
Micro Focus Security

ArcSight ESM

Software Version: 7.0 Patch 1

ESM High Availability Module User's Guide

Document Release Date: August 16, 2018

Software Release Date: August 16, 2018



Legal Notices

Warranty

The only warranties for products and services of Micro Focus and its affiliates and licensors ("Micro Focus") are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Micro Focus shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Except as specifically indicated otherwise, a valid license from Micro Focus is 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

© Copyright 2001-2018 Micro Focus or one of its affiliates.

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

Support

Contact Information

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

Contents

| | |
|---|----|
| Chapter 1: Introduction | 7 |
| Chapter 2: General Environment and System Configuration | 9 |
| Place Systems in a High Availability Environment | 9 |
| Network Requirements | 10 |
| Using the Service IP to Identify the Cluster | 11 |
| Getting the License File | 12 |
| Spectre and Meltdown Patches Required for RHEL 6.9 and 7.4 and CentOS 6.9 and 7.4 | 12 |
| Chapter 3: Installing HA with a New ESM | 13 |
| Hardware Requirements for a New Installation | 13 |
| Software Requirements for a New Installation | 15 |
| System Configuration for a New Installation | 15 |
| Running the HA Module Installation Script | 18 |
| Running the HA Module First Boot Wizard | 19 |
| Install ESM | 21 |
| Verify New HA and ESM Installation | 22 |
| Chapter 4: Installing HA on an Appliance | 23 |
| Appliance Requirements | 23 |
| Appliance Configuration | 23 |
| Running the HA Module Installation Script | 25 |
| Running the HA Module First Boot Wizard | 26 |
| Verify New HA and ESM Installation | 29 |
| Chapter 5: Installing HA with an Existing ESM | 30 |
| Hardware Requirements when Installing HA on an Existing ESM | 30 |
| Planning for the Initial Disk Synchronization | 32 |
| Software Requirements when Adding HA to an Existing ESM | 32 |
| System Configuration for Adding HA to Existing ESM | 34 |

| | |
|--|----|
| Running the HA Module Installation Script | 36 |
| Running the HA Module First Boot Wizard | 37 |
| Verify HA Module on an existing ESM | 40 |
| Chapter 6: Upgrading ESM and the HA Module | 42 |
| Verifying the HA and ESM Upgrade | 43 |
| Chapter 7: Uninstalling Software Components | 45 |
| Uninstalling both ESM and HA Module | 45 |
| Uninstalling HA Module Only | 45 |
| Chapter 8: An Example HA Implementation | 47 |
| Server Configuration | 47 |
| Initial Setup and Installation | 48 |
| Hardware | 48 |
| DNS Setup | 48 |
| Operating System Installation | 48 |
| Disk Partition Setup | 49 |
| Interconnect Cable Setup | 50 |
| Set Up Connected Hosts | 50 |
| Install ArcSight Software | 50 |
| Increase Disk Space | 51 |
| Chapter 9: Maintain and Monitor the Cluster System | 53 |
| The arcsight_cluster Script | 53 |
| Command Syntax | 53 |
| clusterParameters | 54 |
| diagnose | 55 |
| increaseDisk | 55 |
| offline | 56 |
| online | 57 |
| status | 57 |
| Status Output Example | 57 |
| Status Output Explanation | 58 |
| tuneDiskSync | 60 |
| Log Output | 61 |
| Changing Hostname, IP Address, or Service IP | 61 |

| | |
|---|----|
| Changing the Cluster's Service IP Address | 62 |
| Changing the Secondary Hostname or IP Address only | 63 |
| Changing the Primary Hostname or IP Address Only | 64 |
| Changing Both Server Hostnames or IP Addresses | 64 |
| Changing the Interconnect IP Address | 66 |
| Replacing a Server | 66 |
| Changing Mount Options | 67 |
| Setting Configurable HA ModuleProperties | 67 |
| Chapter 10: Troubleshooting the Systems | 68 |
| Installation Issues and Solutions | 68 |
| General Problems | 71 |
| Changing ESM to IPv6 | 71 |
| Audit Events | 72 |
| highavailability:100 | 72 |
| highavailability:200 | 72 |
| highavailability:300 | 73 |
| highavailability:500 | 73 |
| Failover Triggers | 73 |
| Processes Killed During Failover | 74 |
| System does not Failover | 74 |
| System Fails Over for no Reason | 74 |
| Network Interface Commands Stall Disk Mirroring | 74 |
| No ESM Uninstall Links on the Primary | 75 |
| Stopping the Network on the Secondary Kills ESM | 75 |
| Disks on Cluster System Fail to Connect | 75 |
| Appendix A: The highavail.properties File | 77 |
| Appendix B: Upgrade HA Appliance Operating System | 78 |
| Verify Operating System Upgrade File | 78 |
| Upgrade HA Operating System | 78 |
| Appendix C: An overview of the Failover-Check Operation | 80 |
| How Failover Check Works | 80 |

| | |
|-------------------------------------|----|
| Failover Parameter Guidelines | 81 |
| Send Documentation Feedback | 82 |


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: General Environment and System Configuration

This section describes the hardware, software, networking, and other requirements that are needed before the installation begins. This information applies to all installation scenarios. 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. See the *ESM Installation Guide* or the *ESM Upgrade Guide* for the specific requirements to install or upgrade ESM. You will use these documents together to plan your ESM and HA Module cluster installation.

Note: If you are planning to install ESM in a distributed correlation implementation, note that HA is supported only on the persistor node in the distributed correlation cluster. HA is not supported on any non-persistor node in a distributed correlation cluster.

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.

Refer to ["An Example HA Implementation" on page 47](#) for a specific example of 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 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.

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.
- Connected Hosts may be IPv4, IPv6, or a combination of the two.
- The two HA systems must be part of the same IPv4 or IPv6 subnet. If you change the ESM subnet after the HA Module is installed, for example by changing from IPv4 to IPv6, you have to uninstall and reinstall the HA Module. This will require about 30 minutes of downtime, and also require a resync of the data.
- The primary and secondary machines must be close enough together that the cable 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. These may be IPv4 or IPv6 addresses. Note: You can use private IP addresses if you are certain that ESM will not route communication to these addresses.
 - One IP address is the "Service IP" address that is assigned to the ESM cluster. You will specify the host IP addresses and the "Service IP" address when 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. The Service IP must be on the same subnet as the host IP addresses.
- 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. For example, using DNS enables you to manage the Service IP address in relation to the numerous connectors, Consoles, and Command Centers associated with a specific ESM Manager. Also, using DNS enables you to keep the IPs or host names consistent for the primary and secondary machines in your cluster. However, you would not want to use DNS to track the IP addresses for the primary and secondary cables; there is no benefit from using DNS in this case.
- The HA Module uses ports 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 firewall, 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 | <ul style="list-style-type: none">◦ the primary IP address◦ the secondary IP address | N/A | <ul style="list-style-type: none">◦ the primary IP address◦ the secondary IP address◦ the Service IP address◦ to the Connected Host | N/A |
| TCP | <ul style="list-style-type: none">◦ the primary crossover cable◦ the secondary crossover cable | Any | <ul style="list-style-type: none">◦ the primary system cable◦ the secondary system cable | 7789 |
| UDP | <ul style="list-style-type: none">◦ the primary IP address◦ the primary crossover cable◦ the secondary IP address◦ the secondary crossover cable | Any | <ul style="list-style-type: none">◦ 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 or host name during ESM installation, you will specify the Service IP address and not an individual host IP address. The Manager host name should resolve to the Service IP address.

When the ArcSight Console connects to the Manager, it will use the Service IP address. Also, 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.

Best Practice: It is recommended that you configure a hostname during an ESM installation. Host name changes are easier to manage using DNS, and hostnames are required for IPv6 systems.

Note: Using HA ESM and ArcSight Event Broker:

In an HA ESM environment, if there is a mismatch between the actual ESM hostname/IP address and the one listed under the ESM consumer group on the Event Broker Manager, this could be due

to the underlying third party library in the Event Broker preferring the ESM primary or secondary hostname/IP address instead of service IP.

This has no impact on ESM and Event Broker functionality. Refer to the *ArcSight Data Platform Event Broker Administrator's Guide* for additional information.

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 Micro Focus."

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 Guide* for detailed information about installing ESM.

If you are upgrading from ESM 6.11.0 and HA Module 6.11.0, you do not need a new license file.

Spectre and Meltdown Patches Required for RHEL 6.9 and 7.4 and CentOS 6.9 and 7.4

For HA, you must have the Spectre and Meltdown patches installed on RHEL 6.9 or 7.4, or on CentOS 6.9 or 7.4.

Chapter 3: Installing HA with a New ESM

This section describes how to configure your machines and then run the ESM High Availability Module installation wizard and First Boot Wizard. This section is for the case where you are installing both ESM and the High Availability module for the first time.

Note: Be sure to install ESM after HA has completed disk synchronization. Attempting to install ESM while HA synchronization is in process can cause the ESM installation to fail.

Hardware Requirements for a New Installation

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.

- 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.

See the *ESM Installation 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 | 6 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. |

- 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 is calculated as: size (in mebibytes)= (P/32) + 1 P = shared_disk_partition size in gibibytes |
| / (root) partition | 20 GiB (generous) | | An operating system recommended partition. |
| swap | 8 GiB (minimum) | | An operating system recommended partition. |
| temp | 10 GiB (or more) | /tmp | An operating system recommended partition. |

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.
- 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 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 32](#) 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.

Software Requirements for a New Installation

- The HA Module version, ESM version, and operating system version must be compatible. See the *ESM Support Matrix* for ESM, HA Module, and Operating System version compatibility.
- The cluster systems must run either RHEL or CentOS. Both systems must have the same operating system and version installed.

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.

- 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.
- 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 operating system (OS) vendor, a repository created from the OS ISO or CD, or a directory location on the local system. See the vendor-specific documentation for information about configuring Yum and connecting to Yum repositories.
- 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. Note that 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 49](#) for an example of how to do this.
- Download the compatible ESM and HA Module files from the download site to the primary system. The files are:
 - ArcSightESMSuite-7.0.0.xxxx.1.tar
 - ArcSight-Highavail-7.0.0.xxxx.1.tar
- Unpack both tar files. Do NOT unpack them into what will later be the shared directory (generally */opt*), because files there are deleted during installation. You install ESM and the HA Module on the primary system only. After installation, the HA Module synchronizes the secondary system with the primary.

System Configuration for a New Installation

The primary and secondary appliances must be set up so that they are nearly identical. The following steps must be performed as directed on the primary, the secondary, or both appliances to ensure that

they are configured properly to run the HA Module. The HA Module installation scripts check the configuration and return an error message if dependencies are not met.

1. Make sure that both systems have the correct version of the operating system timezone package installed. This is a requirement for ESM. For instructions, refer to the *ESM Installation Guide*, specifically the topic "Install Time Zone Package" in the "Installing ESM" chapter.
2. 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.
3. Connect the two servers with crossover cables. Configure the interfaces with the appropriate IPv4 or IPv6 addresses. They must both use the same IP version. Ping from one system to the other over the configured interfaces to be certain they are configured properly.
4. 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. 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 to `/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.

5. If the mirrored disks are SSD drives, such as Fusion, make sure you have TRIM support configured on both the primary and secondary systems.
6. 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.
7. 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:

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

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

See ["Disk Partition Setup" on page 49](#) 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 metadata partition will be a physical volume (for example, `/dev/sda8`), then create it now. If the metadata partition will be a logical volume (for example, `/dev/mapper/vg00-meta`), then you only need to ensure that at least "size" free disk space is available in a volume group. The `prepareHA.sh` script will create the metadata volume.

8. 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.

9. As user *root* untar the ESM install tarball on the primary if you have not already done so, and then as *root* run:

```
cp -r Tools /tmp  
cd /tmp/Tools/highavail  
cp template.properties highavail.properties  
chmod 644 highavail.properties
```
10. Edit the file `/tmp/Tools/highavail/highavail.properties`, and fill in the empty fields.
 - `service_hostname=` (The hostname of ESM in the HA cluster.)
 - `shared_disk=` (The mount point of the disk to be mirrored across both systems.)
 - `metadata_volume=` (The volume name for the metadata volume, for example, `/dev/mapper/vg00-meta`.)
 - `primary_cable_ip=` (The IP Address of interface to the cable connected to the secondary.)
 - `primary_hostname=` (The hostname of the primary.)
 - `secondary_cable_ip=` (The IP Address of interface to the cable connected to the primary.)
 - `secondary_hostname=` (The hostname of the secondary.)

11. On the primary, as user *root*, run:
`/tmp/Tools/highavail/prepareHA.sh`
 - The script will ask you to confirm the names of the primary and the secondary. Answer yes to continue.
 - If the metadata partition does not exist, and it will be a logical volume, the script will offer to create it. Answer yes.
 - The script will ask for a password for the arcsight user.

If there are any errors, correct them, and rerun `prepareHA.sh`. Continue until `prepareHA.sh` runs without errors.

12. Run the following command as user *root*:

```
scp -r /tmp/Tools <secondary hostname>:/tmp
```

It is important that the `highavail.properties` file gets copied over with the other files, and that the file permissions are preserved. The above command does both of these things.
13. Reboot the primary.
14. On the secondary, as user *root*, run:
`/tmp/Tools/highavail/prepareHA.sh`

If there are any errors, correct them, and rerun `prepareHA.sh`. Continue until `prepareHA.sh` runs without errors.
15. Reboot the secondary.

At this point it is assumed that you have already completed all the required tasks for the primary and secondary machines as described above.

You can run the HA 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. 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, change to the directory where you extracted the file and execute either:

```
./ArcSight-Highavail-7.0.0.xxxx.1.bin -i console for console mode
```

or

```
./ArcSight-Highavail-7.0.0.xxxx.1.bin for GUI mode
```

2. At the **Introduction** prompt, either click **Next** (GUI) or press **Enter** (console mode). The rest of these instructions document console mode.
3. 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. In GUI mode the radio button is grayed out until you scroll to the bottom of the license agreement.
In Console mode, press **Enter** at each prompt to scroll to each page of the license agreement.
4. The installer displays a **Pre-installation Summary**. Press **Enter** to continue. The installation shows its progress.
5. If you ran the installer in console mode, you will be prompted to start the First Boot Wizard when the installer is complete.
If you ran the installer in GUI mode, the First Boot Wizard starts automatically.

Running the HA Module First Boot Wizard

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

To run the First Boot Wizard, change user to *arcsight*, and then type (all on one line):

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

2. At the **Welcome to the First Boot Wizard** prompt, click Next (GUI) or enter **yes** (console).
3. At the **License File** prompt, supply the path to your ArcSight license file (either the zip file or the .lic file that is in 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.
4. The **Properties File** prompt offers the opportunity to load the `highavail.properties` file that defines the cluster configuration. It is the file you created at `/tmp/Tools/highavail/highavail.properties`.
On an appliance, leave this field blank.
5. At the **Hostname Inputs** prompt, you enter the hostnames and some other configuration parameters, as described in the following table.

| Field | Description |
|-----------------|---|
| Shared Disk | <p>Enter the mount point of the disk shared between the primary and secondary. (/opt)The options provided include all relevant mount points.</p> <p>Except on an appliance, the installation does not support bind mounts and it flags them as an error. Use symbolic links instead.</p> <p>The installation <i>completely erases</i> the contents of the shared disk on the secondary, so make sure it contains no data of any value.</p> <p>Make sure no process on the primary or secondary is using this file system or the installation will exit with errors.</p> <p>You cannot change this value on subsequent runs of the First Boot Wizard.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>shared_disk</code>.</p> |
| Metadata Volume | <p>Enter the volume containing disk-synchronization metadata. This volume is expected to start with <code>/dev</code>. On an appliance it is <code>/dev/sda6</code>.</p> <p>The contents of the metadata volume on both the primary and the secondary will be removed.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>metadata_volume</code>.</p> |

| Field | Description |
|--|---|
| Metadata Volume for ESM Express or EMS Appliance | The metadata partition for ESM Express or ESM Appliance is <code>/dev/sda6</code> . |
| 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 <code>hosts</code> file or DNS to resolve the hostnames. This value is identified in the <code>highavail.properties</code> file as <code>service_hostname</code> . |
| 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 <code>highavail.properties</code> file as <code>secondary_hostname</code> . |
| 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 <code>highavail.properties</code> file as <code>primary_cable_ip</code> . |
| 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 <code>highavail.properties</code> file as <code>secondary_cable_ip</code> . |

Click **Next** or type **Yes** and press **Enter** to continue.

- At the **Parameter Configuration** prompt, enter the following information:

| Field | Description |
|---------------------------|---|
| 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 seconds. |
| 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. |

A summary screen of your hostname inputs and other configuration parameters is displayed. IP Addresses are resolved to hostnames, and host names are resolved to IP addresses. The wizard decides whether to use IPv4 or IPv6 for the Service IP, and explains its reasons. If you do not like its choice you may be able to force it to choose IPv6 by entering an IPv6 address instead of a hostname, or to choose IPv4 by entering an IPv4 address instead of a hostname.

For more information in how these settings affect Failover, see ["An overview of the Failover-Check Operation" on page 80](#).

Click **Next** or type **Yes** and press **Enter** to continue.

7. At the "root password" prompt, enter the password for user *root*, and then continue.

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 in 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* user's password.
9. If your shared disk is empty, the wizard assumes that this is a fresh installation. It prompts you with additional information about the duration of the remaining processes. Click **Next** (GUI) or press **Enter** (console) to continue.

The installation displays the status of each operation as it runs. The status is displayed at the console in console mode, or in a special window in GUI mode. 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 68](#).

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
```

Install ESM

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

After you have installed the HA module on the primary system, complete these steps:

1. 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`.

2. On the primary system, install ESM. See the *ESM Installation Guide* for details. Note that if HA is already installed, then there is no need to run `prepare_system.sh` while 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.

3. Complete the post-installation steps described in ["Verify New HA and ESM Installation" on page 29](#).

Verify New HA and ESM Installation

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 Guide*, specifically the chapter titled "Post-Installation Considerations". After the ESM post-installation configuration is complete:

1. Make sure that both systems have the correct version of the operating system timezone package installed. This is a requirement for ESM. For instructions, refer to the *ESM Installation Guide*, specifically the topic "Install Time Zone Package" in the "Installing ESM" chapter. The timezone package is not installed within the shared directory, so you have to install it separately on the secondary.
2. 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.

3. 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 53](#) for details about the command arguments available.

Chapter 4: Installing HA on an Appliance

This section describes how to configure your appliances and then run the ESM High Availability Module installation wizard and First Boot Wizard. This section is for the case where you are installing the High Availability module on an ESM Express or ESM on an appliance.

Appliance Requirements

- Download the compatible HA Module and HA support packages files from the download site to the primary system. The files are:
 - `ArcSight-Highavail-7.0.0.xxxx.1.tar`
 - `esm_ha_support_rpms_rhel74.tar.gz` (required only if migrating an appliance from a single installation to an HA Module cluster)
- Unpack the tar file. Do NOT unpack it into what will be the shared directory (generally `/opt/arcsight`), because files there are deleted during installation. You install ESM and the HA Module on the primary system only. After installation, the HA Module synchronizes the secondary system with the primary.

Appliance Configuration

Two appliances are required. The primary and secondary appliances must be set up so that they are nearly identical. The following steps must be performed as directed on the primary, the secondary, or both appliances to ensure that they are configured properly to run the HA Module. The HA Module installation scripts check the configuration and return an error message if dependencies are not met.

1. Disable the first boot scripts, as described below. New appliances include setup scripts that run at first boot to install the operating system and ESM. Do not install ESM on the secondary. 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. Allow ESM install on the primary and let the installation complete.
 - b. When you boot the secondary machine (the appliance), let the operating system installation script run normally.
 - c. At the prompt: "Install Anywhere will guide you through the installation of ArcSight ESM 7.0 Suite", enter the text: `quit`
 - d. To prevent the appliance from installing ESM after it is rebooted, edit the `.bash_profile` in the root user's home directory on both the primary and the secondary. Remove the lines from `'# run OS configuration and ESM installer if not yet done'` to the end of the file.

- e. Reboot the appliance on the primary and the secondary to ensure that the ESM installation script does not run.
2. Make sure that both systems have the correct version of the operating system timezone package installed. This is a requirement for ESM. For instructions, refer to the *ESM Installation Guide*, specifically the topic "Install Time Zone Package" in the "Installing ESM" chapter.
3. 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.
4. Connect the two servers with crossover cables. Configure the interfaces with the appropriate IPv4 or IPv6 addresses. They must both use the same IP version. Ping from one system to the other over the configured interfaces to be certain they are configured properly.
5. On the primary server, as user *root*, stop ESM by running:

```
/opt/arcsight/manager/bin/remove_services.sh
```
6. Run the following commands as *root* on both the primary and the secondary:

```
systemctl stop hp-asrd hp-health hp-snmp-agents  
mv /opt/hp /usr/local  
ln -s /usr/local/hp /opt  
umount /opt  
ln -s /usr/local/hp /opt  
mount /opt  
systemctl start hp-asrd hp-health hp-snmp-agents
```
7. 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.
8. 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_rpms_rhel74.tar.gz` file to the `/tmp` partition. You downloaded this file from the SSO download site in an earlier step.
 - c. Run following commands to install the supporting packages:

```
cd /tmp  
tar xzf esm_ha_support_rpms_rhel74.tar.gz  
cd install  
./install_ha_support_pkgs.sh
```
9. (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 a 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.

10. On both the primary and secondary systems, as the user `root`, run the following commands:

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

At this point it is assumed that you have already completed all the required tasks for the primary and secondary machines as described above.

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 and then 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. 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, change to the directory where you extracted the file and execute either:

```
./ArcSight-Highavail-7.0.0.xxxx.1.bin -i console for console mode
```

or

`./ArcSight-Highavail-7.0.0.xxxx.1.bin` for GUI mode

2. At the **Introduction** prompt, either click **Next** (GUI) or press **Enter** (console mode). The rest of these instructions document console mode.
3. 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. In GUI mode the radio button is grayed out until you scroll to the bottom of the license agreement.
In Console mode, press **Enter** at each prompt to scroll to each page of the license agreement.
4. The installer displays a **Pre-installation Summary**. Press **Enter** to continue. The installation shows its progress.
5. If you ran the installer in console mode, you will be prompted to start the First Boot Wizard when the installer is complete.
If you ran the installer in GUI mode, the First Boot Wizard starts automatically.

Running the HA Module First Boot Wizard

1. If you ran the installation wizard in GUI mode, the First Boot Wizard starts automatically and you can skip to the next step.
To run the First Boot Wizard, change user to *arcsight*, and then type (all on one line):
`/usr/lib/arcsight/highavail/bin/arcsight firstBootWizard --console`
2. At the **Welcome to the First Boot Wizard** prompt, click Next (GUI) or enter **yes** (console).
3. At the **License File** prompt, supply the path to your ArcSight license file (either the zip file or the .lic file that is in 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.
4. The **Properties File** prompt offers the opportunity to load the `highavail.properties` file that defines the cluster configuration. It is the file you created at `/tmp/Tools/highavail/highavail.properties`.
On an appliance, leave this field blank.
5. At the **Hostname Inputs** prompt, you enter the hostnames and some other configuration parameters, as described in the following table.

| Field | Description |
|--|---|
| Shared Disk | <p>Enter the mount point of the disk shared between the primary and secondary. (/opt)The options provided include all relevant mount points.</p> <p>Except on an appliance, the installation does not support bind mounts and it flags them as an error. Use symbolic links instead.</p> <p>The installation <i>completely erases</i> the contents of the shared disk on the secondary, so make sure it contains no data of any value.</p> <p>Make sure no process on the primary or secondary is using this file system or the installation will exit with errors.</p> <p>You cannot change this value on subsequent runs of the First Boot Wizard.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>shared_disk</code>.</p> |
| Metadata Volume | <p>Enter the volume containing disk-synchronization metadata. This volume is expected to start with <code>/dev</code>. On an appliance it is <code>/dev/sda6</code>.</p> <p>The contents of the metadata volume on both the primary and the secondary will be removed.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>metadata_volume</code>.</p> |
| Metadata Volume for ESM Express or EMS Appliance | <p>The metadata partition for ESM Express or ESM Appliance is <code>/dev/sda6</code>.</p> |
| Service Hostname | <p>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 <code>hosts</code> file or DNS to resolve the hostnames.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>service_hostname</code>.</p> |
| Secondary Hostname | <p>Enter the hostname of the secondary machine. This is the host name assigned specifically to the machine.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>secondary_hostname</code>.</p> |
| Primary Cable IP | <p>Select the IP address of the interface connected to the interconnect cable on the primary system.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>primary_cable_ip</code>.</p> |
| Secondary Cable IP | <p>Enter the IP address of the interface connected to the interconnect cable on the secondary system.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>secondary_cable_ip</code>.</p> |

Click **Next** or type **Yes** and press **Enter** to continue .

- At the **Parameter Configuration** prompt, enter the following information:

| Field | Description |
|---------------------------|---|
| 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 seconds. |
| 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. |

A summary screen of your hostname inputs and other configuration parameters is displayed. IP Addresses are resolved to hostnames, and host names are resolved to IP addresses. The wizard decides whether to use IPv4 or IPv6 for the Service IP, and explains its reasons. If you do not like its choice you may be able to force it to choose IPv6 by entering an IPv6 address instead of a hostname, or to choose IPv4 by entering an IPv4 address instead of a hostname.

For more information in how these settings affect Failover, see ["An overview of the Failover-Check Operation" on page 80](#).

Click **Next** or type **Yes** and press **Enter** to continue.

- At the "root password" prompt, enter the password for user *root*, and then continue.

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.

- If you are running in console mode,
 - you are prompted about whether to hide the input for private parameters from the screen. Press **Enter** to hide these parameters.
 - you are prompted to verify the *root* user's password.

- If your shared disk is empty, the wizard assumes that this is a fresh installation. It prompts you with additional information about the duration of the remaining processes. Click **Next** (GUI) or press **Enter** (console) to continue.

The installation displays the status of each operation as it runs. The status is displayed at the console in console mode, or in a special window in GUI mode. This may take an hour or so depending on whether you are upgrading an existing ESM.

- 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 68](#).

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
```

Verify New HA and ESM Installation

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 Guide*, specifically the chapter titled "Post-Installation Considerations". After the ESM post-installation configuration is complete:

1. Make sure that both systems have the correct version of the operating system timezone package installed. This is a requirement for ESM. For instructions, refer to the *ESM Installation Guide*, specifically the topic "Install Time Zone Package" in the "Installing ESM" chapter. The timezone package is not installed within the shared directory, so you have to install it separately on the secondary.
2. 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.

3. 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 53](#) for details about the command arguments available.

Chapter 5: Installing HA with an Existing ESM

This section describes how to configure your systems and then run the ESM High Availability Module installation wizard and First Boot Wizard. This section is for the case where you are installing the High Availability module where you already have a previous version of ESM installed.

Hardware Requirements when Installing HA on an Existing ESM

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.

- 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.

See the *ESM Installation 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 | 6 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 1 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. |

- 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 is calculated as: size (in mebibytes)= (P/32) + 1 P = shared_disk_partition size in gibibytes |
| / (root) partition | 20 GiB (generous) | | An operating system recommended partition. |
| swap | 8 GiB (minimum) | | An operating system recommended partition. |
| temp | 10 GiB (or more) | /tmp | An operating system recommended partition. |

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.
- 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 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 the next page](#) 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.

Planning for the Initial Disk Synchronization

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 1 Gb cable. In such cases, providing a 10 Gb 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.

Software Requirements when Adding HA to an Existing ESM

- The HA Module version, ESM version, and operating system version must be compatible. See the on *ESM Support Matrix* for ESM, HA Module, and Operating System version compatibility.
- 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. 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.
- The cluster systems must run either RHEL or CentOS. Both systems must have the same operating system and version installed.

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.

- If this is a new HA installation on an existing ESM earlier than ESM 7.0, upgrade to ESM 7.0 *before* installing HA.
- If you plan to convert the system from IPv4 to IPv6, do it after you upgrade to ESM 7.0 and before installing HA.
- The mirrored disk mount point (for example, /opt) must be the same on the secondary as it is on the primary. The mounted volume name (for example, /dev/sda5 or /dev/mapper/vg00-opt) must also be the same as the primary. That means that if the primary uses a physical volume for the mirrored disk, then the secondary must do so as well.
- 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 operating system vendor, a repository created from the OS ISO or CD, or a directory location on the local system. See the vendor-specific documentation for information about configuring Yum and connecting to Yum repositories.
- Create a secondary that has the same volumes as the primary. Use physical volume management on the secondary if that is what is done on the primary.
- Download the compatible ESM and HA Module files from the Software Entitlements Portal download site to the primary system. The files are:
 - ArcSightESMSuite-7.0.0.xxxx.1.tar (If you installed ESM 7.0 earlier and then removed this file, you will need it again for this procedure.)
 - ArcSight-Highavail-7.0.0.xxxx.1.tar
- Unpack the tar files. Do NOT unpack them into what will be the shared directory (generally /opt/arcsight), because files there are deleted during installation. You install ESM and the HA Module on the primary system only. After installation, the HA Module synchronizes the secondary system with the primary.
- 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.
- On the primary system, as the user *root* stop ESM by running:

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

Note: At this point, you must run `sshSetup` again if ESM in distributed mode. When you run `remove_services`, it removes the configuration of `sshSetup` after you install HA and run `setupServices`. See "Set Up Key-Based Passwordless SSH" in the *ESM Administrator's Guide* for details.

- (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 reuse 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.

System Configuration for Adding HA to Existing ESM

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

1. Make sure that both systems have the correct version of the operating system timezone package installed. This is a requirement for ESM. For instructions, refer to the *ESM Installation Guide*, specifically the topic "Install Time Zone Package" in the "Installing ESM" chapter.
2. 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.
3. Connect the two servers with crossover cables. Configure the interfaces with the appropriate IPv4 or IPv6 addresses. They must both use the same IP version. Ping from one system to the other over the configured interfaces to be certain they are configured properly.
4. On both the primary and secondary systems, select the partitions to be mirrored between the two servers. On the primary, use the command:

```
df /opt/arcsight
```

The mount point for /opt/arcsight will be shown under the "Mounted on" column.

Create and mount a volume on the secondary at this mount point. Give it the following characteristics with respect to the primary:

- same volume name
- same size
- same filesystem type

If /opt/arcsight is a symbolic link on the primary, the installer should create the same symbolic link on the secondary.

5. If the mirrored disks are SSD drives, such as Fusion, make sure you have TRIM support configured on both the primary and secondary systems.
6. 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.

7. 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:

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

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

See ["Disk Partition Setup" on page 49](#) 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 metadata partition will be a physical volume (for example, `/dev/sda8`), then create it now. If the metadata partition will be a logical volume (for example, `/dev/mapper/vg00-meta`), then you only need to ensure that at least "size" free disk space is available in a volume group. The `prepareHA.sh` script will create the metadata volume.

8. 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.
9. As user `root` untar the ESM install tarball on the primary if you have not already done so, and then as `root` run:

```
cp -r Tools /tmp
cd /tmp/Tools/highavail
cp template.properties highavail.properties
chmod 644 highavail.properties
```
10. Edit the file `/tmp/Tools/highavail/highavail.properties`, and fill in the empty fields.
 - `service_hostname`= [The hostname of ESM in the HA cluster]
 - `shared_disk`= [The mount point of the disk to be mirrored across both systems.]
 - `metadata_volume`= [The volume name for the metadata volume, for example, `/dev/mapper/vg00-meta`]
 - `primary_cable_ip`= [The IP Address of interface to the cable connected to the secondary]
 - `primary_hostname`= [The hostname of the primary]
 - `secondary_cable_ip`= [The IP Address of interface to the cable connected to the primary]
 - `secondary_hostname`= [The hostname of the secondary]
11. On the primary, as user `root`, run:

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

- The script will ask you to confirm the names of the primary and the secondary. Answer yes to continue.
- If the metadata partition does not exist, and it will be a logical volume, the script will offer to create it. Answer yes.
- The script will ask for a password for the arcsight user.

If there are any errors, correct them, and rerun `prepareHA.sh`. Continue until `prepareHA.sh` runs without errors.

12. Run the following command as user `root`:

```
scp -r /tmp/Tools <secondary hostname>:/tmp
```

It is important that the `highavail.properties` file gets copied over with the other files, and that the file permissions are preserved. The above command does both of these things.

13. Reboot the primary.

14. On the secondary, as user `root`, run:

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

If there are any errors, correct them, and rerun `prepareHA.sh`. Continue until `prepareHA.sh` runs without errors.

15. Reboot the secondary.

At this point it is assumed that you have already completed all the required tasks for the primary and secondary machines as described above.

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. 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, change to the directory where you extracted the file and execute either:

```
./ArcSight-Highavail-7.0.0.xxxx.1.bin -i
```

 for console mode
or

```
./ArcSight-Highavail-7.0.0.xxxx.1.bin
```

 for GUI mode
2. At the **Introduction** prompt, either click **Next** (GUI) or press **Enter** (console mode). The rest of these instructions document console mode.
3. 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. In GUI mode the radio button is grayed out until you scroll to the bottom of the license agreement.
In Console mode, press **Enter** at each prompt to scroll to each page of the license agreement.
4. The installer displays a **Pre-installation Summary**. Press **Enter** to continue. The installation shows its progress.
5. If you ran the installer in console mode, you will be prompted to start the First Boot Wizard when the installer is complete.
If you ran the installer in GUI mode, the First Boot Wizard starts automatically.

Running the HA Module First Boot Wizard

1. If you ran the installation wizard in GUI mode, the First Boot Wizard starts automatically and you can skip to the next step.
To run the First Boot Wizard, change user to *arcsight*, and then type (all on one line):

```
/usr/lib/arcsight/highavail/bin/arcsight firstBootWizard --console
```
2. At the **Welcome to the First Boot Wizard** prompt, click Next (GUI) or enter **yes** (console).
3. At the **License File** prompt, supply the path to your ArcSight license file (either the zip file or the .lic file that is in 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.
4. The **Properties File** prompt offers the opportunity to load the `highavail.properties` file that defines the cluster configuration. It is the file you created at `/tmp/Tools/highavail/highavail.properties`.
On an appliance, leave this field blank.
5. At the **Hostname Inputs** prompt, you enter the hostnames and some other configuration parameters, as described in the following table.

| Field | Description |
|--|---|
| Shared Disk | <p>Enter the mount point of the disk shared between the primary and secondary. (/opt)The options provided include all relevant mount points.</p> <p>Except on an appliance, the installation does not support bind mounts and it flags them as an error. Use symbolic links instead.</p> <p>The installation <i>completely erases</i> the contents of the shared disk on the secondary, so make sure it contains no data of any value.</p> <p>Make sure no process on the primary or secondary is using this file system or the installation will exit with errors.</p> <p>You cannot change this value on subsequent runs of the First Boot Wizard.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>shared_disk</code>.</p> |
| Metadata Volume | <p>Enter the volume containing disk-synchronization metadata. This volume is expected to start with <code>/dev</code>. On an appliance it is <code>/dev/sda6</code>.</p> <p>The contents of the metadata volume on both the primary and the secondary will be removed.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>metadata_volume</code>.</p> |
| Metadata Volume for ESM Express or EMS Appliance | <p>The metadata partition for ESM Express or ESM Appliance is <code>/dev/sda6</code>.</p> |
| Service Hostname | <p>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 <code>hosts</code> file or DNS to resolve the hostnames.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>service_hostname</code>.</p> |
| Secondary Hostname | <p>Enter the hostname of the secondary machine. This is the host name assigned specifically to the machine.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>secondary_hostname</code>.</p> |
| Primary Cable IP | <p>Select the IP address of the interface connected to the interconnect cable on the primary system.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>primary_cable_ip</code>.</p> |
| Secondary Cable IP | <p>Enter the IP address of the interface connected to the interconnect cable on the secondary system.</p> <p>This value is identified in the <code>highavail.properties</code> file as <code>secondary_cable_ip</code>.</p> |

Click **Next** or type **Yes** and press **Enter** to continue .

- At the **Parameter Configuration** prompt, enter the following information:

| Field | Description |
|---------------------------|---|
| 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 seconds. |
| 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. |

A summary screen of your hostname inputs and other configuration parameters is displayed. IP Addresses are resolved to hostnames, and host names are resolved to IP addresses. The wizard decides whether to use IPv4 or IPv6 for the Service IP, and explains its reasons. If you do not like its choice you may be able to force it to choose IPv6 by entering an IPv6 address instead of a hostname, or to choose IPv4 by entering an IPv4 address instead of a hostname.

For more information in how these settings affect Failover, see ["An overview of the Failover-Check Operation" on page 80](#).

Click **Next** or type **Yes** and press **Enter** to continue.

- At the "root password" prompt, enter the password for user *root*, and then continue.

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.

- If you are running in console mode,
 - you are prompted about whether to hide the input for private parameters from the screen. Press **Enter** to hide these parameters.
 - you are prompted to verify the *root* user's password.

- If your shared disk is empty, the wizard assumes that this is a fresh installation. It prompts you with additional information about the duration of the remaining processes. Click **Next** (GUI) or press **Enter** (console) to continue.

The installation displays the status of each operation as it runs. The status is displayed at the console in console mode, or in a special window in GUI mode. This may take an hour or so depending on whether you are upgrading an existing ESM.

- 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 68](#).

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
```

Verify HA Module on an existing ESM

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>`

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

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

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

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

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.
12. Make sure that both systems have the current version of the operating system timezone package installed. This is a requirement for ESM. For instructions, refer to the *ESM Installation Guide*, specifically the topic "Install Time Zone Package" in the "Installing ESM" chapter. The timezone package is not installed within the shared directory, so you have to install it separately on the secondary.

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 6.11.0 and ESM 6.11.0 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 all supported operating systems.

Note: The ESM upgrade and ESM with HA upgrade will cause ESM downtime for the duration of the upgrade.

You can upgrade from ESM 6.11.0, if you have the operating system upgraded properly as well.

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 HA Module installation tar file onto the primary system. the file name is `ArcSight-Highavail-7.0.0.xxxx.1.tar`, where the X's are the build number. Do NOT place the installation binary or unpacked content on the shared disk partition (generally `/opt/arcsight`), because it will be deleted during the upgrade process.
2. As user `arcsight`, untar the file:
`tar xf ArcSight-Highavail-7.0.0.xxxx.1.tar`
3. The file `preUpgrade.sh` is in the tar file. Copy it to the secondary system.
4. As user `root`, run `preUpgrade.sh` on the secondary system.
5. If an operating system upgrade is needed, upgrade the operating system version on the secondary. If this is a software installation, see the operating system vendor's documentation for upgrade instructions. For an appliance installation, see ["Upgrade HA Appliance Operating System" on page 78](#). If the operating system is upgraded, be sure to download the HA support packages for that operating system and install them. Otherwise, skip to step 7.
6. Reboot the secondary system.
7. Download the ESM tar file to the primary system. The file name is `ArcSightESMSuite-7.0.0.xxxx.1.tar`.
8. As user `arcsight`, untar the file:
`tar xf ArcSightESMSuite-7.0.0.xxxx.1.tar`
9. On the primary, as user `root`, run `Tools/stop_services.sh` from the tar file. This will shut down ESM.
10. As user `root`, run `preUpgrade.sh` on the primary system.
11. If an operating system upgrade is needed, upgrade the operating system version on the primary. If this is a software installation, see the operating system vendor's documentation for upgrade

instructions. For an appliance installation, see "[Upgrade HA Appliance Operating System](#)" on [page 78](#). If the operating system is upgraded, be sure to download the HA support packages for that operating system and install them. Otherwise, skip to step 13.

12. Reboot the primary system.
13. On the primary system, as user *arcsight*, execute the file `ArcSight-Highavail-7.0.0.1104.1.tar` 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.
14. 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`.
15. On the primary system, upgrade to the supported ESM version. See the *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/stop_services.sh` to stop the ArcSight services.

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

16. 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.
17. As the root user, start the ArcSight services by executing the command:

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

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

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

19. 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.

Verifying the HA and ESM Upgrade

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

1. Make sure that both systems have the current version of the operating system timezone package installed. This is a requirement for ESM. For instructions, refer to the *ESM Installation Guide*, specifically the topic "Install Time Zone Package" in the "Installing ESM" chapter. The timezone

package is not installed within the shared directory, so you have to install it separately on the secondary.

2. 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.

3. 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 53](#) for details about the command arguments available.

Chapter 7: 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 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 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 8: An Example HA Implementation

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

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

Server Configuration

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

- 2 TiB of RAID 10 storage is provided via 15K RPM disks.
- The network interface runs at 1 GB.
- One 1 GB interface on each server will be interconnected by a cable.
- RedHat 7.4 is used with ESM 7.0 software with the HA Module.
- The company's internal DNS server is used for name-to-address translation for the cluster. This is generally the best choice, because there can be thousands of connectors, and dozens of ESM clients. Changing the ESM hostnames on this many machines would be difficult.
- Linux configuration files are used to define the hostname, the IP addresses for each interface, DNS server addresses, and the default route. In a corporate environment, a more common choice would be to set these values via DHCP. For the purposes of this example it is convenient to configure these on the machine directly, so what is going on can be seen. In any case, it is likely that the interconnect ports would be statically defined, since they connect to each other, and do not have access to a DHCP server.
- The shared disk partition and the metadata partition are allocated space via the Logical Volume Manager (LVM). This is strongly recommended that you use Logical Volume Manager (LVM) tools to manage disk space. It will be much easier for you to increase the disk space later using LVM tools.

Initial Setup and Installation

Hardware

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

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

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

DNS Setup

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

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

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

Operating System Installation

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

The root (/), swap, and boot partitions should be physical partitions allocated during installation. Allocate 20 GiB (generous) for root, 8 GiB (minimum) for swap, and 2 GiB for boot. The remaining disk space can be put into a single LVM volume group (vg00) for later allocation to support ESM.

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

Disk Partition Setup

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

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

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

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

To mount the /tmp partition, run:

```
mount /tmp
```

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

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

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

Create this partition with the following command:

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

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

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

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

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

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

Note that we use the `inode64` option here. That is a good idea for very large file systems – but probably this filesystem is large enough to benefit. In any case, if you have any special mount options you want, mount your filesystem with them if you want them to be used after the HA installation.

Interconnect Cable Setup

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

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

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

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

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

Set Up Connected Hosts

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

Install ArcSight Software

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

Install HA Module

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

| Parameter | Value |
|--------------------|----------------------------------|
| Shared Disk | /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 |
| 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:

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

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

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

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

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

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

```
./arcsight_cluster increaseDisk
```

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

```
xfs_growfs /dev/drbd1
```

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

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

Chapter 9: Maintain and Monitor the Cluster System

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

- [The arcsight_cluster Script53](#)
- [Log Output61](#)
- [Changing Hostname, IP Address, or Service IP61](#)
- [Replacing a Server66](#)
- [Changing Mount Options67](#)
- [Setting Configurable HA ModuleProperties 67](#)

The arcsight_cluster Script

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

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

Command Syntax

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

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

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

clusterParameters

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

| | |
|---|---|
| <ul style="list-style-type: none"> connected hosts | <ul style="list-style-type: none"> ping attempts |
| <ul style="list-style-type: none"> connectivity down timeout | <ul style="list-style-type: none"> ping timeout |

diagnose

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

- Checks for communication problems between the nodes.
- Suggests ways to bring nodes that are offline to online mode.
 - a. Detects if `arcsight_cluster offline` has been used to take a node offline, and if so, recommends using `arcsight_cluster online`.
 - b. Suggests that you run `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 75](#) for additional steps to evaluate whether that is the case.

increaseDisk

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

To increase the size of disk:

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

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

$$\text{size}=(P/327)+1$$

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

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

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

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

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

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

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

3. On the primary system run:

```
./arcsight_cluster increaseDisk
```

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

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

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

offline

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

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

To bring it back on line use the online action.

online

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

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

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

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

status

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

Status Output Example

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

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

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

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

```
Started ESM  
Started 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
```

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

If the secondary was offline or its heartbeat function stopped, these lines would look like this:

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

Disk Status

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

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

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

If the secondary was offline or its 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. |

| Disk State | Description |
|------------|--|
| 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, its 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 80](#). (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 55](#) 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 the previous page, 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" on the next page](#)

["Changing the Secondary Hostname or IP Address only" on page 63](#)

["Changing the Primary Hostname or IP Address Only" on page 64](#)

["Changing Both Server Hostnames or IP Addresses" on page 64](#)

["Changing the Interconnect IP Address" on page 66](#)

Changing the Cluster's Service IP Address

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

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

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

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

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

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

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

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

```
./arcsight managersetup
```

This opens the Manager's setup wizard.

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

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

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

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

Run this command about once a minute. Go on to the next step when you see the line "manager service is available".

6. Make sure you can start the ArcSight Command Center by browsing to the following URL:
`https://<hostname>:8443/`
Where <hostname> is the new hostname (note that hostnames with underscores do not work on IE, so use the IP address.)
7. Import the Manager's newly-generated certificate on all clients (ArcSight Console and connectors) that access the Manager. Use `keytoolgui`. See the "SSL Authentication" section of the ESM Administrator's Guide for details about this tool. Use `runcertutil` if you are running ESM using FIPS mode. See "Tools Used to Configure Components in FIPS" in the ESM Administrator's Guide for details about the `runcertutil` tool.
8. Test to make sure that:
 - The clients can connect to the Manager.
 - Peer configuration works as expected. If not, redo the peer configuration.

Changing the Secondary Hostname or IP Address only

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

1. Run the following 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.

Changing Both Server Hostnames or IP Addresses

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

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

```
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). Wait for the shared disk to complete its sync. When run this command as user *root*:
`/usr/lib/arcsight/highavail/bin/arcsight_cluster status`
the `Disk` line in the status information should read:
`Disk: Connected UpToDate/UpToDate`
Note that it may take some time for the sync to complete.
5. Run the following command on the primary (System A) as user *root*:
`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. As user *root* on the secondary system, run either `systemctl stop heartbeat` or `service heartbeat stop`.
2. Change to the `/etc/sysconfig/network-scripts` directory.
3. Select and edit the file for the network interface that you want to change by changing the `IPADDR` value. For example the file might be `ifcfg-eth1`.
4. Run the `ifdown` and `ifup` commands (for example, `ifdown eth1; ifup eth1`).
5. Run the First Boot Wizard on the primary system and specify the new interconnect cable IP address (es).

Replacing a Server

This topic describes how to use the First Boot Wizard to replace a server (for example, if it has hardware 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 ["Installing HA with an Existing ESM" on page 30](#). 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 32](#) 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
```

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 10: Troubleshooting the Systems

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

- [Installation Issues and Solutions](#)68
- [General Problems](#)71
- [Audit Events](#)72
- [Failover Triggers](#)73
- [Processes Killed During Failover](#)74
- [System does not Failover](#)74
- [System Fails Over for no Reason](#)74
- [Network Interface Commands Stall Disk Mirroring](#)74
- [No ESM Uninstall Links on the Primary](#)75
- [Stopping the Network on the Secondary Kills ESM](#)75
- [Disks on Cluster System Fail to Connect](#)75

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. |
| 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. |

| Installation Message | Description |
|---|--|
| No arcsight user on secondary. Please create one identical to that on primary | Create a user <i>arcsight</i> 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. Follow the instructions for that error as well. |
| <metadata_vol> appears to be in use. See the following output from <pre>blkid -o export <metadata_vol></pre> --- blkid output here --- If this volume is not in use, run <pre>dd if=/dev/zero of=<metadata_vol></pre> as user root on hostname <hostname> 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. |

| Installation Message | Description |
|---|--|
| 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 <code>./arcsight_cluster diagnose</code> 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> is 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. |
| 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, " <code>./arcsight_cluster diagnose</code> ", 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. |
| unsupported kernel version <version> | The kernel version on this server does not correspond to an operating system supported by HA. Upgrade the operating system to a supported version. |
| ERROR installing RPMs - Please check the log file <logfile> on <hostname> for details about the error. | RPMs failed to install. Check the log for a detailed message, and correct the problem described there. |

| Installation Message | Description |
|---|--|
| 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. |
| <host> does not resolve to <IP>/ | Correct DNS or /etc/hosts so <host> resolves to <IP> on the server. |
| 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). |
| remove added message of the day or login string from root logins. Expected "test" saw <returnedString> | A <i>message of the day</i> string has been detected. This may cause problems with SSH communication. Disable the SSH banner by creating an empty hushlogin file in the root user's home directory: # touch /root/.hushlogin. |
| 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.

Changing ESM to IPv6

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

Audit Events

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

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

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

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

highavailability:100

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

Severity: 3

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

highavailability:200

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

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

Severity: 7

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

highavailability:300

This event occurs when the Distributed Replicated Block Device (DRBD) storage system begins the process of synchronizing the primary and secondary hard drives and continues every five minutes (by default) until the synchronization is complete. Each event includes the amount of data between the two systems that has been synchronized as a percentage until it reaches 100 percent. You can change the interval using the `highavailability.notification.interval` property as described in ["Setting Configurable HA ModuleProperties" on page 67](#).

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 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 its services without triggering a failover. For example, if you change a property in the `server.properties` file and have to start the Manager again, 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 80](#).

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 67](#), 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

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 Guide* topic entitled “Uninstalling ESM.”

Stopping the Network on the Secondary Kills ESM

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

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

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

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

Disks on Cluster System Fail to Connect

In this scenario, the disk status will be either `WFCConnection` or `Standalone` on both systems. The command `./arcsight_cluster diagnose` will clear this condition in simple cases (see details about ["diagnose" on page 55](#)). 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 highavail.properties File

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

If you are installing the HA Module for the first time, this file will not exist. If you want to use it with the First Boot Wizard or `prepareHA.sh` script, you must create it with a text editor. Copy and rename the `template.properties` file, located in the "Tools/highavail" directory where you unpacked the ESM 7.0 Installation Package. The following example provides guidance about how to define each property value. The actual values will be unique to your deployment environment.

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

Appendix B: Upgrade HA Appliance Operating System

This Appendix provides information on how to upgrade the ESM High Availability Module (HA) from RHEL 7.3 to RHEL 7.4 on an appliance.

If you are running software ESM on your own hardware, consult your operating system vendor for information on how to upgrade the operating system.

Verify Operating System Upgrade File

Download files from <https://softwaresupport.softwaregrp.com/>.

You will need to download the upgrade script and the HA Support packages files, which are `esm_osupgrade_rhel74_20180727112006.tar.gz` and `esm_ha_support_rpms_rhel74.tar.gz`, respectively.

After you download the software, contact support to verify that the signed software you received is indeed from Micro Focus and has not been manipulated by a third party.

Upgrade HA Operating System

Perform this procedure each appliance when you are instructed to upgrade the Operating System (see the procedure in Chapter 6).

1. Log in to the system as user *root*.
2. Create a temporary directory for the download. (Do not put it in the `/opt` directory. An example is `/downloadtmp`.
`mkdir /downloadtmp`
3. Change to the directory you created. For example::
`cd /downloadtmp`
4. As user *root*, transfer the files `esm_osupgrade_rhel74_20180727112006.tar.gz` and `esm_ha_support_rpms_rhel74.tar` to the target system and place it in the directory you created.
5. From the directory where you put the archives in step 4, extract `esm_osupgrade_rhel74_20180727112006.tar.gz` and `esm_ha_support_rpms_rhel74.tar.gz` as follows:
`/bin/tar xzf esm_osupgrade_rhel74_20180727112006.tar.gz`
6. Change directory:
`cd esm-rhel74upgrade`

7. Run the following command to make the script executable:
`chmod 0700 osupgrade`
8. Run the following command to start the upgrade:
`./osupgrade 2>&1 | tee osupgrade.log`
9. Make sure the system is rebooted after the script completes.
10. Check the operating system version by running the following command:
`cat /etc/redhat-release`

The result of this command should be:

```
Red Hat Enterprise Linux Server release 7.4
```

The RHEL 7.4 upgrade is now completed on this appliance.

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 not 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:

$$([Ping\ Attempts] + [Ping\ Timeout]) * (1 + \# \text{ of Connected Hosts}) < 90$$

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 computer, click the link above and an email window opens with the following information in the subject line:

Feedback on ESM High Availability Module User's Guide (ESM 7.0 Patch 1)

Just add your feedback to the email and click send.

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

We appreciate your feedback!