the following steps to update its "primary-down" information and, if necessary, initiate a fail over.

1. Ping the Service Hostname or Service IP address.
   a. If this succeeds, it removes the existing record that the primary ping failed and skips the remaining steps in this process. It repeats this step in two minutes.

   b. If it fails, it performs step 2.

2. Since the Ping failed, it checks to see if there is a record indicating that the previous ping failed also.
   a. If there is no record, it creates a new record indicating that this attempt failed and then skips the remaining steps. It repeats step 1 in two minutes.

   b. If there is a record of a previous failure, it performs step 3.

3. Because the ping attempt failed and there was a previous failure, it checks to see if the time between the first failure and the current time is less than that defined by the Connectivity Down

Timeout parameter.

a. If it is less, then it skips the next step. It repeats step 1 in two minutes.

b. If it is more than that timeout, it performs step 4.

4. It attempts to ping each of the hosts on the Connected Hosts list. If any of these attempts succeed, this indicates that the secondary system has network access, but the primary does not. A failover is initiated to the secondary.

Note that if there is a network failure that affects both the primary and the secondary, a failover will no occur.

Failover Parameter Guidelines

The Connected Hosts list should be representative hosts in your network that can respond to ping. If your network does not support ping, you can leave this value empty – but this will have the effect of disabling the Failover-Check feature and the system will not failover when the primary gets disconnected from the intranet.

The Connected Hosts list should be chosen as a test of whether the network is working properly. If the network is down, there is little point in doing a failover. For that reason, the First Boot Wizard and the Cluster Parameters Wizard disallow the use of the following hosts:

- Primary
- Secondary
- Service Hostname or Service IP
- Primary Cable IP
- Secondary Cable IP
- localhost

The Connectivity Down Timeout value must be longer than 120 seconds, which is the polling period used by the Failover-Check. If it were 120 seconds, a single, failing ping may cause the system to failover. The default, 180 seconds, is a good choice.

The results of the Failover-Check described in the previous section are ignored by the system if the check takes longer than 90 seconds. The First Boot Wizard and the Cluster Parameters Wizard limit the number of connected hosts, and the values of Ping Timeout, and ping attempts by the formula (below) so that the check never takes this long. The longest time a ping check on a single host can take is [Ping Attempts] + [Ping Timeout] seconds, since the attempts are sent out within [Ping Attempts] seconds, and then the last ping times out after [Ping Timeout]. At most, Failover-Check pings the primary and the hosts on the Connected Hosts list. So the following inequality must be met:

([Ping Attempts] + [Ping Timeout]) * (1 + # of Connected Hosts) < 90

The left side (of the <) represents the longest time the operation may take, and the right hand side is the longest the system will wait for the operation to complete.

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 system, 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 High Availability Module 6.9.1)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to arc-doc@hp.com.

We appreciate your feedback!