



# HP ArcSight ESM High Availability Module

Software Version: 1.0

## ESM High Availability Module User's Guide

December 23, 2014



## Legal Notices

### Warranty

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

The information contained herein is subject to change without notice.

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

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

This document is confidential.

### Restricted Rights Legend

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

### Copyright Notice

© Copyright 2015 Hewlett-Packard Development Company, L.P.

Follow this link to see a complete statement of copyrights and acknowledgements:

<http://www.hpenterprisesecurity.com/copyright>

## Support

### Contact Information

<b>Phone</b>	A list of phone numbers is available on the HP ArcSight Technical Support Page: <a href="https://softwaresupport.hp.com/documents/10180/14684/esp-support-contact-list">https://softwaresupport.hp.com/documents/10180/14684/esp-support-contact-list</a>
<b>Support Web Site</b>	<a href="https://softwaresupport.hp.com">https://softwaresupport.hp.com</a>
<b>Protect 724 Community</b>	<a href="https://protect724.hp.com">https://protect724.hp.com</a>



# Contents

Chapter 1: Introduction .....	6
General Requirements .....	6
Chapter 2: Preparing for Installation .....	7
Independent Steps to Ensure Availability .....	7
Hardware Requirements .....	7
Crossover cables .....	8
Other Hosts to Ping .....	8
Software requirements .....	8
Port Usage .....	9
Ports and Protocols to Keep Open .....	9
Licensing .....	9
ESM Is Already Installed .....	10
Preparation .....	10
Initial-Synchronization Speed .....	10
Primary and Secondary Changes .....	11
Primary-Only Changes .....	13
ESM is Not Already Installed .....	13
Primary and Secondary Changes .....	14
Primary-Only Changes .....	16
Chapter 3: HA Module Installation .....	17
Introduction .....	17
Running the installation Script .....	17
Running the First Boot Wizard .....	18
Verify HA Installation .....	25
After Installing the HA Module .....	26
ESM is Not Installed .....	26
ESM was Already Installed .....	26
Configurable Properties .....	28
Uninstalling the HA Module .....	28
Uninstalling ESM and HA .....	29
Uninstalling HA Only .....	29



Chapter 4: Maintenance and Monitoring .....	31
arcsight_cluster Script .....	31
Command Syntax .....	31
clusterParameters .....	33
diagnose .....	33
forcePrimary .....	34
increaseDisk .....	34
offline .....	35
online .....	36
prefer .....	36
status .....	36
Status Output Example .....	36
Status Output Explanation .....	37
tuneDiskSync .....	39
Log Output .....	40
Changing Hostname, IP Address, or Service IP .....	41
Changing the Cluster's Service IP Address .....	41
Changing the Secondary Hostname or IP Address Only .....	43
Changing the Primary Hostname or IP Address Only .....	43
Changing Both Server Hostnames or IPs .....	44
Changing the Interconnect IP Address .....	44
Replacing a Server .....	45
Changing Mount Options .....	45
Chapter 5: Example HA Implementation .....	46
Requirements .....	46
Initial Setup and Installation .....	47
Hardware .....	47
DNS Setup .....	47
Operating System Installation .....	47
Disk Partition Setup .....	48
Interconnect Cable Setup .....	48
Set Up Connected Hosts .....	49
Install ArcSight Software .....	50
Increase Disk Space .....	51
Chapter 6: Troubleshooting .....	53
Installation Issues and Solutions .....	53



- General Problems .....56
- Audit Events .....57
  - highavailability:100 .....57
  - highavailability:200 .....57
  - highavailability:300 .....58
  - highavailability:500 .....58
- Failover Triggers ..... 58
- Processes Killed During Failover .....59
- System does not Failover ..... 59
- Network Interface Commands Stall Disk Mirroring .....59
- No ESM Uninstall Links on the Primary .....60
- Neither Server will Come Up as Primary .....60
- Stopping the Network on the Secondary Kills ESM .....60
- Send Documentation Feedback .....61



# 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 communications or operational problems.

The HA Module is installed on the primary of adjacent machines connected by an Ethernet crossover cable (or more than one, if you are using bonded interfaces). The HA module replicates the installation and all data by mirroring the hard disk partition to the secondary machine. Both ESMs are configured to use the same Service IP Address.

Ordinarily one ESM runs on the primary machine and selected hard-disk writes are mirrored to the secondary machine. The HA Module monitors the health of ESM. When a failover is triggered, the HA option starts the secondary ESM, which takes over. During the failover, events are cached at the connectors, so that no data is lost.

Risk Insight only supports English and ESM and its components must be set the same way.

## General Requirements

The following are the general requirements for the HA Module. Specifications, model numbers, version numbers and capacities are specified in the ESM High Availability Module Product Lifecycle Document.

- The supported operating system is specified in the ESM High Availability Module Product Lifecycle Document (PLD).

**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 will not work properly. Do not upgrade to a newer version of your operating system until there is a version of HA that supports it.

- The network interface card should be at 1 GB and a cable that supports it, but we recommend a 10 GB card.
- The primary and secondary machines must be close enough together that the internet connection between them requires no intervening routers or switches.

Additional details on these and other pre-installation requirements are covered in ["Preparing for Installation" on page 7](#).



## Chapter 2: Preparing for Installation

This section covers the preparations and planning required for the ESM High Availability Module which is somewhat different for a fresh ESM installation than it is when installing the HA Module and upgrading an existing ESM to the latest version. There are hardware considerations, optional devices, and tasks to perform on the servers you are using.

Independent Steps to Ensure Availability .....	7
Hardware Requirements .....	7
Software requirements .....	8
ESM Is Already Installed .....	10
ESM is Not Already Installed .....	13

Refer to ["Example HA Implementation" on page 46](#) for an example of how to plan and set up an HA installation.

If you already have ESM and are licensed for the existing High Availability solution, this is a new High Availability module. It requires a new ESM license that supports it. The new High Availability module uses software to manage failovers and a different hardware configuration. Read this guide and follow the instructions for upgrading ESM and installing the HA module as if you are a new HA user.

This sections covers the hardware and software requirements that need to be in place before the installation begins.

### Independent Steps to Ensure Availability

There are several things you can do independently of the HA module to help ensure continued availability for ESM.

- For the primary and secondary machines, provide separate power supplies that rely on different circuitry and power sources.
- Network redundancy (but with no intervening routers or switches) will enable you to fail over if the primary loses its connection to the network.
- Make use of application management software to notify you of any issues with the primary or secondary servers themselves.

### Hardware Requirements

The HA Module requires two identical machines that conform to the latest ESM version hardware and software requirements.



- The network interface card should be at 1 GB or higher and a cable that supports it.
- The primary and secondary machines must be close enough together that the internet connection between them requires no intervening routers or switches.

These servers must be part of the same IPV4 subnet.

## Crossover cables

The machines must be connected using one or more crossover cables with at least 1 Gigabyte ethernet cards. You can set up more than one crossover cable if you are using bonded interfaces. The use of bonded interfaces are recommended to provide a redundant link between the two servers. If your servers have very high speed disk subsystems you may see improved performance with one or more 10 GB interfaces. The mirrored disk performance is limited by the slower of the disk write throughput and the throughput on the crossover link.

See ["Disk Partition Setup" on page 48](#) for an example of how to configure bonded interfaces.

## Other Hosts to Ping

Set up some host on the network or designate existing hosts for the HA cluster to ping as a check for network connectivity. You specify these host IP addresses when you run the First Boot Wizard.

## Software requirements

The primary and secondary machines must be substantially the same. There is some variation in the procedure depending on whether you are upgrading a previous version of ESM or installing ESM and the HA Module for the first time. See ["General Requirements" on page 6](#).

Both of the servers must be on the same IPV4 subnet. In addition, a third IP Address on the same subnet should be reserved to act as the IP Address of the ESM. This is the "Service IP."

Both hosts must be configured to access a yum server. This can either be one of the standard yum servers provided by the OS vendor, or a local server created from the OS ISO or CD. A YUM server automatically determines the required libraries, so no specific list needs to be provided here.

We recommend that you use Logical Volume Manager (LVM) to manage volumes and partitions in the HA cluster.

We recommend you use DNS to manage IP addresses and host names for all the components in the HA cluster.



## Port Usage

Each of the HA servers uses the following ports in addition to those used by ESM (see the *ESM Installation and Configuration Guide*, Chapter 2, "Installing ESM"):

- 694 (udp)
- 7789 (tcp)

## Ports and Protocols to Keep Open

The protocols and ports listed below are used by the HA system, and must not be blocked. This can be done by running `service iptables stop` on both the primary and the secondary HA machines, and by setting 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 HA to verify that it is still on the network.

Protocol	From IPs	From Port	To IPs	To Port
ICMP	Primary Secondary	N/A	Primary Secondary service connected hosts	N/A
TCP	Primary cable Secondary cable	Any	Primary cable Secondary cable	7789
UDP	Primary Primary cable Secondary Secondary cable	Any	Primary Primary cable Secondary Secondary cable	Any

## Licensing

The license file for HA is an ESM license file with the HA Module included. If you already have ESM installed, obtain a new ESM license that includes the HA Module. After upgrading ESM, install the new ESM/HA license as described in the *ESM Administrator's Guide*, Chapter 2, "Configuration." The topic is "Installing New License Files Obtained from HP."

If ESM is not already installed, you specify the same ESM/HA license file when you install the HA Module and then after that, when you install ESM. Refer to the *ESM Installation and Configuration Guide* for information for ESM. For HA installation, see ["HA Module Installation" on page 17](#).



## ESM Is Already Installed

If ESM is already installed on the Primary, do not install it on the secondary machine. The following is an overview of the procedure:

1. Upgrade to the latest version of ESM that supports this HA module and the appropriate operating system version, if necessary. Refer to the ESM Upgrade Guide for details.
2. Install the ArcSight ESM HA Monitoring foundation package in the ArcSight Console. See the Standard Content Guide for installation instructions. Installing this package before installing the HA Module is essential to acquiring up-to-date HA status information from the outset. If you install this package later, there is no status information available until an HA event occurs, which could be a long time.
3. Complete the preparations, and primary and secondary, and primary-only changes described in the topics that follow.
4. Install the HA Module, as described in ["HA Module Installation" on page 17](#).
5. Complete the post-install steps in ["After Installing the HA Module" on page 26](#).

After you install the HA Module, it mirrors the ESM installation to the secondary. Many of the preparations listed in ["Primary and Secondary Changes" on the next page](#) and ["Primary-Only Changes" on page 13](#) are part of a normal ESM installation and may already be set correctly on the primary machine, but you need to complete these on the secondary machine.

## Preparation

This topic assumes you have upgraded ESM and the operating system to the latest supported versions, before you install the HA Module.

The network interface cards used for the interconnection of the two servers should run at 1 or 10 GB. The benefit of 10 GB is seen for the most part only when synchronizing a primary on which ESM is already installed and has data that also must be synchronized. When the HA module is installed, it synchronizes the shared volumes, including the existing ESM data and files as well as empty space. This may take two or more days, if you have 8 TB of data and depending on the speed of your network interface card. You can run ESM on the primary during this time, but the secondary is not ready to take over until the synchronization is complete.

## Initial-Synchronization Speed

If you install HA on an existing ESM, the entire mirrored partition on the existing ESM primary has to be synchronized to the secondary. This may take two days or more if you have 8 TB of data and depending on the speed of your network interface card. The synchronization speed is determined by the slower of the disk I/O rate and the data transfer rate across the cable.



Typical ESM installations use very fast server class disks, which can be much faster than a 1G cable. In such cases providing multiple bonded 1G interfaces, or one or more bonded 10G interfaces may lead to noticeable reductions in the time required for the initial synchronization.

SSD drives (Fusion, for example) provide an additional speed-up of the synchronization above and beyond what can be expected because they are fast. SSD drives require and support TRIM to manage their free space. The HA disk synchronization is TRIM-aware; it can use TRIM to identify free blocks of 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 HA, then the TRIM information is passed to the SSD drives by way of the HA disk synchronization. the disk synchronization uses this information to detect which blocks are free and skips these blocks. For this example, only 4 TB of data would need to be synchronized, instead of 12.

## Primary and Secondary Changes

Make the following changes on the primary and on the secondary machines:

- Both servers must be running either RHEL 6.5 or CentOS 6.5. The HA Module does not support SUSE Linux.
- Once ESM 6.8c is installed on the primary, it should have the correct time zone data package installed. Install this package manually on the secondary. For instructions, refer to the topic "Install Time Zone Package" in the "Installing ESM" chapter of the *ESM installation and Configuration Guide*.
- Set up both servers to run the Network Time Protocol (NTP) so that the system time is kept synchronized between them.
- Connect the two servers with crossover cables. Configure the interfaces with IPv4 addresses. (HA does not support IPv6.) It is recommended that you set the operating systems up to use a bonded interface to connect the two servers. Linux allows you to bind multiple network interfaces into one channel. Use the bonding kernel module and a channel-bonding network interface. This allows the network interfaces to act together for redundancy and added bandwidth. Consult the operating system documentation for information on setting this up.
- The password for user *root* must be the same on each machine during installation, but you may change the root passwords after installation.
- If the shared disk is an SSD drive such as Fusion, make sure you have TRIM support configured.
- Create the *arcsight* user on the secondary using the same home directory, UID, and GID as already exist on the primary.

On the primary, run this command to get the UID and GID for user *arcsight*.

```
id arcsight
```

In the output, make a note of the numerals that follow `uid=` and `gid=`.



On the secondary, run these commands as user *root* to create a user and group called *arcsight*. For <GID> and <UID> use the values from the `id arcsight` command you ran on the primary.

```
groupadd -g <GID> arcsight
```

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

Be sure to specify the same password that you specified on the primary.

- The execute permission for `/sbin` should be set on the secondary as it is on the primary:

```
chmod 777 /sbin
```

- The HA Module mount the file system on the secondary exactly as you mounted it on the primary, so make sure all the options are set the way you want them. The mount options, or even whether the file system is mounted on the secondary, do not matter for installation.
- Meet the ESM `/tmp` partition size requirements on both machines (currently 3 GB).
- RHEL (or CentOS) 6.5 TCP buffer sizes are tuned for links with speeds of 10 GB/sec or less. If you are using bonded interfaces with a total speed of greater than 10 GB/sec, you will need to add lines to `/etc/sysctl.conf` to increase these values. The following lines, which approximately double the maximum TCP buffer sizes will support bonded interfaces with speeds up to 20 GB/sec:

```
net.core.rmem_max = 250000  
net.core.wmem_max = 250000
```

Run `sysctl -p` as user *root* to make these changes take affect immediately.

- The user process limit should already be increased on the primary, repeat the procedure on the secondary:
  - a. If you do not already have a file called `90-nproc.conf` in the `/etc/security/limits.d` folder, create it (and the `limits.d` directory, if necessary). If the file already exists, delete all entries in the file.
  - b. Add the following lines:

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

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

- c. Reboot the secondary machine.



- d. Log in as user *arcsight*.
- e. Run the following command to verify the new settings:

```
ulimit -a
```

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

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

- The file system where ESM is installed can be EXT4 or XFS. You cannot change it while installing the HA Module or during an ESM upgrade.
- Create a metadata partition

Create a small partition on each server for disk-synchronization metadata. The size in mebibytes (MiB, 1,048,576 bytes) can be calculated as  $\text{size}=(P/32768)+1$ , where P is the size of the mirrored partition in mebibytes. For example, if the mirrored partition size is 1 TiB (P=1,048,576 MiB), the metadata partition size would be 33 MiB. See ["Disk Partition Setup" on page 48](#) for an example of how to do this.

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

## Primary-Only Changes

When installing over an existing ESM, after it is upgraded, stop it before installing the HA module:

1. As the user *root*, stop ESM by running:

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

2. As the user *root*, create the folder */usr/lib/arcsight* and make user *arcsight* the owner:

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

## ESM is Not Already Installed

If ESM is not installed on either the primary or the secondary, do not install it first. When ESM is not already installed, the shared volumes are empty and the initial synchronization step is not needed. When you install ESM later, the new files and folders are mirrored as they are installed. Perform the installation tasks in the following order:



1. Complete the preparations and primary and secondary changes described below.
2. Install the HA Module, as described in the next chapter.
3. Install ESM.  
Be sure to include the Foundation package called "ArcSight ESM HA Monitoring," when you get to that step in the Configuration Wizard. Installing this content package with ESM is essential to acquiring up-to-date HA status information from the outset. If you install this package later, there is no status information available until an HA event occurs, which could be a long time.

Make sure that when the ESM installer asks you for the Manager host name, that you enter the service host name for ESM and NOT the host name of the actual machine.

For the HA Module, you make the following preparations before you install the HA Module. Most of the following steps are included as part of the ESM installation procedure, so when you start installing ESM, you can skip any steps that you perform in the next two topics.

## Primary and Secondary Changes

Make the following changes on the primary and on the secondary machines:

- Both servers must be running either RHEL 6.5 or CentOS 6.5. The HA Module does not support SUSE Linux.
- Once ESM 6.8c is installed on the primary, it should have the correct time zone data package. Install this package manually on the secondary. For instructions, refer to the topic "Install Time Zone Package" in the "Installing ESM" chapter of the *ESM installation and Configuration Guide*.
- Set up both servers to run the Network Time Protocol (NTP) so that the system time is kept synchronized between them.
- Connect the two servers with crossover cables. Configure the interfaces with IPv4 addresses. (HA does not support IPv6.) It is recommended that you set the operating systems up to use a bonded interface to connect the two servers. Linux allows you to bind multiple network interfaces into one channel. Use the bonding kernel module and a channel-bonding network interface. This allows the network interfaces to act together for redundancy and added bandwidth. Consult the operating system documentation for information on setting this up.
- The password for user root must be the same on each machine during installation, but you may change the root passwords after installation.
- If the shared disk is an SSD drive such as Fusion, make sure you have TRIM support configured.
- Create the *arcsight* user on the secondary machine using the same home directory, UID, and GID as on your existing ESM server. Do this as user *root*:



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

If 500 is not available for user and group ids, pick any free ID. The user and group ids must be the same on both machines, or the installation will fail, so when you install ESM on the primary, make sure you use these same ID values.

- The execute permission for `/sbin` should be set on both machines.

```
chmod 777 /sbin
```

- Meet the ESM `/tmp` partition size requirements on both machines (currently 3 GB).
- Increasing the user process limit should be done on both machines as follows:
  - a. If you do not already have a file called `90-nproc.conf` in the `/etc/security/limits.d` folder, create it (and the `limits.d` directory, if necessary). If the file already exists, delete all entries in the file.

- b. Add the following lines:

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

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

- c. Reboot the machine.
- d. Log in as user *arcsight*.
- e. Run the following command to verify the new settings:

```
ulimit -a
```

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

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

- Select the partition to be mirrored between the two servers. The file system of this partition can be either EXT4 or XFS. Typically this is the partition mounted as `/opt` for your ESM installation. Use the command `df /opt/arcsight` to obtain the partition. This partition must exist on both the primary and secondary and must have the same device name, be mounted at the same location, and be the



same size.

If the partition is not mounted as `/opt` or `/opt/arcsight`, then create a symbolic link from `/opt` or `/opt/arcsight` on both the primary and the secondary.

Note that after installation, this partition is only mounted on the primary, and only that primary can make changes to it.

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. See ["Disk Partition Setup" on page 48](#) for an example of how to do this.

- Create a metadata partition

Create a small partition on each server for disk-synchronization metadata.

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

$size=(P/32768)+1$

... where P is the size of the mirrored partition in mebibytes. For example, if the mirrored partition size is 1 TiB (P=1,048,576 MiB), the metadata partition size would be 33 MiB. See ["Disk Partition Setup" on page 48](#) for an example of how to do this.

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

## Primary-Only Changes

When installing ESM for the first time create the folders `/opt/arcsight` and `/usr/lib/arcsight`. Set set the ownership for both to user `arcsight`. before you install anything.

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

Create these folders logged in as user `root` on the primary server. This change is mirrored to the secondary after the HA Module is installed, assuming your mount point for the mirroring is either `/opt` or `/opt/arcsight`.



## Chapter 3: HA Module Installation

This section covers installing the ESM High Availability Module. The First-Boot Wizard enables you to configure the HA Module. It runs automatically following the HA Module installation or you can run it by itself at any time to make certain changes to your HA Module configuration.

Introduction .....	17
Running the installation Script .....	17
Running the First Boot Wizard .....	18
Verify HA Installation .....	25
After Installing the HA Module .....	26
Configurable Properties .....	28
Uninstalling the HA Module .....	28

### Introduction

It is assumed that you have already completed all the required tasks for the primary and secondary machines as described in ["Primary and Secondary Changes" on page 11](#).

The First Boot Wizard automatically runs when the initial installation steps finish, so it might appear as a seamless operation. When the First Boot Wizard finishes, it invokes a script that checks that the configuration is complete and correct, reports any inconsistencies with the location of logs to aid fixing them. It completes the installation with the specified configuration, if there are no inconsistencies. You can run the First Boot Wizard separately at any time, to make changes.

The degree to which the two servers match is important and the installation examines every relevant characteristic 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.

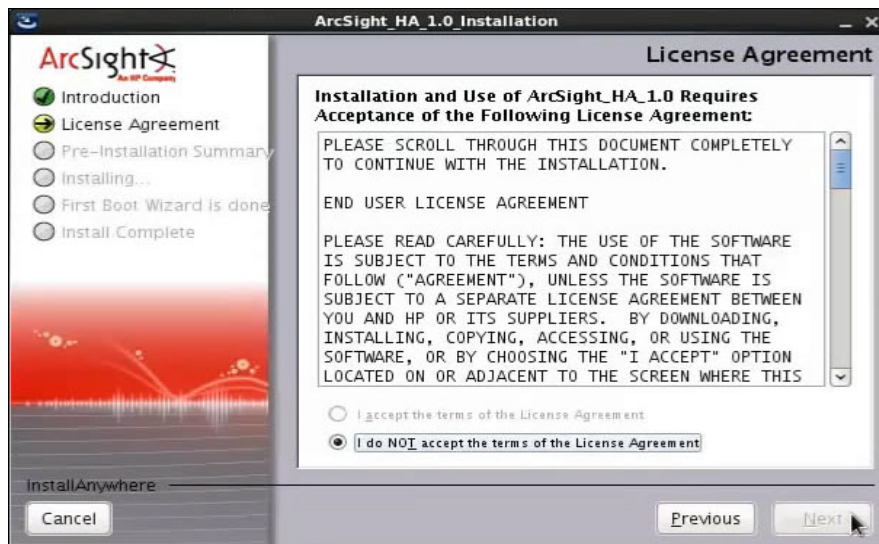
**Note:** The contents of the shared disk on the secondary machine will be removed during installation, so make sure it does not contain data of any value.

### Running the installation Script

To run the installation



1. Download the HA Module installation binary file, `ArcSight-Highavail-1.0.0.1039.0.bin` to the `/home/arcsight` directory. (Do not download it to the shared disk.)
2. Download the License file that includes the HA Module.
3. Run the installer. For example, if you downloaded it to a Downloads directory, run  
`Downloads/ArcSight-Highavail-1.0.0.1039.0.bin`
4. Click Next on the **Introduction** dialog.
5. On the license dialog the radio button for agreeing to the license is grayed out until you scroll to the bottom of the license agreement.



Scroll down, reading the license agreement, select **I accept the terms of the License Agreement**, at the bottom, and click **Next**.

The installer displays a summary and then a dialog that shows the progress of the installation. When this part of the installation completes the First Boot Wizard starts automatically.

## Running the First Boot Wizard

1. If you started with the installer, described in the previous topic, this dialog appears automatically, and you can skip this step.

Run the First Boot Wizard as follows:

```
cd /usr/lib/arcsight/highavail/bin
```



```
./arcsight firstBootWizard [--console]
```

If you specify `--console`, the First Boot Wizard runs in console mode, if not it runs in GUI mode, which is the mode described in this topic. If you run it in console mode, it requires the same inputs described for the GUI mode.

Unless otherwise noted, fields are required and omitting them will result in an error message.

The **Welcome to the First Boot Wizard** dialog appears.

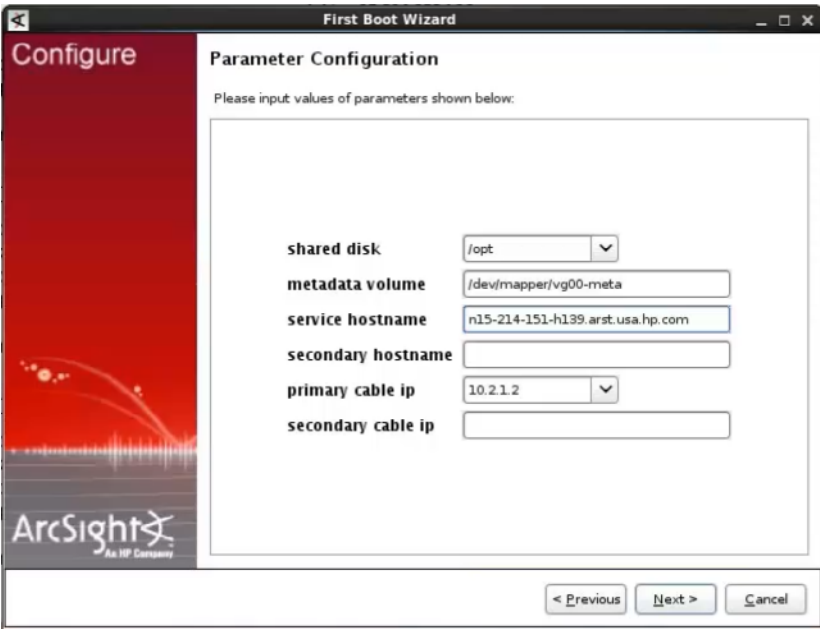


2. At the **Welcome to the First Boot Wizard** dialog click **Next**.
3. On the **License File** dialog, click the browse button (...) and navigate to the directory to which you downloaded the license file for the HA Module and select it. Click **Next** to continue.





4. On the **Parameter Configuration** Page enter the requested information and click Next.





Field	Description
Shared Disk	<p>Enter the mount point of the disk shared between the primary and secondary. The items in the drop-down includes all relevant mount points.</p> <p>Bind mounts are not supported. and are flagged as an error by the installation. Use symbolic links.</p> <p>The contents of the shared disk on the secondary will be <b>completely erased</b>, so make sure it contains no data of any value.</p> <p>Make sure no process on the primary or secondary is using this filesystem or the installation will exit with errors.</p> <p>You cannot change this value on subsequent runs of the First Boot Wizard.</p>
Metadata Volume	<p>Enter the volume containing disk-synchronization metadata. This volume is expected to start with /dev/.</p> <p>The contents of the metadata volume on both the primary and the secondary will be removed.</p>
Service Hostname	<p>Enter the service hostname of the service IP. The IP address works, but we recommend using the service hostname. You can use the <code>hosts</code> file or DNS. This is a virtual hostname that is used to connect to ESM regardless of which physical computer it is running on.</p>
Secondary Hostname	<p>Enter the hostname of the secondary machine. This is the "real" IP for this machine.</p>
Primary Cable IP	<p>Select the IP of the interface connected to the interconnect cable on the primary from the list.</p>
Secondary Cable IP	<p>Enter the IP of the interface connected to the interconnect cable on the secondary</p>

Click **Next** when done.

5. On the second Parameter Configuration dialog, enter the following information:



Field	Description
Preferred Primary	Enter the hostname of the machine you would prefer to be primary. This field is not required. Leave it blank if either machine may be the primary. If you supply a value, HA will always fail over from the non-preferred primary to the preferred primary when both machines are running properly. This may cause additional, unneeded failovers.
Connected Hosts	These hosts are other machines in the network that HA can ping to verify that it is connected to the network. Enter a space-separated list 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 120.
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.

6. Enter the password for user *root* and click **Next**.

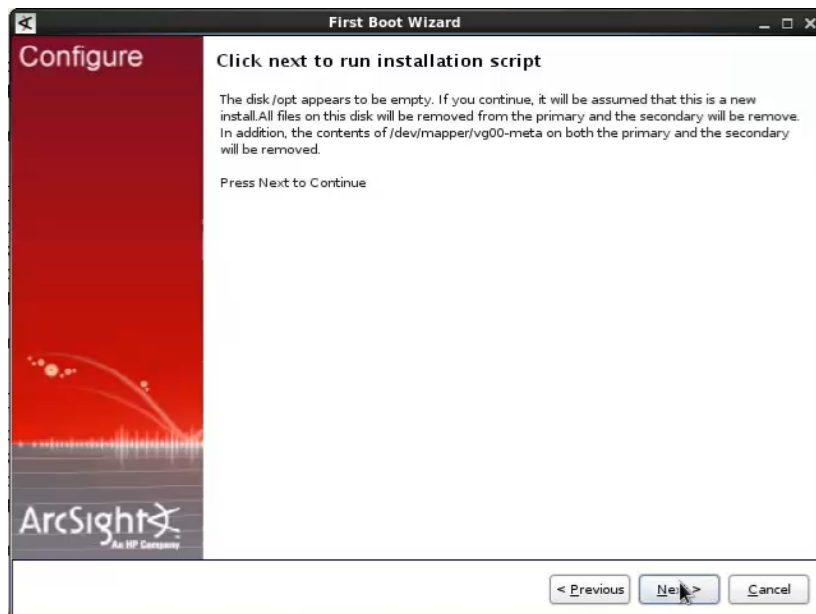




Supplying the password for the *root* user enables the HA installation to handle components and actions that have to be performed as the *root* user. The password must be the same on both machines. This password is not stored permanently.

You may change this password after the installation completes.


7. If your shared disk is empty, you get the following screen:



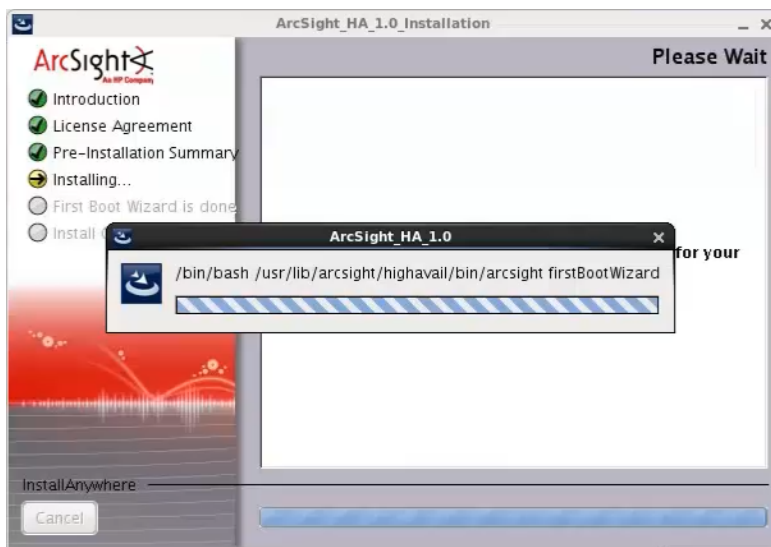


The screen remains while the installation script runs. This may take an hour or so depending on whether you are upgrading an existing ESM. A progress bar indicates how far along it is.

8. When the installation script has finished checking for consistency and ready to run, the following screen appears. Click **Finish**.



The installer runs:



9. When the First Boot Wizard finished it displays the First Boot Wizard is done screen, and shows any relevant messages. Click **Next**.
10. When the Installation is finished it displays the **Install Complete** screen, and shows any relevant



messages. Click **Done**.

11. If there are errors, check both servers for log files. See "[Installation Issues and Solutions](#)" on [page 53](#).

Fix any errors noted in these logs. Then uninstall HA by running the following command as user *root*:

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

Then re-run the installation script to install HA.

## Verify HA Installation

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

1. Check the shared disk partition:

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

Sample output is shown below:

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

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

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

```
./arcsight_cluster status
```

Example output is shown below:

```
Tue Dec 16 14:04:04 PST 2014 OK  
n15-214-132-h127.arst.usa.hp.com: online  
n15-214-132-h16.arst.usa.hp.com: online Primary
```

```
Disk: Connected UpToDate/UpToDate
```

```
OK Network-n15-214-132-h127.arst.usa.hp.com  
OK Network-n15-214-132-h16.arst.usa.hp.com
```

```
Started Failover-Check-n15-214-132-h127.arst.usa.hp.com  
Started Failover-Check-n15-214-132-h16.arst.usa.hp.com  
Started Filesystem  
Started Ping-n15-214-132-h127.arst.usa.hp.com  
Started Ping-n15-214-132-h16.arst.usa.hp.com  
Started STONITH-SSH-n15-214-132-h127.arst.usa.hp.com
```



```
Started STONITH-SSH-n15-214-132-h16.arst.usa.hp.com
Started Service-IP
```

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

## After Installing the HA Module

Now that the HA Module is installed, attend to ESM. The tasks depend on whether ESM is already installed or you are about to install it for the first time.

### ESM is Not Installed

If ESM was not already installed, install it now. See the *ESM Installation and Configuration Guide*.

Be sure to specify the *Service* hostname or IP address. Do not use the ESM default, which is the hostname or IP address of the Manager server.

**Note:** After the installation is complete, 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.

### ESM was Already Installed

If ESM was installed before you installed the HA module, now is the time to switch to the new service hostname or service IP address .

1. Set up the ESM services by running this command as user *root*:

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

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

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

```
fstrim <shared disk>
```

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

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

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



4. Stop ArcSight Web by running (as user *arcsight*):

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

You may get error messages from this command indicating *arcsight\_web* was not stopped. This is normal and may be ignored.

5. 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.
6. Start the Manager by running (as user *arcsight*):

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

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

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

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

8. While logged in as user *arcsight*, run the following to start the setup program for ArcSight Web from the */opt/arcsight/web/bin* directory:

```
./arcsight websetup
```

This will open the ArcSight Web setup wizard.

- a. Enter the new hostname (that you set earlier in this procedure) in every field in which the old hostname occurs.
  - b. When the certificate from the manager is displayed, check the option "Trust the certification from the manager."
  - c. Select the self-signed keypair option when prompted and enter the required information to generate the self-signed certificate with the new hostname .
9. Start ArcSight Web by running (as user *arcsight*):

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



10. Wait two minute to ensure that `arcsight_web` has started.
11. Make sure you can start the ArcSight Command Center by browsing to the following URL:  
  
`https://<hostname>:8443/`  
  
Where `<hostname>` is the new hostname (Hostnames with underscores do not work on IE, so use the IP address.)
12. 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.
13. Test to make sure that:
  - The clients can connect to the Manager.
  - Peer configuration works as expected. If not, redo the peer configuration.

**Note:** 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.

## Configurable Properties

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 Manager for them to take effect. For more information on editing ESM properties files refer to the "Configuration" chapter of the ESM Administrator's Guide.

## Uninstalling the HA Module

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



## Uninstalling ESM and HA

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

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

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

## Uninstalling HA Only

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

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

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

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

3. After the HA uninstall is complete, all the files you need to run ESM are on both of your servers. Select one of those machines as the new ESM server. Run the following command as user *root* on the selected server:

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

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

4. Update the ARP cache:

```
arping -U -s <service_ip> <default_gateway_ip>
```

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

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

At this point ESM is running on its new 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.

6. To make sure the ESM service IP comes up at reboot on the selected server, change the appropriate scripts in */etc/sysconfig/network-scripts/* on that server.

If you would rather use the individual IP or the hostname of this server, this is a good time to do it. Use the method for changing the IP Address of an ESM that is described in the ESM Installation and Configuration Guide. In this case, no changes to the network scripts would be required.



If you use the individual host IP address, you have to change the Manager IP for every connector and Console that connects to this Manager as well as the URL used for ArcSight Command Center.



## Chapter 4: Maintenance and Monitoring

This section covers tasks related to maintaining the HA Module's primary/secondary cluster, and guidelines for monitoring it's health.

arcsight_cluster Script .....	31
Log Output .....	40
Changing Hostname, IP Address, or Service IP .....	41
Replacing a Server .....	45
Changing Mount Options .....	45

### arcsight\_cluster Script

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

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

### Command Syntax

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

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



Actions	<code>clusterParameters [--console]</code>	Update the cluster parameters using the Cluster Parameters Wizard. Only run this on the primary. The <code>--console</code> option displays in console mode. GUI mode is the default.
	<code>diagnose</code>	Checks the system health. If any problems are found it corrects them or suggests how the user can correct them. After correcting a problem, run it again to see if there are any other problems.
	<code>help (or -h)</code>	Provides command usage and HA version.
	<code>forcePrimary</code>	Forces the server where you run this command to become the primary. Typically only used in emergency situations, because it may result in loss of data collected on the other server, but not yet mirrored to this one.
	<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>prefer [hostname]</code>	System uses this server as the primary if both are eligible. If left blank the system only switches primary if the other server becomes ineligible (reduces fail-overs).
	<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 ./arcsight_cluster online myfirstesm.mydomain.com ./arcsight_cluster prefer p12-345-678-q90</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:

• preferred primary	• ping timeout
• connected hosts	• ping attempts
• connectivity down timeout	

## diagnose

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

- Checks for communication problems between the nodes.
- Suggests ways to bring nodes that are offline to online mode.
  - a. Detects if `arcsight_cluster offline` has been used to take a node offline, and if so, recommends using `arcsight_cluster online`.
  - b. Suggests that you run `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.



## forcePrimary

The forcePrimary action provides a means of getting ESM to run when neither server will start up as the primary and run ESM on its own. Such situations are rare. The most common case is when an ESM server has crashed, and cannot be brought up immediately. In that situation, there is a small amount of data that has not been mirrored from the crashed machine to the other. A failover cannot occur if the disks are not in sync.

Under such circumstances, the best option is to make the other machine the primary anyway. Any data that was on the crashed machine but not mirrored, is lost.

The forcePrimary action works only when neither server will come up automatically as the primary, but one server could be used as the primary to run ESM. Specifically:

- There is currently no primary.
- The server you want to be primary is online.
- The local disk is working properly.
- That server can communicate with the ESM clients. The forcePrimary action warns if there are communication problems, but does not prevent the server from becoming the primary.
- The other server is offline or not connected to the cluster.

Because of the potential for data loss, forcePrimary asks the user to confirm that he understands the consequences before proceeding. After that, the server on which you run this command action becomes the primary.

## 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  $\text{size} = (P/32768) + 1$ , where P is the size of the mirrored partition in mebibytes. For example, if the mirrored partition size is 1 TiB ( $P = 1,048,576$  MiB), the metadata partition size would be 33 MiB.

If you ever need to increase the size of the mirrored partition, increase the size of the metadata partition accordingly. Decreasing the size of the mirrored 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.

If you need to increase the size, 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 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 resize it. 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.

In order to take advantage of this increased disk space, you probably need to increase the size of the ESM Default Storage Group. You can do this from the ArcSight Command Center (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 an "Are you sure?" 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. However, an offline server cannot act as secondary while it is off line. That is, even if it is operating normally, it cannot take over as primary in a failover.

To bring it back on line use the online action.



## online

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

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

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

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

## prefer

The prefer action allows you to specify on which server to run ESM, if both are eligible. By default there is no preference; the first available server is used. You can use the prefer action to change the server running ESM. If there is a preferred primary, when it goes down the other server takes over. When it comes back up, it resumes being primary by taking over again. The second failover would not happen if there was no preferred primary set.

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

## 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
```

```
p12-345-678-q90.test.hp.com: online
p12-345-678-q33.test.hp.com: online Primary Preferred
```

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

```
OK Network-p12-345-678-q90.test.hp.com
OK Network-p12-345-678-q33.test.hp.com
```

```
Started ESM
Started Failover-Check-p12-345-678-q90.test.hp.com
Started Failover-Check-p12-345-678-q33.test.hp.com
```



```
Started Filesystem Started Ping-p12-345-678-q33.test.hp.com
Started Ping-p12-345-678-q33.test.hp.com
Started STONITH-SSH-p12-345-678-q90.test.hp.com
Started STONITH-SSH-p12-345-678-q33.test.hp.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 system is OK. In the case above, FAIL indicates that the HA system is not OK. In this case, the secondary disk is out-of-date (primary status/secondary status). FAIL appears if one or more of the following cases apply:

- The heartbeat function is down.
- One of the servers is not online.
- The disk communication state is other than Connected.
- One or more of the heartbeat 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 "service heartbeat start" to restart it.
Disk: Unconfigured
```

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

### Server Status

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

```
p12-345-678-q90.test.hp.com: online
p12-345-678-q33.test.hp.com: online Primary Preferred
```

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



If the secondary was offline or it's heartbeat function stopped, and there was no preferred primary, these lines would look like this:

```
p12-345-678-q90.test.hp.com: offline
p12-345-678-q33.test.hp.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.

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-p12-345-678-q90.test.hp.com
OK Network-p12-345-678-q33.test.hp.com
```

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

### Resource Status

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

- **ESM** is the ESM instance on the primary (ESM services). The Started status begins when the startup process begins. ESM takes several minutes to complete the startup process and become accessible. During this interval, ESM is not available, even though the status is Started. Wait a few minutes and try again.
- **Failover-Check-<hostname>** is a program that checks if a failover is needed. An instance of it runs on each machine. (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 33](#) action. This action will restart the resource.

## tuneDiskSync

The tuneDiskSync 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.

If you are increasing the interconnect cable speed to over 10 GB/sec, and you have not already done so, you will need to increase the maximum TCP buffer size. Follow the instructions given in ["Primary and Secondary Changes" on page 11](#).

## Log Output

HA produces log output of two types, syslogs and HA logs.

**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. HA 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`. Where that is logged 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

There are five procedures for these changes:

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

["Changing the Secondary Hostname or IP Address Only" on page 43](#)

["Changing the Primary Hostname or IP Address Only" on page 43](#)

["Changing Both Server Hostnames or IPs" on page 44](#)

["Changing the Interconnect IP Address" on page 44](#)

## 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."

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. Stop ArcSight Web by running (as user *arcsight*):

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

You may get error messages from this command indicating *arcsight\_web* was not stopped. This is normal and may be ignored.

4. 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.

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

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

6. 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".

7. While logged in as user *arcsight*, run the following to start the setup program for ArcSight Web from the `/opt/arcsight/web/bin` directory:

```
./arcsight websetup
```

This will open the ArcSight Web setup wizard.

- a. Enter the new hostname (that you set earlier in this procedure) in every field in which the old hostname occurs.
  - b. When the certificate from the manager is displayed, check the option "Trust the certification from the manager."
  - c. Select the self-signed keypair option when prompted and enter the required information to generate the self-signed certificate with the new hostname .

8. Start ArcSight Web by running (as user *arcsight*):

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

9. Wait two minute to ensure that `arcsight_web` has started.

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

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

Where `<hostname>` is the new hostname (Hostnames with underscores do not work on IE, so use the IP address.)

11. 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.
12. 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*:

```
service heartbeat stop
```

2. Change the hostname and/or IP address of the secondary as required.
3. On the primary, as user *arcsight*, run:

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

In the First Boot Wizard, specify the new hostname or IP address for the secondary.

When the FBW 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 as user *root*:

```
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. On the primary, as user *arcsight*, run:

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

5. In the First Boot Wizard, specify the new hostname or IP address for the secondary.

When the FBW completes the heartbeat restarts.

You can stop here. You may want to use `arcsight_cluster prefer <secondary hostname>` to fail back to the original server, but it is not necessary. We recommend that you have no preferred server in



order to reduce failovers. In that case, run `arcsight_cluster prefer` with no arguments. This leaves ESM running on the new primary without an additional failover interruption.

## Changing Both Server Hostnames or IPs

1. Run the following commands on the secondary as user *root*:

```
service heartbeat stop
```

2. Change the hostname and/or IP address of the secondary as required.
3. Run the following command on the primary as user *root*:

```
service heartbeat stop
```

There is no failover, because the HA heartbeat is offline on both machines.

4. On the same machine as the previous step, change the hostname and/or IP address as required.
5. 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 FBW completes the heartbeat restarts.

You can stop here. You may want to use `arcsight_cluster prefer <secondary hostname>` to fail back to the original server, but it is not necessary. We recommend that you have no preferred server in order to reduce failovers. In that case, run `arcsight_cluster prefer` with no arguments. This leaves ESM running on the new primary without an additional failover interruption.

## Changing the Interconnect IP Address

Use the following procedure to change the IP address of the interconnect cable:

1. Change to the `/etc/sysconfig/network-scripts` directory.
2. Select and edit the file for the network interface that you want to change by changing the `IPADDR` value. For example the file might be `ifcfg-eth1`.
3. Run the `ifdown` and `ifup` commands (for example, `ifdown eth1; ifup eth1`).
4. Run the First Boot Wizard on the primary 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 if, for example, it has hardware problems. 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 ["Preparing for Installation" on page 7](#). 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, that are described in the "Preparing for Installation" topics.
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.
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.

## Changing Mount Options

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

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

Where:

- <file system type> must be ext4 or xfs, and *cannot be changed*.
- <new mount options> are the new options you want.
- <shared disk> is where the shared disk is mounted, which *cannot be changed* (typically /opt or /opt/arcsight).
- /dev/drbd1 is the name of the mirrored volume.

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

```
arcsight_cluster tuneDiskSync
```



# Chapter 5: Example HA Implementation

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

- ["Requirements"](#) outlines requirements for the cluster – what it must be able to do.
- ["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.

## Requirements

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

- 2TiB of RAID 10 storage is provided via 15K RPM disks.
- The network interface runs at 1G.
- Two 1G interfaces are bonded into a single 2G interconnect connection.
- RedHat 6.5 is used with ESM 6.8c 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 recommended. It will be seen how much easier it is to increase the disk space later using LVM.



## 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 1G 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, and eth2 (port 3) is similarly interconnected. 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
<b>Primary</b>	ha1.internal.<yourcompany>.com	10.10.10.2
<b>Secondary</b>	ha2.internal.<yourcompany>.com	10.10.10.3
<b>Service</b>	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/vfstab 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 2.2TiB (it will of course be a little smaller than that), the metadata partition must be at least  $(2.2\text{TiB}/32778) + 1 = 72$  MiB.

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/vfstab and mount /opt with the command mount /opt. The vfstab 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 set up a bonded interface between two interfaces. For this example eth1 and eth2 will be bonded into the bond0 interface. Pick some IP addresses for the interconnect



interfaces – a private subnet that is not routed to other nodes is a good choice. In this case we will use subnet 192.168.10.0/24 – 192.168.10.2 will be the primary cable IP, and 192.168.10.3 will be the secondary IP.

To set this up, first modify the interface scripts `ifcfg-eth1` and `ifcfg-eth2`, and create a new script `ifcfg-bond0`. All of these files are in `/etc/sysconfig/network-scripts`. You make these changes on both machines.

An example of an `ifcfg-eth*` script after changes is shown below:

```
DEVICE=eth2
HWADDR=12:34:56:78:90:AB
UUID=3835e99d-2ef2-422b-9455-75697e092689
TYPE=Ethernet
ONBOOT=yes
BOOTPROTO=none
MASTER=bond0
SLAVE=yes
USERCTL=no
NM_CONTROLLED=no
```

The first three lines come from the original file that was created when the Operating System was installed. Delete any other lines from the original file. The remaining lines are the same for all such files – you may copy them in.

The `ifcfg-bond0` file for the primary contains the following:

```
DEVICE=bond0
IPADDR=192.168.10.2
NETMASK=255.255.255.0
ONBOOT=yes
BOOTPROTO=none
USERCTL=no
NM_CONTROLLED=no
BONDING_OPTS="miimon=100 mode=0"
```

This file should contain the secondary cable IP (192.168.10.3) on the secondary. Otherwise the two files should be the same.

To bring up the connection, run `ifup bond0` 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

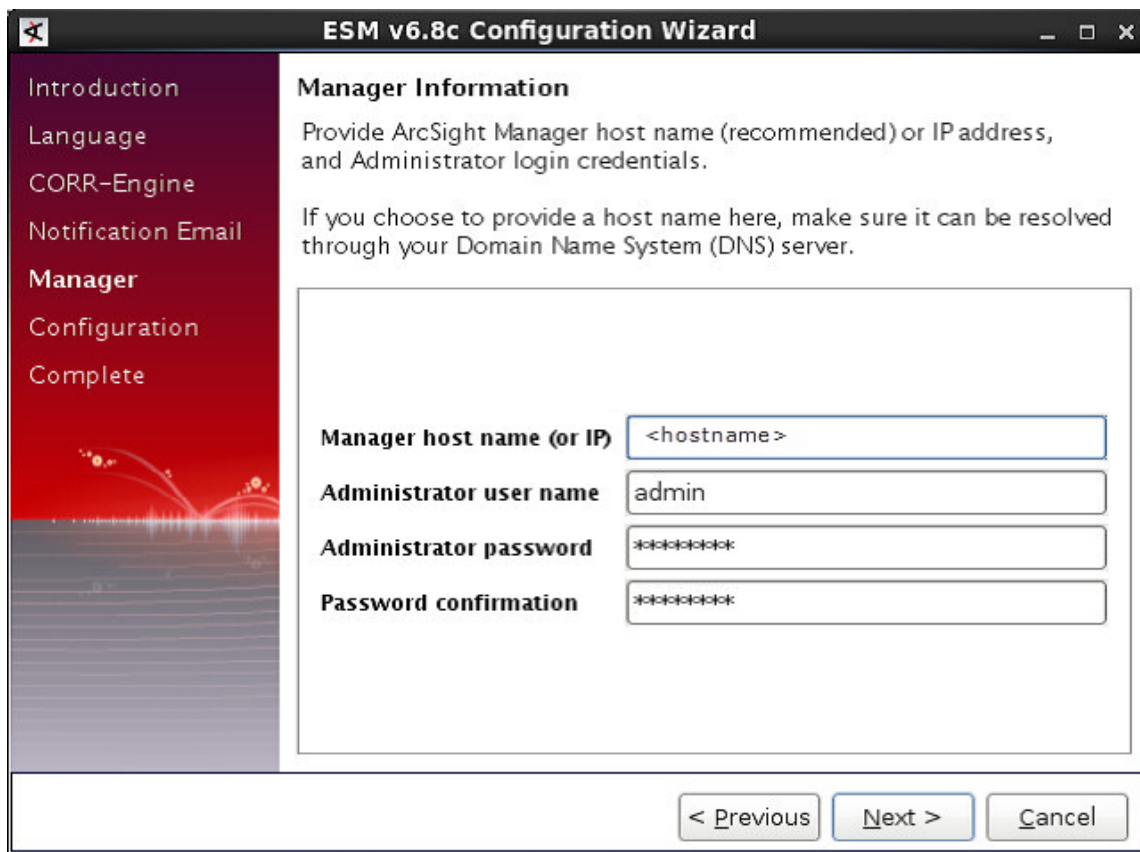
HA is installed on ha1.internal.acme.com . Here are the parameters to use to install HA:

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

### Install ESM

ESM is installed as described in the ESM Installation Guide. The only special step is at the window shown below:





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

## 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. 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 = 3.3TiB of usable disk space. 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  $(5.5\text{TiB}/32778) + 1 = 177$  MiB.

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

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



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

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

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

```
./arcsight_cluster increaseDisk
```

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

```
xfs_growfs /dev/drbd1
```

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

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



## Chapter 6: Troubleshooting

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 .....	53
General Problems .....	56
Audit Events .....	57
Failover Triggers .....	58
Processes Killed During Failover .....	59
System does not Failover .....	59
Network Interface Commands Stall Disk Mirroring .....	59
No ESM Uninstall Links on the Primary .....	60
Neither Server will Come Up as Primary .....	60
Stopping the Network on the Secondary Kills ESM .....	60

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



Installation Message	Description
Incorrect root password for <hostname> - please enter correct one.	You entered an incorrect password. Enter the correct one.
Failed to set up SSH access. See <log file> for details.	SSH access didn't work. See the specified log for suggestions.
No arcsight user on secondary. Please create one identical to that on primary	Create a user <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.
<b>Crossover Cable Issues</b>	
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.
<b>Shared Disk Issues</b>	
Unmount of <shared_disk> failed. Fix the problem, and re-run this script.	Fix the problem, uninstall HA, and reinstall HA. Do not re-run the script.



Installation Message	Description
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 get the next error as well - so follow the instructions there.
<metadata_vol> appears to be in use. See the following output from blkid <metadata_vol> --- blkid output here ---  If this volume is not in use, run dd if=/dev/zero of=<metadata_vol> to clear this volume and then re-run the installation.	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-install.
Disk status must be Connected to reconfigure cluster.	The HA module is already installed on both machines, so this call to the FBW must be to reconfigure the installation. This can only be done if the disk status is Connected (normal).  Run arcsight_cluster diagnose and then try re-running the installation.
Please mount <shared_disk partition>, and re-run installation.	Mount the shared disk.
Size of metadata volume <metadata_vol> is less than required minimum of <megabytes>M	The metadata volume is too small to support shared_disk. Increase the size of the metadata volume.
the size of <volume> on the secondary is <megabytes>M. It must be the same as the primary - <megabytes>M.	This could refer either to the shared disk volume or the metadata volume. The size of each must be the same on each server (rounded to the nearest Mbyte). Change the sizes to make them match.
<volume> not a valid disk volume.	Either the shared disk or the metadata volume is not really a volume. Check to see if there is a typographical error in the name you specified.



Installation Message	Description
Found <megabytes>M disk space used on <shared_disk>. The installation would delete these files.	The installation found more than 10M of files on <shared_disk> on the secondary. Most likely, you are trying to do an upgrade install on the secondary. This is hard to check, but it would result in the deletion of the ESM installation on the other server! The installation is terminated. You can either run this on the primary server or delete the files.
<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.
<b>Primary/Secondary Host Issues</b>	
No interface found for <primary_ip> on Primary	The primary IP/hostname must be the first IP on an interface. Configure the primary hostname to correspond to an interface.
No interface found for <secondary_ip> on Secondary	The secondary IP/hostname must be the first IP on an interface. Configure the secondary hostname to correspond to an interface.
Primary IP <primary CIDR> and Secondary IP <secondary IP> must be in the same subnet.	Change host IP addresses so they are in the same subnet.
<hostname> - the hostname of this host does not correspond to the hostname given for either the Primary or the Secondary.	Correct the incorrect hostname.
<hostname> is not a valid hostname or IP address.	Correct the incorrect hostname.
IP for <host> is <found IP> on this server and <other IP> on the other.	The hostname resolves to different IP addresses on the different servers. Make the server configurations consistent. A likely cause is inconsistencies in /etc/hosts.
OS version on primary and secondary are different.	Make them the same.

## General Problems

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

```
arcsight_cluster diagnose
```

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



## Audit Events

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

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

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

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

### highavailability:100

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

Severity: 3

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

### highavailability:200

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

- Failure of either network interface card (NIC)
- Cross-over cable failure or disconnect
- Secondary system failure or shutdown



- Secondary system hard drive failure.
- You reboot the secondary system for any reason

Severity: 7

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

## highavailability:300

This event occurs when the Distributed Replicated Block Device (DRBD) storage system begins the process of synchronizing the primary and secondary hard drives and continues every five minutes (by default) until the synchronization is complete. Each event includes the amount of data between the two systems that has been synchronized as a percentage until it reaches 100 percent. You can change the interval using the `highavailability.notification.interval` property as described in ["Configurable Properties" on page 28](#).

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, 300, or 400) to OK. It may take 30 seconds for this event to generate after the secondary system and high-availability service is restored.

Severity: 3

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

## Failover Triggers

The following occurrences can trigger a failover:

- You select the secondary as the preferred system using the `arcsight_cluster` command. This is the preferred way of forcing a failover.
- You put the primary in offline mode using the `arcsight_cluster` command.
- The primary operating system goes down. In the case of a routine system restart, the machine doing the restart may continue to be primary. This is true when the system starts again before the failover had time to trigger.
- The hard disk on the primary system fails.
- Loss of an internet connection to the primary system. (it may take several minutes.)

Although you might think they would, 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 restart the Manager, it does not trigger a failover.
- If the network switch fails causing a communications failure to both primary and secondary systems, there is no failover. Users would immediately detect that their ArcSight Console or ArcSight Command Center UIs have lost communication with the Manager. The primary continues to run and connectors cache events until communications are restored, at which time the primary ESM continues as usual.
- If the primary system runs out of disk space, the secondary also runs out of space because of the mirroring. No failover is triggered.

## Processes Killed During Failover

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

## System does not Failover

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 45](#), adding the `inode64`.

Your mount command might look something like this:

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

## Network Interface Commands Stall Disk Mirroring

If you use network interface commands such as:

- `ifdown <interface>` followed by `ifup <interface>`,
- `ifconfig <interface> down` followed by `ifconfig <interface> up`, or
- `ip set <interface> down`, followed by `ip set <interface> up`

... the disk mirroring component does not recover automatically.



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

## No ESM Uninstall Links on the Primary

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

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

## Neither Server will Come Up as Primary

There could be a scenario where, machine A, the primary, is healthy and machine B, the secondary has been taken off line, perhaps for maintenance. The disks cannot be mirrored in that situation, and the plan is to allow time for the synchronization to catch up after machine B is brought back on line.

However, if machine A, the Primary, were to crash before the disks are synchronized, machine B cannot take over because the disks are not in sync.

The solution is to run the `arcsight_cluster forcePrimary` on machine B to force it to become the primary, even though some data might be lost.

## Stopping the Network on the Secondary Kills ESM

If you run the command `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 `service network start` by itself does not restore ESM.

To recover from this situation, run `service network start`, if you haven't already. Then run `arcsight_cluster diagnose` on both machines. This command repairs the condition and restarts ESM on the original primary.

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

Avoid using `service network stop` on either machine.



## Send Documentation Feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

**Feedback on ESM High Availability Module User's Guide (ESM High Availability Module 1.0)**

Just add your feedback to the email and click send.

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

We appreciate your feedback!