



High Availability Installation Guide for Cisco Security Manager 4.27

First Published: 2023-09-07

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883



CONTENTS

CHAPTER 1

Preface 1

- Audience 1
- Conventions 3
- Related Documentation 4

CHAPTER 2

Overview 7

- Local Redundancy (HA) Process Overview 8
 - Local Redundancy (HA) Configuration Steps 8
- Geographic Redundancy (DR) Process Overview 9
 - Geographic Redundancy (DR) Configuration Steps 10
- Veritas Products 11

CHAPTER 3

System Requirements 13

- Hardware Requirements for a Single-Node Site 13
- Hardware Requirements for a Dual-Node Site 14
- Software Requirements for a Local Redundancy Configuration 15
- Software Requirements for a Geographic Redundancy (DR) Configuration 16
- Software Requirements for Replication without Clustering 17
- Preinstallation Worksheets 18
 - Local Redundancy Configuration Worksheet 18
 - Geographic Redundancy (DR) Configuration Worksheet 19

CHAPTER 4

Installing the Cisco Security Management Suite High Availability Solution 21

- Making Ethernet Connections 21
- Installing Microsoft Windows Server 22
- Connecting the Servers to External Storage 22

- Installing Veritas Products 23
- Mirroring the Boot Disk (Optional) 23
- Veritas Volume Manager Configuration Tasks 24
 - Primary Server (without Replication) 24
 - Primary Servers (with Replication) 25
 - Secondary Servers and the Primary Server in a Secondary Cluster 25
- Installing Security Manager 26
 - Installing Security Manager on the Primary Server 27
 - Installing Security Manager on Secondary Servers 29
 - Manually Starting Services in Secondary HA Server 32
- Veritas Volume Replicator Tasks 32
- Updating Permissions on the Working Volume 35
 - Updating Permissions when using Shared Storage 35
 - Updating Permissions when using Replication 36
- Veritas Cluster Server Tasks 37
 - Single Local Cluster (Dual-Node) Configuration 37
 - Creating the Cluster 37
 - Creating the Application Service Group 38
 - Creating the ClusterService Group 40
 - Dual Geographic Cluster Configuration 41
 - Creating the Primary and Secondary Clusters 41
 - Creating the ClusterService Group 42
 - Creating the Replication Service Group 43
 - Creating the Application Service Group 45
 - Creating the Cluster Level Configuration 52

CHAPTER 5

- Maintenance Activities 55**
 - Customizing VCS Behavior 55
 - Security Certificates for SSL 56
 - Manually Starting, Stopping, or Failing Over Security Manager 57
 - VCS Case 57
 - Non-VCS Case 58
 - Integrating Cisco Secure ACS with Security Manager 59
 - Upgrading Security Manager 60

Backing Up Security Manager	60
Uninstalling Security Manager	61
Migrating a Non-HA Security Manager to HA	61

CHAPTER 6	High Availability and Disaster Recovery in Virtual Machines	63
	Host-based Failover (Local HA)	63
	Prerequisites for Creating VMware HA Clusters	64
	Configuring Security Manager for Host-based Failover	64
	Limitations	64
	Fault Tolerance	65
	Creating Fault Tolerant Systems	65
	Disaster Recovery	68
	System Requirements	68
	Hardware Requirements	68
	Software Requirements	68
	Configuring VMware Site Recovery Manager	68
	Configuring vCenter	71
	Configuring the Recovery Site	73
	Configuring Replication	74
	Installing Security Manager in Disaster Recovery Environment	74

APPENDIX A	VCS Resource Views for the Reference Configurations	75
	VCS Resource Views for the Reference Configurations	75
	Single Local Cluster (Dual-Node) Configuration	75
	Dual Geographic Cluster (Single-Node) Configuration	76

APPENDIX B	High Availability and Disaster Recovery Certification Test Plan	81
	Manual Switches	81
	IntraCluster Switch	81
	InterCluster Switch	82
	Ethernet/Network Failures	83
	Network Communication Failures	83
	Network Ethernet Failure on Secondary Server, Single Cluster	83
	Network Ethernet Failure on Primary Server, Single Cluster	84

Network Ethernet Failure on Secondary Server, Dual Cluster	85
Network Ethernet Failure on Primary Server, Dual Cluster	87
Cluster Communication Failure	88
Server Failures	90
Standby Server Failure, Single Cluster	90
Primary Server Failure, Single Cluster	91
Standby Server Failure, Dual Cluster	92
Primary Server Failure, Dual Cluster	94
Application Failures	96
Application Failure, Single Cluster	96
Application Failure, Dual Cluster	98



CHAPTER 1

Preface

This document explains how to install Cisco Security Management Suite (Security Manager) in a high availability (HA) or disaster recovery (DR) environment. The Security Manager HA/DR solution is based on Veritas Storage Foundation and High Availability solutions.

For the steps to install Security Manager in a VMware based high availability (HA) or disaster recovery (DR) environment, see [High Availability and Disaster Recovery in Virtual Machines](#), on page 63.

- [Audience](#), on page 1
- [Conventions](#), on page 3
- [Related Documentation](#), on page 4

Audience

The primary audience for this guide is system administrators who are responsible for installing and managing the HA/DR solutions. This guide assumes that you are familiar with the topics in the [Table](#).

Table 1: Topics in this Guide

Configuration	Topics
Local Redundancy	<ul style="list-style-type: none"> • Cisco Security Management Suite • Microsoft Windows Server 2019, Standard and Datacenter Edition, Microsoft Windows Server 2012, Standard and Datacenter Edition, or Microsoft Windows Server 2012 R2, Standard and Datacenter Edition or Microsoft Windows Sever 2016, Standard and Datacenter Edition <p>Note Tested in Microsoft Windows server 2019 and 2016 platforms for CSM 4.27.</p> <ul style="list-style-type: none"> • Veritas Storage Foundation HA for Windows 6.0.1, 6.0.2, 6.1, Veritas InfoScale 7.0, Veritas InfoScale 7.2, Veritas InfoScale 7.4. <ul style="list-style-type: none"> • Veritas InfoScale 7.4.2 and above supports Windows Server 2019 • Windows 2019: Veritas Storage Foundation for Windows Version: 7.4.2 • Windows 2016: Veritas Storage Foundation for Windows Version: 7.4
Geographic Redundancy	<ul style="list-style-type: none"> • Cisco Security Management Suite • Microsoft Windows Server 2019, Standard and Datacenter Edition, Microsoft Windows Server 2012, Standard and Datacenter Edition, or Microsoft Windows Server 2012 R2, Standard and Datacenter Edition or Microsoft Windows Sever 2016, Standard and Datacenter Edition <p>Note Tested in Microsoft Windows server 2019 and 2016 platforms for CSM 4.27.</p> <ul style="list-style-type: none"> • Veritas Storage Foundation HA for Windows 6.0.1, 6.0.2, 6.1, Veritas InfoScale 7.0, Veritas InfoScale 7.2, or Veritas InfoScale 7.4. <ul style="list-style-type: none"> • Veritas InfoScale 7.4.2 and above supports Windows Server 2019 • Windows 2019: Veritas Storage Foundation for Windows Version: 7.4.2 • Windows 2016: Veritas Storage Foundation for Windows Version: 7.4 • Veritas Volume Replicator Option

Configuration	Topics
Geographic Redundancy without Clustering	<ul style="list-style-type: none"> • Cisco Security Management Suite • Microsoft Windows Server 2019, Standard and Datacenter Edition, Microsoft Windows Server 2012, Standard and Datacenter Edition, or Microsoft Windows Server 2012 R2, Standard and Datacenter Edition or Microsoft Windows Server 2016, Standard and Datacenter Edition <p>Note Tested in Microsoft Windows server 2019 and 2016 platforms for CSM 4.27.</p> <ul style="list-style-type: none"> • Veritas Storage Foundation HA for Windows 6.0.1, 6.0.2, 6.1, Veritas InfoScale 7.0, Veritas InfoScale 7.2, or Veritas Infoscale 7.4. <ul style="list-style-type: none"> • Veritas InfoScale 7.4.2 and above supports Windows Server 2019 • Windows 2019: <ul style="list-style-type: none"> Veritas Storage Foundation for Windows Version: 7.4.2 • Windows 2016: <ul style="list-style-type: none"> Veritas Storage Foundation for Windows Version: 7.4 • Veritas Volume Replicator Option

Because the Security Manager HA/DR solution utilizes Veritas Storage Foundation and High Availability Solutions for Windows, we highly recommend the following courses for a local redundancy solution:

- Veritas Storage Foundation for Windows
- Veritas Cluster Server for Windows

For a geographically redundant solution, the following additional courses are highly recommended:

- Veritas Volume Replicator for Windows
- Disaster Recovery Using Veritas Volume Replicator and Global Cluster Option for Windows

Conventions

This document uses the following conventions:

Item	Convention
Commands, keywords, special terminology, and options that should be selected during procedures	boldface font
Variables for which you supply values and new or important terminology	<i>italic font</i>
Displayed session and system information, paths and file names	screen font

Item	Convention
Information you enter	boldface screen font
Variables you enter	<i>italic screen</i> font
Menu items and button names	boldface font
Indicates menu items to select, in the order you select them.	Option > Network Preferences



Tip Identifies information to help you get the most benefit from your product.



Note Means *reader take note*. Notes identify important information that you should reflect upon before continuing, contain helpful suggestions, or provide references to materials not contained in the document.



Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage, loss of data, or a potential breach in your network security.



Warning Identifies information that you must heed to prevent damaging yourself, the state of software, or equipment. Warnings identify definite security breaches that will result if the information presented is not followed carefully.

Related Documentation

Refer to the following Cisco publications for additional information. These documents are available at <http://www.cisco.com/c/en/us/support/security/security-manager/tsd-products-support-series-home.html>.

- *Installation Guide for Cisco Security Manager 4.27*
- *User Guide for Cisco Security Manager 4.27*
- *Release Notes for Cisco Security Manager 4.27*

Refer to the following publications for additional information concerning Veritas Storage Foundation:

- Veritas Storage Foundation™ *and High Availability Solutions Getting Started Guide*
- Veritas Storage Foundation™ *and High Availability Solutions Release Notes*
- Veritas Storage Foundation™ *and High Availability Solutions Installation and Upgrade Guide*
- Veritas Storage Foundation™ *Administrator's Guide*
- Veritas™ *Cluster Server Release Notes*

- Veritas™ *Cluster Server Installation and Upgrade Guide*
- Veritas™ *Cluster Server Bundled Agents Reference Guide*
- Veritas™ *Cluster Server Administrator's Guide*
- Veritas™ *Volume Replicator Administrator's Guide*
- Veritas™ *Volume Replicator Advisor User's Guide*
- *Hardware Compatibility List (HCL) for Veritas Storage Foundation™ and High Availability Solutions for Windows*
- *Software Compatibility List (SCL) for Veritas Storage Foundation™ and High Availability Solutions for Windows*

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPTER 2

Overview

This document explains how to install Cisco Security Management Suite (Security Manager) in a high availability (HA) or disaster recovery (DR) environment. The Security Manager HA/DR solution is based on Veritas Storage Foundation and High Availability solutions. The Security Manager HA/DR solutions described in this guide support the following applications:

- Security Manager 4.27

The HA solution supports both local redundancy (HA) and geographic redundancy (DR) configurations.



Note Cross-launching the Cisco Prime Security Manager (PRSM) application is supported in both HA and DR configurations; however, seamless, direct access to PRSM from Security Manager using the “single sign-on” (SSO) feature is only supported in HA mode.



Note From version 4.21 onwards, Cisco Security Manager terminates whole support, including support for any bug fixes or enhancements, for all Aggregation Service Routers, Integrated Service Routers, Embedded Service Routers, and any device operating on Cisco IOS software, including the following devices:

- Cisco Catalyst 6500 and 7600 Series Firewall Services Modules (EOL8184)
 - Cisco Catalyst 6500 Series Intrusion Detection System Services Module 2 (EOL8843)
 - Cisco Intrusion Prevention System: IPS 4200, 4300, and 4500 Series Sensors (EOL9916)
 - Cisco SR 500 Series Secure Routers (EOL7687 , EOL7657)
 - PIX Firewalls (EOL)
-

This chapter covers the following sections:

- [Local Redundancy \(HA\) Process Overview, on page 8](#)
- [Geographic Redundancy \(DR\) Process Overview, on page 9](#)
- [Veritas Products, on page 11](#)

Local Redundancy (HA) Process Overview


The local redundancy configuration provides an automatic failover solution in the event of software or hardware failures without the need to reconfigure IP addresses or DNS entries on your switched/routed network.

The [figure](#) illustrates the local redundancy HA configuration.



Note The servers in the [figure](#) optionally contain mirrored internal boot disks. We recommend that they be the same make, model, and storage capacity. We recommend a fault-tolerant switched/routed network for communicating with the HA servers.

Figure 1: Local Redundancy HA Configuration



Local Redundancy (HA) Configuration Steps

The following table lists the steps required to configure a locally redundant installation of Cisco Security Manager.

Procedure

	Command or Action	Purpose
Step 1	Make physical connections.	Making Ethernet Connections, on page 21
Step 2	Install Microsoft Windows Server and all necessary drivers.	Installing Microsoft Windows Server, on page 22
Step 3	Make storage connections.	Connecting the Servers to External Storage, on page 22
Step 4	Install and configure the Veritas products and components.	Installing Veritas Products, on page 23
Step 5	Mirror the boot disk.	Mirroring the Boot Disk (Optional), on page 23
Step 6	Setup required volumes on the shared array.	Veritas Volume Manager Configuration Tasks, on page 24
Step 7	Install Cisco Security Manager on the shared volume on the primary server.	Installing Security Manager, on page 26
Step 8	Install Cisco Security Manager on the spare (dummy) volume on the secondary server.	Installing Security Manager, on page 26
Step 9	Update permissions on secondary server.	Updating Permissions on the Working Volume, on page 35
Step 10	Create and configure clusters.	Veritas Cluster Server Tasks, on page 37

Geographic Redundancy (DR) Process Overview


The geographic redundancy configuration provides disaster recovery by replicating application data between two sites. Failover between sites can be initiated manually or performed automatically.

The [figure](#) illustrates a geographic redundancy (DR) configuration.



Note The servers in the [figure](#) optionally contain mirrored internal boot disks. We recommend that they be the same make, model, and storage capacity. We recommend a fault-tolerant switched/routed network for communicating with the servers.

Figure 2: Geographic Redundancy (DR) Configuration



Geographic Redundancy (DR) Configuration Steps

The following table lists the steps required to configure a geographically redundant installation of Cisco Security Manager.

Procedure

	Command or Action	Purpose
Step 1	Make physical connections.	Making Ethernet Connections, on page 21
Step 2	Install Microsoft Windows Server and all necessary drivers.	Installing Microsoft Windows Server, on page 22
Step 3	Make storage connections.	Connecting the Servers to External Storage, on page 22
Step 4	Install and configure the Veritas products and components.	Installing Veritas Products, on page 23
Step 5	Mirror the boot disk.	Mirroring the Boot Disk (Optional), on page 23
Step 6	Setup required volumes on the shared array.	Veritas Volume Manager Configuration Tasks, on page 24
Step 7	Install Cisco Security Manager on the shared volume on the primary server.	Installing Security Manager, on page 26

	Command or Action	Purpose
Step 8	Install Cisco Security Manager on the spare (dummy) volume on the secondary server.	Installing Security Manager, on page 26
Step 9	Configure replication.	Veritas Volume Replicator Tasks, on page 32
Step 10	Update permissions on secondary server.	Updating Permissions on the Working Volume, on page 35
Step 11	Create and configure clusters.	Veritas Cluster Server Tasks, on page 37

Veritas Products

The Security Manager HA/DR solutions described in this document are based on Veritas products. This section gives a brief summary of each specific Veritas application.

- **Veritas Storage Foundation for Windows (VSWF)**
VSWF provides volume management technology, quick recovery, and fault tolerant capabilities to Windows enterprise computing environments. VSWF provides the foundation for VCS and VVR.
- **Veritas Cluster Server (VCS)**
VCS is a clustering solution for reducing application downtime. The Global Cluster Option (GCO) for VCS supports managing multiple clusters (such as used in a DR configuration).
- **Veritas Volume Replicator (VVR)**
VVR provides a foundation for continuous data replication over IP networks, enabling rapid and reliable recovery of critical applications at remote recovery sites.
- **Veritas Enterprise Administrator (VEA GUI) console**
The VEA GUI console window provides a graphical way to view and manipulate all the storage objects in your system.
- **Cluster Manager (Java Console)**
Cluster Manager (Java Console) offers complete administration capabilities for your cluster. Use the different views in the Java Console to monitor clusters and VCS objects, including service groups, systems, resources, and resource types:
 - **Cluster Monitor**
Cluster Monitor displays general information about actual or simulated clusters. Use Cluster Monitor to log on to and off of a cluster, view summary information on various VCS objects, customize the display, use VCS Simulator, and exit Cluster Manager.
 - **Cluster Explorer**
Cluster Explorer is the main window for cluster administration. From this window, you can view the status of VCS objects and perform various operations.



CHAPTER 3

System Requirements



Note There are numerous configurations possible using different hardware setups. Consult the respective Microsoft and Veritas Hardware Compatibility Lists (HCLs).



Note Although we make every attempt to ensure the availability of third-party hardware and software platforms specified for Security Manager, we reserve the right to change or modify system requirements due to third-party vendor product availability or changes that are beyond our control.

This chapter describes reference configurations for installing Security Manager in an HA or DR environment. This chapter contains the following sections:

- [Hardware Requirements for a Single-Node Site, on page 13](#)
- [Hardware Requirements for a Dual-Node Site, on page 14](#)
- [Software Requirements for a Local Redundancy Configuration, on page 15](#)
- [Software Requirements for a Geographic Redundancy \(DR\) Configuration, on page 16](#)
- [Software Requirements for Replication without Clustering, on page 17](#)
- [Preinstallation Worksheets, on page 18](#)

Hardware Requirements for a Single-Node Site

To install Security Manager in a single-node HA environment, you can configure a fault-tolerant storage array or use internal disks.


The following are the server hardware specifications for a single-node site:

- Server which meets the basic processor and RAM requirements as described in the *Installation Guide for Cisco Security Manager 4.27*
- Minimum of one Ethernet interface (two recommended)
- Minimum of two physical drives (six recommended)

The [Figure 3: Ethernet Connection for a Single-Node Site](#) shows using two Ethernet connections from the server to the switch/router network for redundancy. If an Ethernet port or switch fails, communication to the

server is maintained. If this level of network redundancy is not required, you can use a single connection to the switch/router network (that is, Eth 2 and its associated Ethernet switch are optional).

Figure 3: Ethernet Connection for a Single-Node Site



Hardware Requirements for a Dual-Node Site

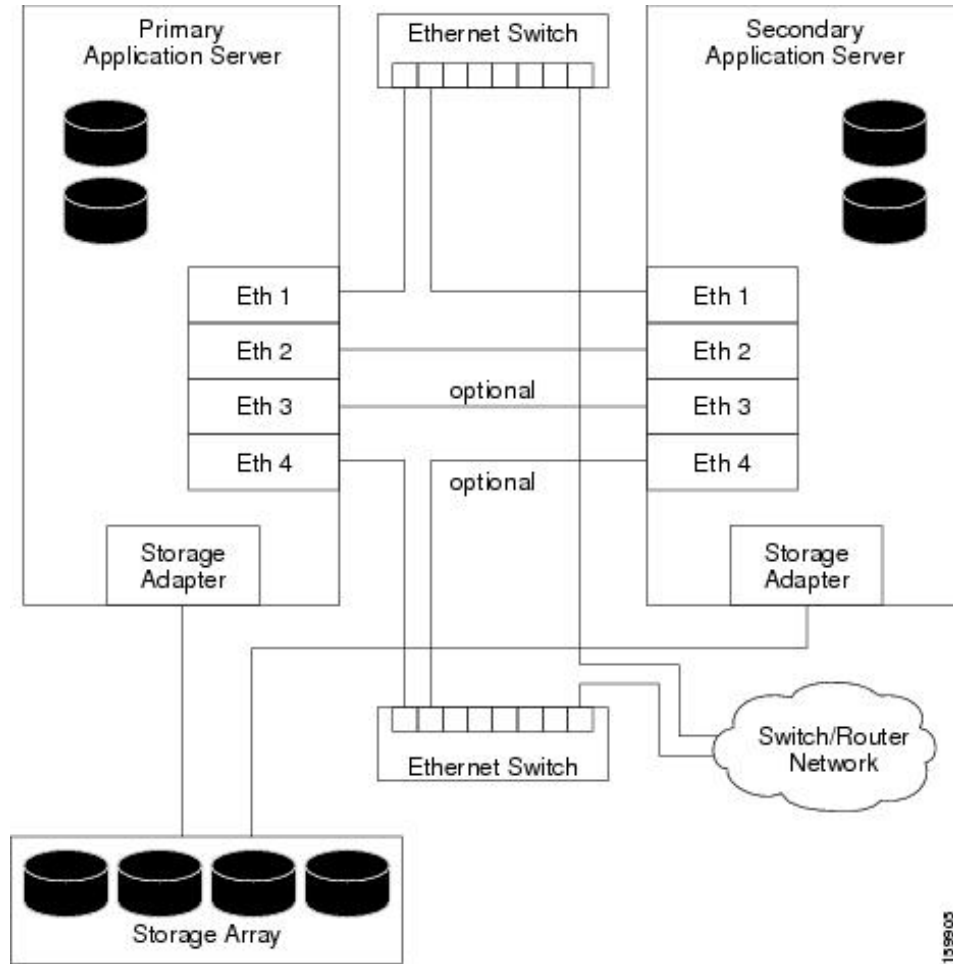
To install Security Manager in a dual-node HA environment, you need two servers that can access a shared storage array.

The following are the server hardware specifications for a dual-node site:

- Servers that meets the basic processor and RAM requirements, as described in the *Installation Guide for Cisco Security Manager 4.26*
- Servers that meets the basic processor and RAM requirements, as described in the *Installation Guide for Cisco Security Manager 4.27*
- Minimum of two Ethernet interfaces (four recommended)
- Minimum of one internal physical drive (two recommended)
- Minimum of one external drive (two recommended; four recommended if using replication)

Figure 4: Ethernet and Storage Connections for a Dual-Node Site depicts the configuration of a dual-node site showing the Ethernet and external storage connections. Two Ethernet connections are used from the server to the switch/router network for redundancy. If an Ethernet port or switch fails, communications to the server is maintained. If this level of network redundancy is not required, you can use a single connection to the switch/router network (that is, Eth 4 and its associated Ethernet switch are optional). Two direct Ethernet connections are made between the servers for cluster heartbeat communications, although second heartbeat connection (Eth 3) is optional.

Figure 4: Ethernet and Storage Connections for a Dual-Node Site



159903

Software Requirements for a Local Redundancy Configuration

The following software is required to install Security Manager in a local redundancy HA configuration:

- Cisco Security Management Suite 4.27
- Microsoft Windows Server 2019, Standard and Datacenter Edition, Microsoft Windows Server 2016, Standard and Datacenter Edition.



Note Starting from version 4.24, Cisco Security Manager supports Microsoft Windows Server 2019



Note Starting from version 4.13, Cisco Security Manager supports Microsoft Windows Server 2016

- Veritas Storage Foundation HA for Windows versions 6.0.1 / 6.0.2 / 6.1/Veritas InfoScale 7.0 / 7.2 / 7.4.
 - Veritas InfoScale 7.4.2 and above supports Windows Server 2019
 - Windows 2019:
 - Veritas Storage Foundation for Windows Version: 7.4.2
 - Windows 2016:
 - Veritas Storage Foundation for Windows Version: 7.4



Note Veritas Infoscale 7.0 does not support Windows Sever 2016. However Veritas Infoscale 7.2 / 7.4 supports Windows Server 2016.



Note Veritas InfoScale 7.4.2 and above supports Windows Server 2019.

- Dynamic Multipathing Option

A Security Manager license is only required for the active server in a HA/DR configuration. Additional licenses for standby servers are not required.

Veritas Storage Foundation HA for Windows is licensed on a per-node basis. In the same local redundancy configuration example, each server needs to have its own license for running Veritas Storage Foundation HA for Windows.

The Veritas Dynamic Multipathing Option is required only if you plan to use external storage with more than one host bus adapter in a server, which provides multiple paths between the server and storage.



Note Beginning with version 4.20, Cisco Security Manager supports Veritas Infoscale 7.4.

Software Requirements for a Geographic Redundancy (DR) Configuration

The following software is required to install Security Manager in a geographic redundancy (DR) configuration:

- Cisco Security Management Suite 4.27
- Microsoft Windows Server 2019, Standard and Datacenter Edition, Microsoft Windows Server 2016, Standard and Datacenter Edition
- Veritas Storage Foundation HA/DR for Windows 6.0.1 / 6.0.2 / 6.1/Veritas InfoScale 7.0 / 7.2 / 7.4
 - Veritas InfoScale 7.4.2 and above supports Windows Server 2019
 - Windows 2019:

Veritas Storage Foundation for Windows Version: 7.4.2

- Windows 2016:

Veritas Storage Foundation for Windows Version: 7.4

- Veritas Volume Replicator Option
- Veritas Dynamic Multipathing Option

Security Manager is licensed per active server in an HA/DR configuration. For example, in a geographic redundancy configuration with a single-node cluster at site A and a single-node cluster at Site B, you only need to purchase one copy of Security Manager, since Security Manager is only active on one server at any given time.

Veritas Storage Foundation HA for Windows is licensed on a per-node basis. In the same geographic redundancy configuration example with two servers (one per cluster), each server needs to have its own license for running Veritas Storage Foundation HA for Windows.

The Veritas Volume Replicator Option is licensed on a per-node basis.

The Veritas Dynamic Multipathing Option is required only if you plan to use external storage with more than one host bus adapter in a server, which provides multiple paths between the server and storage.

Software Requirements for Replication without Clustering

The following software is required to install Security Manager in a geographic redundancy (DR) configuration without clustering:

- Cisco Security Management Suite 4.27
- Microsoft Windows Server 2019, Standard and Datacenter Edition, Microsoft Windows Server 2012, Standard and Datacenter Edition
- Veritas Storage Foundation Basic for Windows 6.0.1 / 6.0.2 / 6.1/Veritas InfoScale 7.0 / 7.2 / 7.4
 - Veritas InfoScale 7.4.2 and above supports Windows Server 2019
 - Windows 2019:
Veritas Storage Foundation for Windows Version: 7.4.2
 - Windows 2016:
Veritas Storage Foundation for Windows Version: 7.4
- Veritas Volume Replicator Option
- Veritas Dynamic Multipathing Option

Security Manager is licensed for each active server in a HA/DR configuration. For example, in a geographic redundancy configuration with replication running between a primary server and a secondary server, you need to purchase only one copy of Security Manager, because Security Manager is active on only one server at any given time.

Veritas Storage Foundation for Windows is licensed on a per-node basis. In the same geographic redundancy configuration example with two servers, each server must have its own license for running Veritas Storage Foundation for Windows.

Veritas Storage Foundation Basic for Windows versions 6.0.1 / 6.0.2 / 6.1 / 7.4/7.4.2 Veritas InfoScale 7.0 / 7.2 / 7.4/7.4.2 work with up to four volumes and are available for free download.

The Veritas Volume Replicator Option is licensed on a per-node basis.

The Veritas Dynamic Multipathing Option is required only if you plan on using external storage with more than one host bus adapter in a server, which provides multiple paths between the server and storage.

Preinstallation Worksheets

Use the preinstallation worksheet to plan your installation and to gather the information you will need during configuration. This section contains the following topics:

- [Local Redundancy Configuration Worksheet, on page 18](#)
- [Geographic Redundancy \(DR\) Configuration Worksheet, on page 19](#)

Local Redundancy Configuration Worksheet

Before you install Security Manager in a local redundancy HA configuration, write down the information outlined in the [table](#) to assist you in completing the installation.

Table 2: Preinstallation Worksheet for a Local Redundancy Configuration

Information	Primary Site	
Shared Disk Group Name	datadg	
Shared Volume Name	cscopx	
Drive Letter for Security Manager Data		
Shared Disk Group Name for Event Data ¹	datadg_evt	
Shared Volume Name for Event Data ²	cscopx_evt	
Drive Letter for Security Manager Event Data ³		
Cluster Name	CSManager_Primary	
Cluster ID	0 ⁴	
Security Manager Virtual IP Address/Subnet mask		
Cluster Service Virtual IP Address/Subnet mask ⁵		
	Primary Server	Secondary Server
Hostname		

Information	Primary Site	
Public Network Interface #1 and IP Address/Subnet Mask		
Public Network Interface #2 ⁶ and IP Address/Subnet Mask		
Private Cluster Interconnect #1		
Private Cluster Interconnect #2		

- ¹ Optional: Use these fields if you want your event data stored separately.
² Optional: Use these fields if you want your event data stored separately.
³ Optional: Use these fields if you want your event data stored separately.
⁴ Must be an integer between 0 and 255 and unique for clusters in the same subnet.
⁵ This is the same value as the Security Manager Virtual IP Address/Subnet mask.
⁶ Required if a second NIC will be used to access the public network for redundancy.

7

Geographic Redundancy (DR) Configuration Worksheet

If you are installing Security Manager in a geographic redundancy (DR) configuration, write down the information outlined in the [table](#) to assist you in completing the installation.

Table 3: Preinstallation Worksheet for a Geographic Redundancy (DR) Configuration

Information	Primary Site	Secondary Site
Disk Group	datadg	datadg
Data Volume	cscopx	cscopx
Drive Letter for Security Manager		
Disk Group for Event Data ⁸	datadg_evt	datadg_evt
Data Volume for Event Data	cscopx_evt	cscopx_evt
Drive Letter for Event Data		
Storage Replicator Log Volume	data_srl	data_srl
Replicated Data Set	CSM_RDS	
Replicated Volume Group	CSM_RVG	

- ⁷ 1. Optional: Use these fields if you want your event data stored separately.
2. Must be an integer between 0 and 255 and unique for clusters in the same subnet.
3. This is the same value as the Security Manager Virtual IP Address/Subnet Mask.
4. Required if a second NIC will be used to access the public network for redundancy.

Information	Primary Site		Secondary Site	
Cluster Name	CSManager_Primary		CSManager_Secondary	
Cluster ID	0 ⁹		1 ¹⁰	
Security Manager Virtual IP Address/Subnet Mask				
Replication Virtual IP Address/Subnet Mask				
Cluster Service Virtual IP Address/Subnet Mask ^{11 12}				
	Primary Server	Secondary Server	Primary Server	Secondary Server
Hostname				
Public Network Interface #1 and IP Address/Subnet Mask				
Public Network Interface #2 and IP Address/Subnet Mask ¹³				
Private Cluster Interconnect #1 ¹⁴				
Private Cluster Interconnect #2 ¹⁵				

⁸ Optional: Use these fields if you want your event data stored separately.

⁹ Must be an integer between 0 and 255 and unique for clusters in the same subnet.

¹⁰ Must be an integer between 0 and 255 and unique for clusters in the same subnet.

¹¹ Required only for clusters using two servers or multiple adapters to access the public network. For a single server cluster with only one network adapter to access the public network, the fixed IP address of this adapter can be used.

¹² This is the same value as the Security Manager Virtual IP Address/Subnet mask.

¹³ Required if you are using a second NIC to access the public network for redundancy.

¹⁴ Required only for clusters using two servers.

¹⁵ Required only for clusters using two servers.

[16](#)

¹⁶ 1. Optional: Use these fields if you want your event data stored separately.

2. Must be an integer between 0 and 255 and unique for clusters in the same subnet.

3. Required only for clusters using two servers or multiple adapters to access the public network. For a single server cluster with only one network adapter to access the public network, the fixed IP address of this adapter can be used.

4. This is the same value as the Security Manager Virtual IP Address/Subnet Mask.

5. Required if a second NIC will be used to access the public network for redundancy.

6. Required only for clusters using two servers.



CHAPTER 4

Installing the Cisco Security Management Suite High Availability Solution

This chapter explains how to install Security Manager in an HA or DR deployment configuration. You should perform these tasks in order, although some tasks are optional or might not apply depending on your specific configuration.

This chapter contains the following topics:

- [Making Ethernet Connections, on page 21](#)
- [Installing Microsoft Windows Server, on page 22](#)
- [Connecting the Servers to External Storage, on page 22](#)
- [Installing Veritas Products, on page 23](#)
- [Mirroring the Boot Disk \(Optional\), on page 23](#)
- [Veritas Volume Manager Configuration Tasks, on page 24](#)
- [Installing Security Manager, on page 26](#)
- [Veritas Volume Replicator Tasks, on page 32](#)
- [Updating Permissions on the Working Volume, on page 35](#)
- [Veritas Cluster Server Tasks, on page 37](#)

Making Ethernet Connections

To make the Ethernet connections required by your HA or DR configuration, follow these steps:

-
- Step 1** Make the Ethernet connections between the servers and switches according to [Figure 3: Ethernet Connection for a Single-Node Site](#) or [Figure 4: Ethernet and Storage Connections for a Dual-Node Site](#), depending on your cluster configuration.

Note Use of a second Ethernet connection to the router/switch network for each server is optional, but it adds an extra level of redundancy in the event of a NIC or local Ethernet switch failure. Veritas Cluster Server (VCS) includes the IPMultiNicPlus agent. This agent allows setting up multiple NIC cards on a server which provides redundant access for the server to the router/switch network. If a NIC card fails, a cable is removed, or some other failure occurs, VCS can detect the failure and reassign the working virtual IP address to another working NIC card on the server. See the Veritas Cluster Server Bundled Agents Reference Guide for details on the IPMultiNicPlus agent. The examples in this document only show the case of a single NIC card for network access. You can also use vendor specific NIC teaming (IEEE 802.3ad link aggregation) solutions as an alternative.

Step 2 In the case of a dual-node cluster, make the Ethernet cluster communication connections between the servers according to [Figure 4: Ethernet and Storage Connections for a Dual-Node Site](#). When connecting directly between servers, you might not have to use a crossover Ethernet cable, depending on whether the interfaces support automatic crossover detection. Most newer Ethernet interfaces support this feature and allow using a straight-through cable when directly connecting to another server.

Installing Microsoft Windows Server

Install the supported Microsoft Windows operating system:

- Microsoft Windows Server 2019, Standard and Datacenter Edition
- Microsoft Windows Server 2012, Standard and Datacenter Edition
- Microsoft Windows Server 2012 R2, Standard and Datacenter Edition
- Microsoft Windows Server 2016, Standard and Datacenter Edition

We recommend that you use the same operating system on all servers.



Note Veritas Storage Foundation HA for Windows version 6.0.1 / 6.0.2 / 6.1 / 7.4 / 7.4.2 Veritas InfoScale 7.0 / 7.2 / 7.4 / 7.4.2 requires that you install the operating system in the same path on all systems. For example, if you install Windows on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.

Connecting the Servers to External Storage

If you are using a dual-node cluster, shared external storage is required. You may use any storage hardware in Hardware Compatibility List for Veritas Storage Foundation & High Availability Solutions for Windows. Either internal or external storage can be used for a single-node cluster.

Installing Veritas Products

Install and configure the Veritas products and components. The products and components required vary depending on whether a single local cluster, dual geographic clusters, or replication without clustering configuration is used. Some components are optional, such as the GUI for Volume Manager (Veritas Enterprise Administrator). See [Table 4: Veritas Software Components](#).

Table 4: Veritas Software Components

Veritas Product/Component	Single Local Cluster	Dual Geographic Clusters	Replication without Clustering
Storage Foundation for Windows	—	—	Required
Veritas Storage Foundation HA for Windows version 6.0.1 / 6.0.2 / 6.1 / 7.4 / 7.4.2 Veritas InfoScale 7.0 / 7.2 / 7.4 / 7.4.2	Required	Required	—
Volume Replicator Option	Not Required	Required	Required
Global Cluster Option	Not Required	Required	—
Dynamic Multipathing Option	See Note ¹⁷	See Note	See Note
Veritas Enterprise Administrator (GUI) ¹⁸	Required	Required	Required
Cluster Manager (GUI) 2	Optional	Optional	—

¹⁷ Required only if you are using external storage with multiple host bus adapters providing multiple paths between the server and disk storage

¹⁸ Can be installed either on the server or a separate client machine.

See the applicable Veritas release notes and installation guides for prerequisites and instructions for installing the Veritas software.



Note One important prerequisite is that you configure the servers as part of a Windows Server domain.

Mirroring the Boot Disk (Optional)

Mirroring the boot disk is optional; however, it provides an extra level of protection for a given server. If the boot disk fails, the machine can be recovered quickly by booting from the mirrored alternate boot disk. Mirroring is accomplished by placing the boot disk in a dynamic disk group under Veritas Volume Manager control and then adding a mirror.

See the section called “Set up a Dynamic Boot and System Volume” in the Veritas Storage Foundation HA for Windows version 6.0.1 / 6.0.2 / 6.1 / 7.4 / 7.4.2 Veritas InfoScale 7.0 / 7.2 / 7.4 / 7.4.2 administrator’s guide for details on this procedure.

Veritas Volume Manager Configuration Tasks

In this section, you configure the necessary disk group and volumes required for the Security Manager application. The configuration varies depending on whether the server involved is the primary server or a secondary server and whether or not replication is involved. You can perform Volume Manager tasks with the VEA GUI or through the command line. For details on using VEA or the command line for these steps, see the Veritas Storage Foundation HA for Windows version 6.0.1 / 6.0.2 / 6.1 / 7.4 / 7.4.2 Veritas InfoScale 7.0 / 7.2 / 7.4 / 7.4.2 administrator's guide.

This section contains the following topics:

- [Primary Server \(without Replication\), on page 24](#)
- [Primary Servers \(with Replication\), on page 25](#)
- [Secondary Servers and the Primary Server in a Secondary Cluster, on page 25](#)

Primary Server (without Replication)

Use this procedure to configure the disk group and volumes required for Security Manager on the primary server in a single-cluster configuration where replication is not involved. In a single-cluster configuration, external shared storage is used, which is accessible to all servers in the cluster.

To configure the disk group and volumes, follow these steps:

Step 1 Create a disk group with the following characteristics:

- Group Name: **datadg**
- Type: **Dynamic (Cluster)**
- Number of Disks: If using software RAID, include at least two disks in the group for mirroring; otherwise, a single logical disk (using hardware RAID) is sufficient. The disks used for this disk group must be accessible to all nodes in the cluster.

Note The use of software RAID 5 is not recommended.

Step 2 Create a volume in the datadg disk group with the following characteristics:

- Volume Name: **cscoptx**
- Assigned Drive Letter: **<Selected Drive Letter>**

Note You can choose any available drive letter; however, the drive letter must be the same on all systems.

- File Type: **NTFS**
-

Primary Servers (with Replication)

Use this procedure to configure the disk group and volumes required for Security Manager on the primary servers in a dual geographic configuration where replication is running between the two clusters. Perform this procedure on the primary server in both the primary and secondary cluster. For each cluster you can use either a single-node cluster or a cluster with multiple nodes using shared storage; however, this document does not cover the case of a multi-node cluster in a dual geographic configuration.

To configure the disk group and volumes, follow these steps:

-
- Step 1** Create a disk group with the following characteristics:
- Group Name: datadg
 - Type: Dynamic (Cluster) (when using VCS), Dynamic (Secondary) (when not using VCS)
 - Number of Disks: If using software RAID, include at least two disks in the group for mirroring; otherwise, a single logical disk (which uses hardware RAID) is sufficient. If this is a multi-node cluster, the disks used for this disk group must be accessible to all nodes in the cluster.

Note The use of software RAID 5 is not recommended.

- Step 2** Create a volume in the datadg disk group with the following characteristics:
- Volume Name: cscopx
 - Assigned Drive Letter: <Selected Drive Letter> (for the primary cluster), None (for the secondary cluster)
 - File Type: NTFS (for the primary cluster), None (for the secondary cluster)
 - Volume Logging: None

- Step 3** Create a volume in the datadg disk group for use as a storage replicator log (SRL) with the following characteristics:
- Volume Name: data_srl
 - Assigned Drive Letter: None
 - File Type: Unformatted
 - Volume Logging: None

Note For information on choosing the proper size of the SRL, see the Volume Replicator administrator's guide.

Secondary Servers and the Primary Server in a Secondary Cluster

Use this procedure to configure the disk group and volumes required for installing Security Manager on secondary servers and on the primary server in a secondary cluster. You must install Security Manager on all secondary servers, as well as the primary server in a secondary cluster. In these cases, you install Security Manager on a spare volume, which is mounted temporarily before installation, then dismounted and not used again until you want to uninstall Security Manager from the server or you want to upgrade Security Manager. You must mount the temporary volume on the same drive letter as the one used for the primary server in the

primary cluster and you must use the same installation path (for example, F:\Program Files\CSCOPx) during the installation.

To configure the disk group and volumes, follow these steps:

Step 1 If you are not creating the spare volume on an existing disk group, create a disk group with the following characteristics:

- Group Name: **datadg_spare**
- Type: **Dynamic (Secondary)**
- Size: **10GB** (The volume only needs to be large enough to install Security Manager)
- Number of Disks: Since this disk group is not used to store application data, a single, nonredundant disk is sufficient.

Step 2 Create a volume in the disk group with the following characteristics:

- Volume Name: **escopx_spare**
- Assigned Drive Letter: **<Selected Drive Letter>**

Note You **must** use the same drive letter that is used for the escopx drive on the primary server.

- File Type: **NTFS**
-

Installing Security Manager

The Security Manager installer detects the presence of Veritas Storage Foundation HA for Windows version 6.0.1 / 6.0.2 / 6.1 / 7.4 / 7.4.2 Veritas InfoScale 7.0 / 7.2 / 7.4 / 7.4.2 and asks you whether you want to install Security Manager in an HA/DR configuration. If you select this option, the only additional information to specify beyond a regular installation is the database password. In a non-HA/DR installation, the database password is autogenerated. However, since the database password must be the same on all servers in the HA/DR configuration, the installer prompts you to specify the password. You must use this same password on all servers in the HA/DR configuration.

The HA/DR installation installs the Cisco Security Manager agent for VCS, so VCS recognizes a new CSManager resource type and can control and monitor Security Manager.

The HA/DR installation also configures the Security Manager and related services in Windows for a Startup Type of Manual instead of Automatic, because the Veritas cluster server instead controls the starting and stopping of Security Manager on each server in the HA/DR configuration. Otherwise, the Security Manager application would try to start on all servers in the HA/DR configuration after any server reboot, when Security Manager should run only on one server at any time.

You must install Security Manager on each server in the HA/DR configuration. However, only the primary instance of Security Manager is used and protected in the HA/DR configuration. Other installations are performed to enable the primary instance to run on any of the secondary servers in the configuration.

This section contains the following topics:

- [Installing Security Manager on the Primary Server, on page 27](#)

- [Installing Security Manager on Secondary Servers, on page 29](#)

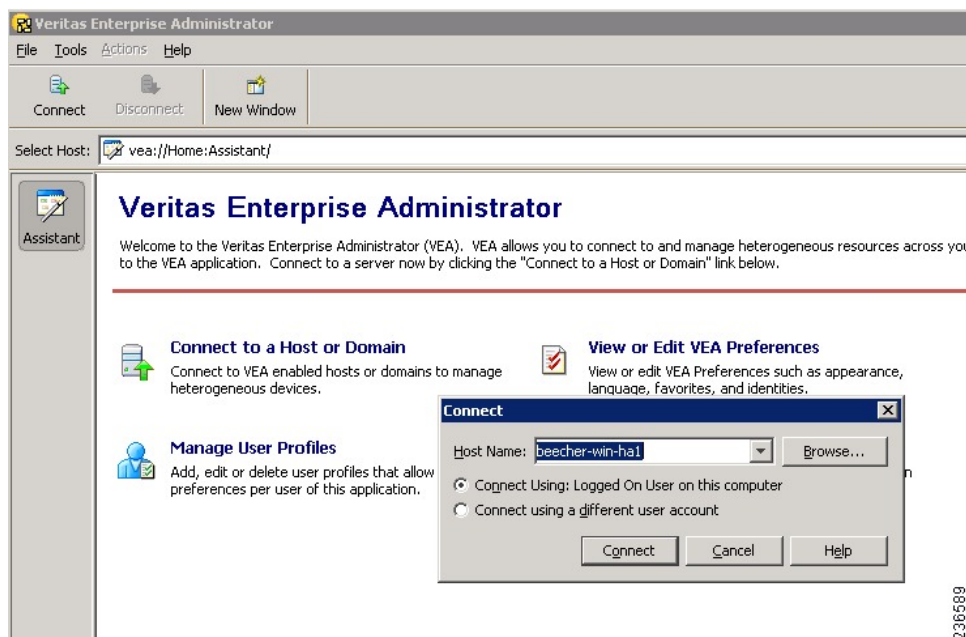
Installing Security Manager on the Primary Server

Use this procedure to install the primary instance of Security Manager that is used in production and is protected by the HA/DR configuration.

To install Security Manager on the primary server, follow these steps:


- Step 1** On the primary server in the cluster, open the Veritas Enterprise Administrator (VEA GUI) application and login.

Figure 236589




236589

- Step 2** Right-click the datadg disk group and select **Import Dynamic Disk Group**.



236582

- Step 3** Make sure the **Import as dynamic disk group** option is selected, and then click **OK**.
- Step 4** Expand the **Volumes** folder under System.
- Step 5** Right-click the cscopx volume and choose **File System > Change Drive Letter and Path**.
- Step 6** Assign the desired drive letter to the cscopx volume and then click **OK**. Refer to the [Local Redundancy Configuration Worksheet, on page 18](#), or [Geographic Redundancy \(DR\) Configuration Worksheet, on page 19](#), for drive assignment.
- Step 7** Install Security Manager according to the Security Manager Installation Guide, while noting the following HA specific items.
- a) When prompted whether to install Security Manager for HA, indicate yes by checking the box.



- b) When prompted for the installation directory, specify: <Selected Drive Letter>\Program Files\CSCOpX.
- c) When prompted to specify the database password, choose an appropriate password and remember it; you will use this password for all Security Manager servers in the HA/DR configuration.

Note Near the end of the Security Manager installation, you might see a message that you are using a multihomed server and that you must update the gatekeeper.cfg file. You can ignore this message, because the agent scripts used in the HA/DR configurations modify this file.

- Step 8** After Security Manager has been installed, reboot the server.
- Step 9** After the system reboots, open the VEA GUI and check to see if the shared disk group is Imported. If the disk group status is Offline, repeat [Step 2](#) through [Step 6](#) to import the disk group and assign the same drive letter used during installation.
- Step 10** Start Security Manager using the online.pl script. For more information, see [Manually Starting, Stopping, or Failing Over Security Manager, on page 57](#).

Note It is necessary to start Security Manager to complete configuration of the Windows registry entries needed for Security Manager to correctly operate.

- Step 11** Allow 5 to 10 minutes for Security Manager to complete startup, then log in to the application's web interface using the following URL: **http://<server hostname or IP address>:1741**. Verify that you can successfully log in.
- Tip** Alternatively, you can use the **pdshow** command to verify that the Cisco Security Manager services are running successfully.
- Step 12** Log out of the application's web interface, then stop Security Manager using the `offline.pl` script. For more information, see [Manually Starting, Stopping, or Failing Over Security Manager, on page 57](#).


Installing Security Manager on Secondary Servers

Use this procedure to install Security Manager on secondary servers. Installing Security Manager on secondary servers is similar to installing it on a primary server, with one important difference. You install Security Manager onto a spare volume (**cscopx_spare**) associated with the specific secondary server, which is used again only if you want to upgrade or uninstall Security Manager. This spare volume must be large enough to hold the Security Manager application with an empty database (~2 GB). You can create the spare volume on the **datadg** disk group if enough space is available or, preferably, on a separate disk group.

To install Security Manager on a secondary server, follow these steps:


- Step 1** On the secondary server, open the Veritas Enterprise Administrator (VEA GUI) application and log in.

Figure 236588



- Step 2** Right-click the `datadg_spare` disk group and select **Import Dynamic Disk Group**.
- Step 3** Make sure the **Import as dynamic disk group** option is selected, and then click **OK**.

Figure 236583




Step 4 Expand the **Volumes** folder under System.

Step 5 Right-click the `cscopx_spare` volume and choose **File System > Change Drive Letter and Path**.

Step 6 Assign the desired drive letter to the `cscopx_spare` volume and then click **OK**. Refer to the [Local Redundancy Configuration Worksheet, on page 18](#), or [Geographic Redundancy \(DR\) Configuration Worksheet, on page 19](#), for drive assignment.

Step 7 Install Security Manager according to the Security Manager Installation Guide, while noting the following HA-specific items.

a) When prompted whether to install Security Manager for HA, indicate yes by checking the box.



- b) When prompted for the installation directory, specify: <Selected Drive Letter>:\Program Files\CSCOpX.
- c) When prompted to specify the database password, choose the same password you chose for the primary server.

Note Near the end of the Security Manager installation, you might see a message that you are using a multihomed server and that you must update the gatekeeper.cfg file. You can ignore this message, because the online script used in the HA/DR configurations modifies this file.

Step 8 After Security Manager has been installed, reboot the server.

Step 9 After the system reboots, open the VEA GUI and check to see if the shared disk group is Imported. If the disk group status is Offline, repeat [Step 2](#) through [Step 6](#) to import the disk group and assign the same drive letter used during installation.

Step 10 Start Security Manager using the online.pl script. For more information, see [Manually Starting, Stopping, or Failing Over Security Manager, on page 57](#).

Note It is necessary to start Security Manager to complete configuration of the Windows registry entries needed for Security Manager to correctly operate.

Step 11 Allow 5 to 10 minutes for Security Manager to complete startup, then log in to the application's web interface using the following URL: **http://<server hostname or IP address>:1741**. Verify that you can successfully log in.

Tip Alternatively, you can use the **pdshow** command to verify that the Cisco Security Manager services are running successfully.

Step 12 Log out of the application's web interface, then stop Security Manager using the offline.pl script. For more information, see [Manually Starting, Stopping, or Failing Over Security Manager, on page 57](#).

Step 13 After installation is complete, unassign the drive letter from the spare volume.

Manually Starting Services in Secondary HA Server

If in Security Manager version 4.13, one or more services do not start up in secondary HA server in DR mode, follow these steps:

Step 1 Run the following command to reset the casuser password:

```
<NMSROOT>\setup\support\resetcasuser.exe
```

Example: C:\Progra~2\CSCOpX\setup\support\resetcasuser.exe

Step 2 Of the two options displayed, on screen, choose option 2 -Enter casuser password. You will be prompted to enter a password for casuser and then to reenter the password for confirmation.

Step 3 If local security policy is configured, add the casuser account to the 'Log on as a service' operation in the local security policy.

Note The following five permissions are assigned and set, automatically, at the time of Security Manager installation: Access this computer from network- casusers, Deny access to this computer from network-casuser, Deny logon locally-casuser, Log on as a batch job-casuser casusers, and Log on as a service- casuser.

Step 4 Run the following command to apply the casuser permission to NMSROOT:

```
C:\Windows\System32\cacls.exe "<NMSROOT>" /E /T /G Administrators:F casusers:F
```

Example: C:\Windows\System32\cacls.exe "C:\Progra~2\CSCOpX" /E /T /G Administrators:F casusers:F

Step 5 Run the following command to set the casuser to database services.

```
<NMSROOT>\bin\perl <NMSROOT>\bin\ChangeService2Casuser.pl casuser <casuserpassword>
```

Example: C:\Progra~2\CSCOpX\bin\perl C:\Progra~2\CSCOpX\bin\ChangeService2Casuser.pl casuser admin123

Veritas Volume Replicator Tasks

Use this procedure to configure replication for a dual geographic cluster configuration where replication is running between the clusters.

To configure replication, follow these steps:


Step 1 Using VEA GUI, connect to the primary and secondary hosts.

Step 2 Make sure that the *datadg* disk group is imported on both the primary and secondary server.

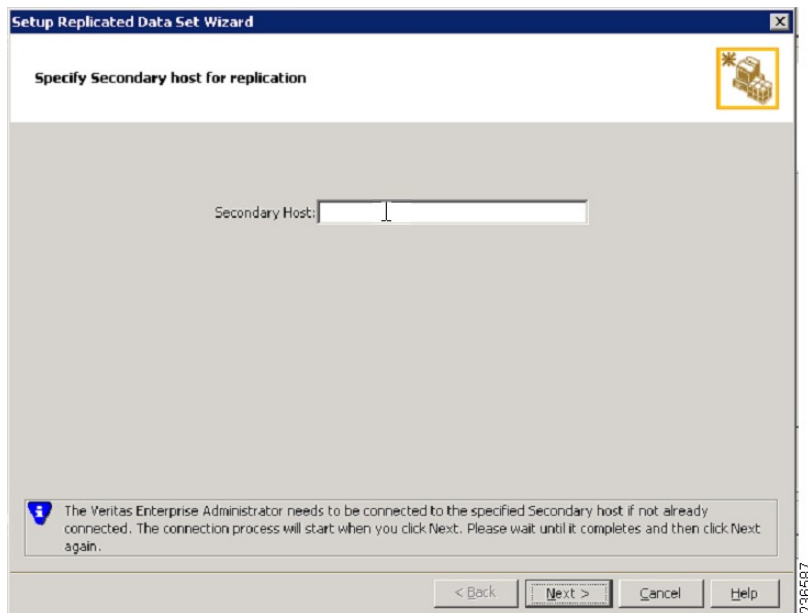
Step 3 Choose **View > Connection > Replication Network**.

Step 4 Select Replication Network from the tree, select the **Setup Replicated Data Set** wizard from the toolbar, and then specify the following on the first panel of the wizard:

- Replicated Data Set Name: **CSM_RDS**
- Replicated Volume Group name: **CSM_RVG**
- Select the primary host from the drop-down list.




- Step 5** Click Next, and on the Select Dynamic Disk Group and volumes to be replicated panel of the wizard, specify the following:
- Dynamic Disk Group: datadg
 - Volumes: cscopx
- Step 6** Click Next. If data_srl is the only other available volume, it will automatically be selected as the storage volume for the replicator log. If more than one additional volume is available, the Storage Replicator Log panel appears. Specify the following:
- Volume for the Replicator Log: data_srl
- Step 7** Click Next, review the summary information, and then click Create Primary RVG to create the RVG.
- Step 8** After successfully creating the Primary RVG, click Yes when prompted to add a secondary host to the RDS.
- Step 9** On the Specify Secondary host for replication panel, enter the name or IP address of the secondary host in the Secondary Host field.



Step 10 Click Next and on the edit replication settings panel specify the following:

Note For the primary and secondary side IP addresses you can specify the fixed IP addresses of the NIC cards. However, if you use Veritas Cluster Server, you must go back later and update the IP address to use virtual IP addresses under VCS control. Do this from VEA by selecting the secondary RVG in the tree and then choosing **Actions > Change Replication Settings**.

- Primary side IP: <IP address of the primary server>
- Secondary side IP: <IP address of the secondary server>
- Replication Mode: **Synchronous Override**
- Replicator Log Protection: <Choose from **Off**, **Fail**, **DCM**, **AutoDCM** (Default), **Override**>. See the Volume Replicator administrator's guide for descriptions of each choice.



Step 11 Click Next to start replication with the default settings. Select Synchronize Automatically and make sure Start Replication is checked.

Step 12 Click **Next** to display the Summary page, and then click Finish.

Updating Permissions on the Working Volume

When Security Manager is installed, it creates a special local user (casuser) and group (casusers) for running Security Manager. To run the protected instance of Security Manager on secondary servers, you must add the local casusers group permissions to the cscopx volume.

This section contains the following topics:

- [Updating Permissions when using Shared Storage, on page 35](#)
- [Updating Permissions when using Replication, on page 36](#)

Updating Permissions when using Shared Storage

To add the local casusers group permissions for a secondary server when using shared storage, follow these steps:

- Step 1** If it is running on the primary server, stop Security Manager using the offline.pl script. For more information, see [Manually Starting, Stopping, or Failing Over Security Manager, on page 57](#).
- Step 2** Deport the **datadg** disk group from the primary server.
- Step 3** Import the **datadg** diskgroup onto the secondary server.
- Step 4** Assign the primary volume (cscopx) to the selected drive letter using either the VEA GUI or the command line.

- Step 5** From Windows Explorer, right-click the <Selected Drive Letter>:\Program Files\CSCOPx folder and choose the Sharing and Security menu item.
- Step 6** The folder properties dialog box appears. Select the **Security** tab, and then click Add.
- Step 7** In the Select Users or Groups dialog box, click **Location**, and then select the local server from the selection tree.
- Step 8** Enter casusers in the enter object names text box, and then click **Check Names**. The text box should then display <ServerName>\casusers . Click **OK**.
- Step 9** Making sure casusers is selected, check the **Full Control** check box under Allow to grant the casusers group full control.
- Step 10** Click **Advanced**.
- Step 11** Under Advanced Settings, check the **Replace permission entries on all child objects with entries shown here that apply to child objects** check box.
- Step 12** Click **Apply** and wait for the permissions to propagate to all child objects under the CSCOPx directory.
- Step 13** When propagation is complete, click **OK**.
- Step 14** Click **OK** to close the CSCOPx Properties dialog box.
- Step 15** Unassign the drive letter from the cscopx volume.
- Step 16** Deport the datadg disk group from the secondary server.
- Step 17** Import the datadg diskgroup onto the primary server.
- Step 18** Assign the primary volume (cscopx) to the selected drive letter using either the VEA GUI or the command line.

Updating Permissions when using Replication

To add the local casusers group permissions for a secondary server when using replication, follow these steps:

- Step 1** If it is running on the primary server, stop Security Manager using the offline.pl script. For more information, see [Manually Starting, Stopping, or Failing Over Security Manager, on page 57](#).
- Step 2** Unassign the drive letter from the cscopx volume.
- Step 3** Migrate the replication primary to the secondary.
- Step 4** Assign the selected drive letter to the cscopx volume on the secondary server.
- Step 5** From Windows Explorer, right-click the <Selected Drive Letter>:\Program Files\CSCOPx folder and choose the **Sharing and Security** menu item.
- Step 6** The folder properties dialog box appears. Select the **Security** tab and click **Add**.
- Step 7** In the **Select Users or Groups** dialog box click **Location**, and select the local server from the selection tree.
- Step 8** Enter **casusers** in the enter object names text box, and then click **Check Names**. The text box should then display <ServerName>\casusers . Click **OK**.
- Step 9** Making sure casusers is selected, check the Full Control check box under Allow to grant the casusers group full control.
- Step 10** Click **Advanced**.
- Step 11** Under the Advanced Settings, check the **Replace permission entries on all child objects with entries shown here that apply to child objects** check box.
- Step 12** Click **Apply** and wait for the permissions to be propagated to all child objects under the CSCOPx directory.
- Step 13** When propagation is complete, click **OK**.

Note While the permissions are being updated you may encounter an error dialog with the title “Error Applying Security” with the message “An error occurred applying security information to: <Selected Drive Letter>:\Program Files\CSCOpX\log\dcr.log. Access is denied.” You can safely ignore this error and click **Continue** on the error dialog to complete the process of updating permissions.

- Step 14** Click **OK** to close the CSCOpX Properties dialog box.
- Step 15** Unassign the drive letter from the cscopx volume.
- Step 16** Migrate the replication back to the primary server.
- Step 17** Assign the selected drive letter to the cscopx volume on the primary server.
-

Veritas Cluster Server Tasks

This section describes the process for setting up and configuring the Veritas cluster(s). There are two specific scenarios described:

[Single Local Cluster \(Dual-Node\) Configuration, on page 37](#)

[Dual Geographic Cluster Configuration, on page 41](#)

Single Local Cluster (Dual-Node) Configuration

This section covers the setup and configuration of a single, local cluster with two nodes in the cluster (primary and secondary).

This section contains the following topics:

- [Creating the Cluster, on page 37](#)
- [Creating the Application Service Group, on page 38](#)
- [Creating the ClusterService Group, on page 40](#)

Creating the Cluster

To create the cluster, follow these steps:

- Step 1** Create a cluster using the VCS Cluster Configuration wizard, where:
- Cluster Name = CSManager_Primary
 - Cluster ID = 0

Include the primary and secondary servers in the definition of the cluster. Part of the cluster definition in the wizard is to specify the NICs for the private network. VCS uses a private network for communications between cluster nodes for cluster maintenance. You can also assign one of the network Ethernet interfaces to act as low-priority cluster communications interface in case all the dedicated cluster communication interfaces fail.

- Step 2** To start the Cluster Manager, choose **Start > All Programs > Veritas Cluster Manager - Java Console** and log in to the cluster.

Step 3 Using the Cluster Manager, import the CSManager resource type by choosing **File > Import Types**. Browse to the CSManagerTypes.cf file located under %VCS_ROOT%\cluster server\conf\config and click Import.

Creating the Application Service Group

To create the application service group, follow these steps:



Note Perform these steps on both the primary and secondary clusters.

Step 1 Add a service group called APP.

Step 2 Right-click the **APP** service group and select **Add Resource**.

Add the RVG primary resource:

- Resource name = **APP_RVGPrimary**
- Resource Type = **RVGPrimary**
- RvgResourceName = **APP_RVG**

Step 3 Right-click the **APP** service group and select **Add Resource**.

Add the MountV resource:

- Resource name = **APP_MountV**
- Resource Type = **MountV**
- Mount Path = <Selected Drive Letter>:\
- Volume Name = **cscopx**
- VMDg Resource Name = **datadg**

Step 4 Right-click the **APP** service group and select **Add Resource**.

Add the RegRep resource and check the **Critical** and **Enabled** check boxes.

- Resource name = **APP_RegRep**
- MountResName = **APP_MountV**
- ReplicationDirectory = **\REGREP\DEFAULT**
- Keys = **HKLM\Software\JavaSoft\Prefs\vms**

Note Security Manager stores client user preferences in the server registry under HKEY_LOCAL_MACHINE\SOFTWARE\JavaSoft\Prefs\vms. The registry replication agent (RegRep) monitors changes to the specified registry location on the active server and synchronizes these changes to a secondary server in the event of a failover.

Step 5 Right-click the **APP** service group and select **Add Resource**.

Add the Proxy resource:

- Resource name = **APP_NIC_Proxy**
- Resource Type = **Proxy**
- TargetResName = **NIC**

Step 6 Right-click the **APP** service group and select **Add Resource**.

Add the IP resource:

- Resource name = **APP_IP**
- Resource Type = **IP**
- Address = <Virtual IP address allocated for the application>
- SubNetMask = <subnet mask>
- MACAddress = <MAC address of the corresponding NIC card>

Step 7 Right-click the **APP** service group and select **Add Resource**.

Add the CSManager resource:

- Resource name = **APP_CSManager**
- Resource Type = **CSManager**
- PathName = <Selected Drive Letter >:\Program Files\CSCOpX
- EventIPAddress = The same IP address as you used in APP_IP
- CertificateDir = See [Security Certificates for SSL, on page 56](#), for an explanation of this attribute.

Step 8 Link the resources as defined in the following table (see [Figure 12: Resource View: APP Group \(Dual Cluster, Single Node\)](#), on page 77).

Parent Resource	Child Resource
APP_MountV	APP_RVGPrimary
APP_RegRep	APP_MountV
APP_CSManager	APP_RegRep
APP_IP	APP_NIC_Proxy
APP_CSManager	APP_IP

To link resources, follow these steps:

- Right-click the parent resource, and then select **Link**.
The Link Resources dialog box appears.
- Select the child resource, and then click **OK**.

The selected resources are linked.

Creating the ClusterService Group

To create the ClusterService group, follow these steps:



Note Perform these steps on both the primary and secondary clusters.



Tip You can use the VCS Configuration wizard as an alternate method to the procedures in this section for creating the ClusterService group and wac resource for intercluster communications. You can also configure the optional Cluster Manager (Web Console) and Notification components with the VCS Configuration wizard. See the Veritas Cluster Server administrator's guide.

Step 1 Right-click the **CSManager** resource and select **Add Service Group**.

Add a service group called **ClusterService**.

Step 2 Right-click the **ClusterService** service group and select **Add Resource**.

Add the NIC resource:

- Resource name = **NIC**
- Resource Type = **NIC**
- MACAddress = <MAC Address of the NIC card>

Note You can find the MAC address associated with each Ethernet interface using the DOS-level command **ipconfig -all**.

Step 3 Right-click the **ClusterService** service group and select **Add Resource**.

Add the IP resource

- Resource name = **VCS_IP**
- Resource Type = **IP**
- Address = <Virtual IP address allocated for the cluster>
- SubNetMask = <subnet mask>
- MACAddress = <MAC Address of the corresponding NIC card>

Step 4 Right-click the ClusterService service group and select **Add Resource**.

Add the wac resource:

- Resource name = **wac**

- Resource Type = **Process**
- StartProgram = **C:\Program Files\Veritas\Cluster Server\bin\wac.exe**
- StopProgram = **C:\Program Files\Veritas\Cluster Server\bin\wacstop.exe**
- MonitorProgram = **C:\Program Files\Veritas\Cluster Server\bin\wacmonitor.exe**

Step 5 Link the resources as defined in the following table (see [Figure 12: Resource View: APP Group \(Dual Cluster, Single Node\)](#), on page 77).

Parent Resource	Child Resource
wac	VCS_IP
VCS_IP	NIC

To link resources, follow these steps:

- Right-click the parent resource, and then select **Link**.
The Link Resources dialog box appears.
- Select the child resource, and then click **OK**.
The selected resources are linked.

Dual Geographic Cluster Configuration

This section covers the setup and configuration of two clusters geographically separated with a single node in each cluster.



Note You can also create dual geographic cluster configurations with multiple nodes within one or both clusters.

This section contains the following topics:

- [Creating the Primary and Secondary Clusters](#), on page 41
- [Creating the ClusterService Group](#), on page 40
- [Creating the Replication Service Group](#), on page 43
- [Creating the Application Service Group](#), on page 38
- [Creating the Cluster Level Configuration](#), on page 52

Creating the Primary and Secondary Clusters

To create the primary and secondary clusters, follow these steps:

-
- Step 1** Create a cluster on the primary server (in the primary cluster) using the VCS Cluster Configuration wizard, where:
- Cluster Name = CSManager_Primary
 - Cluster ID = 0
- Step 2** Create a cluster on the primary server (in the secondary cluster) using the VCS Configuration wizard, where:
- Cluster Name = CSManager_Secondary
 - Cluster ID = 1
- Step 3** In the primary cluster, start the Cluster Manager by choosing **Start > All Programs > Veritas Cluster Manager - Java Console** and log in to the cluster.
- Step 4** Using the Cluster Manager, import the CSManager resource type by choosing **File > Import Types**. Browse to the CSManagerTypes.cf file located under %VCS_ROOT%\cluster server\conf\config and click Import.
- Step 5** Repeat Steps 3 and 4 for the secondary cluster.
-

Creating the ClusterService Group

To create the ClusterService group, follow these steps:



Note Perform these steps on both the primary and secondary clusters.



Tip You can use the VCS Configuration wizard as an alternate method to the procedures in this section for creating the ClusterService group and wac resource for intercluster communications. You can also configure the optional Cluster Manager (Web Console) and Notification components with the VCS Configuration wizard. See the Veritas Cluster Server administrator's guide.

-
- Step 1** Right-click the **CSManager** resource and select **Add Service Group**.
Add a service group called **ClusterService**.
- Step 2** Right-click the **ClusterService** service group and select **Add Resource**.
Add the NIC resource:
- Resource name = **NIC**
 - Resource Type = **NIC**
 - MACAddress = <MAC Address of the NIC card>

Note You can find the MAC address associated with each Ethernet interface using the DOS-level command **ipconfig -all**.

Step 3 Right-click the **ClusterService** service group and select **Add Resource**.

Add the IP resource

- Resource name = **VCS_IP**
- Resource Type = **IP**
- Address = <Virtual IP address allocated for the cluster>
- SubNetMask = <subnet mask>
- MACAddress = <MAC Address of the corresponding NIC card>

Step 4 Right-click the ClusterService service group and select **Add Resource**.

Add the wac resource:

- Resource name = **wac**
- Resource Type = **Process**
- StartProgram = **C:\Program Files\Veritas\Cluster Server\bin\wac.exe**
- StopProgram = **C:\Program Files\Veritas\Cluster Server\bin\wacstop.exe**
- MonitorProgram = **C:\Program Files\Veritas\Cluster Server\bin\wacmonitor.exe**

Step 5 Link the resources as defined in the following table (see [Figure 12: Resource View: APP Group \(Dual Cluster, Single Node\)](#), on page 77).

Parent Resource	Child Resource
wac	VCS_IP
VCS_IP	NIC

To link resources, follow these steps:

- a) Right-click the parent resource, and then select **Link**.
The Link Resources dialog box appears.
- b) Select the child resource, and then click **OK**.
The selected resources are linked.

Creating the Replication Service Group

To create the replication service group, follow these steps:



Note Perform these steps on both the primary and secondary clusters.

-
- Step 1** Right-click the **CSManager** resource and select **Add Service Group**.
Add a service group called APPrep.
- Step 2** Right-click the **APPrep** service group and select **Add Resource**.
Add the Proxy resource:
- Resource name = **VVR_NIC_Proxy**
 - Resource Type = **Proxy**
 - TargetResName = **NIC**
- Step 3** Right-click the APPrep service group and select **Add Resource**.
Add the IP resource:
- Resource name = **VVR_IP**
 - Resource Type = **IP**
 - Address = <Virtual IP address allocated for replication>
 - SubNetMask = <subnet mask>
 - MACAddress = <MAC address of the corresponding NIC card>
- Step 4** Right-click the APPrep service group and select **Add Resource**.
Add the VMDg resource:
- Resource name = **datadg**
 - Resource Type = **VMDg**
 - DiskGroupName = **datadg**
- Step 5** Right-click the **APPrep** service group and select **Add Resource**.
Add the VvrRvg resource:
- Resource name = **APP_RVG**
 - Resource Type = **VvrRvg**
 - RVG = **CSM_RVG**
 - VMDGResName = **datadg**
 - IPResName = **VVR_IP**
- Step 6** Link the resources as defined in the following table (see [Figure 13: Resource View: APPrep Group \(Dual Cluster, Single Node\)](#), on page 78).

Parent Resource	Child Resource
VVR_IP	VVR_NIC_Proxy
APP_RVG	VVR_IP
APP_RVG	datadg

To link resources, follow these steps:

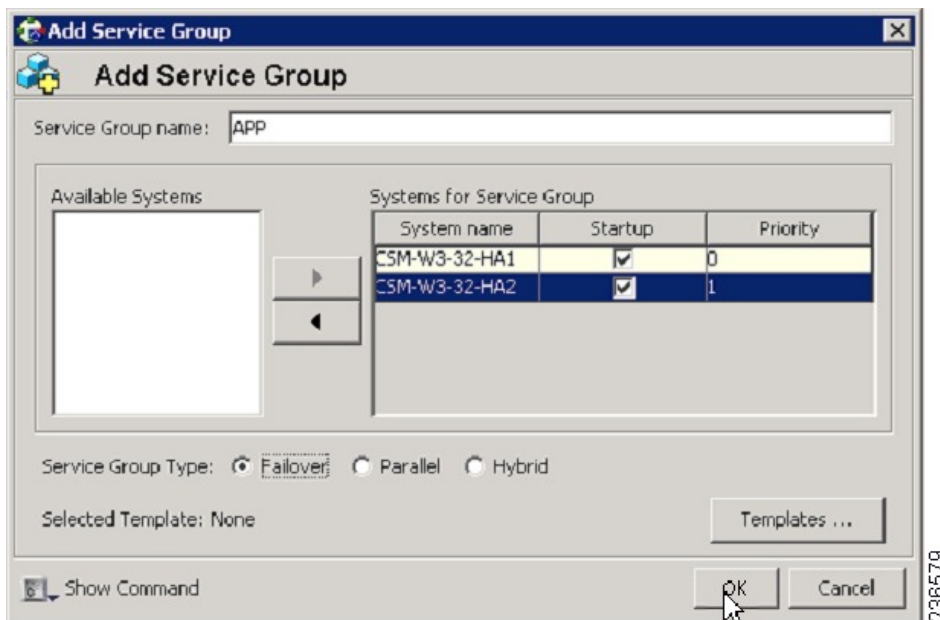
- a) Right-click the parent resource, and then select **Link**.
The Link Resources dialog box appears.
- b) Select the child resource, and then click **OK**.
The selected resources are linked.

Creating the Application Service Group

To create the application service group, follow these steps:

Step 1 Right-click the **CSManager** resource and select **Add Service Group**.

Add a service group called APP, and include both servers for this service group with the Startup option checked for each server and the service group type of Failover.

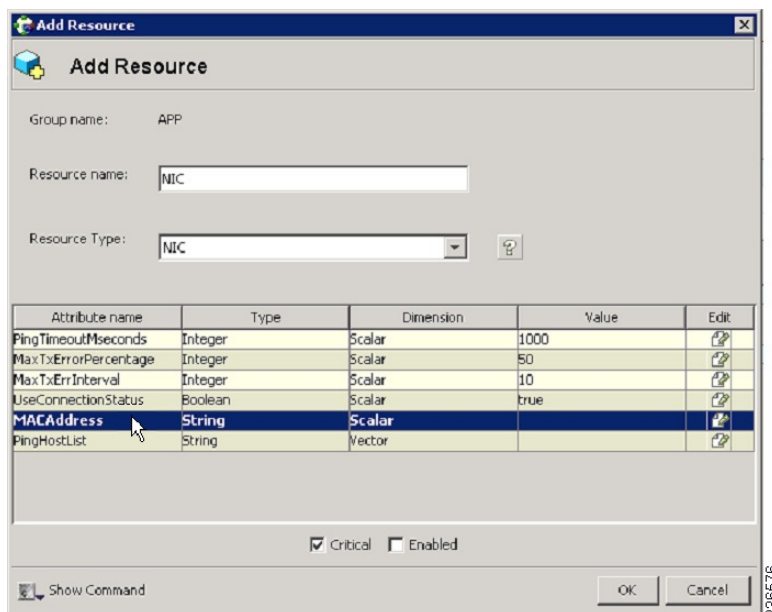


Step 2 Right-click the **APP** service group and select **Add Resource**.

Add the NIC resource and check the Critical and Enabled check boxes.

- Resource name = NIC
- Resource Type = NIC
- MACAddress = <MAC address of the NIC used for accessing the Security Manager application>, which is defined uniquely for each server in the cluster.

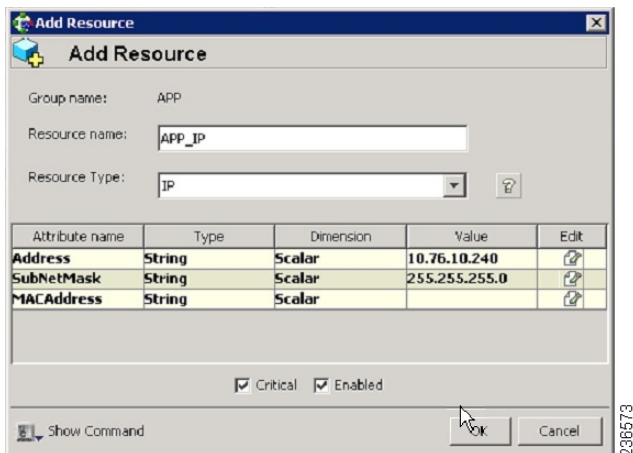
Note You can find the MAC address associated with each Ethernet interface using the DOS-level command `ipconfig -all`.



Step 3 Right-click the **APP** service group and select **Add Resource**.

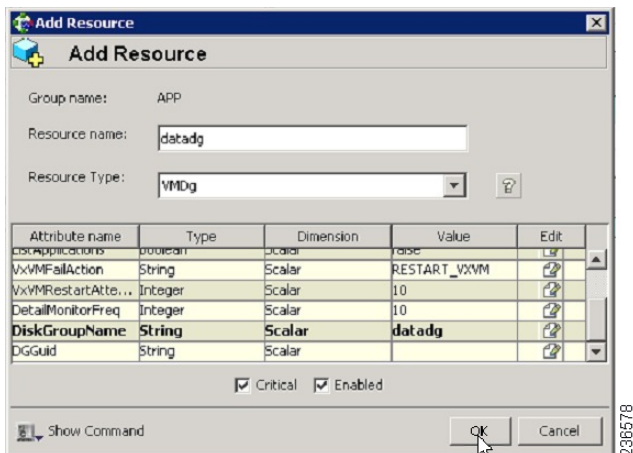
Add the IP resource and check the Critical and Enabled check boxes.

- Resource name = APP_IP
- Resource Type = IP
- Address = <Virtual IP address allocated for use by the Security Manager application>(defined as a Global attribute)
- SubNetMask = <subnet mask>(defined as a Global attribute)
- MACAddress = <MAC Address of the NIC used for accessing the Security Manager application>, (defined for each server in the cluster)



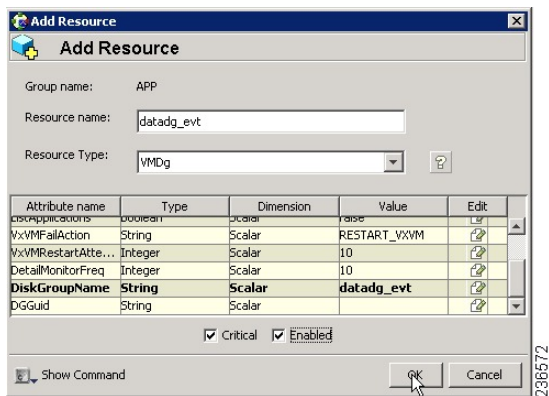
236573

- Step 4** Right-click the **APP** service group and select **Add Resource**.
Add the VMDg Resource and check the Critical and Enabled check boxes.
- Resource name = datadg
 - Resource Type = VMDg
 - DiskGroupName = datadg(defined as a Global attribute)



236578

- Step 5** Right-click the **VMDg** resource group and select **Add Resource**.
Add the datadg_evt resource and check the Critical and Enabled check boxes.
- Resource name = datadg_evt
 - Resource Type = VMDg
 - DiskGroupName = datadg_evt(defined as a Global attribute)



Step 6 Right-click the **APP** service group and select **Add Resource**.

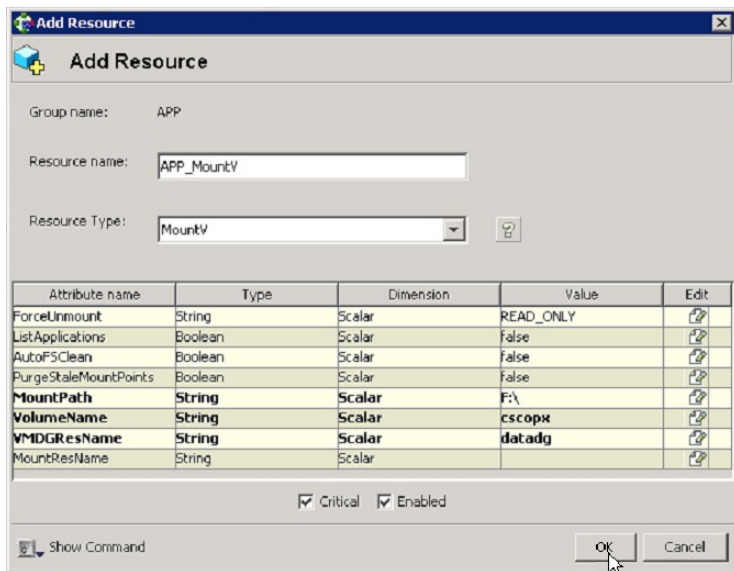
Add the MountV Resource and check the Critical and Enabled check boxes.

- Resource name = APP_MountV
- Resource Type = MountV
- MountPath = <Selected Drive Letter>:\(defined as a Global attribute)
- VolumeName = cscopx(defined as a Global attribute)
- VMDGResName = datadg(defined as a Global attribute)
- ForceUnmount = {NONE, READ-ONLY, ALL}

Defines whether the agent unmounts the volume forcibly when it is being used by other applications. The following choices are available:

- **NONE**: The agent does not unmount the volume if an application is accessing it.
- **READ-ONLY**: The agent unmounts the volume if applications are accessing it in a read-only mode.
- **ALL**: The agent unmounts the volume regardless of the type of access an application has.

The default is NONE. If the volume cannot be unmounted, automatic failover to the secondary server might be prevented, so you might want to select a value of READ-ONLY or ALL.



Step 7 Right-click the **MountV** resource group and select **Add Resource**.

Add the MountV_evt Resource and check the Critical and Enabled check boxes.


- Resource name = APP_MountV_evt
- Resource Type = MountV
- MountPath = <Selected Drive Letter>:\(defined as a Global attribute)
- VolumeName = cscopx_evt(defined as a Global attribute)
- VMDGResName = datadg_evt(defined as a Global attribute)
- ForceUnmount = {NONE, READ-ONLY, ALL}

Defines whether the agent unmounts the volume forcibly when it is being used by other applications. The following choices are available:

- NONE: The agent does not unmount the volume if an application is accessing it.
- READ-ONLY: The agent unmounts the volume if applications are accessing it in a read-only mode.
- ALL: The agent unmounts the volume regardless of the type of access an application has.

The default is NONE. If the volume cannot be unmounted, automatic failover to the secondary server might be prevented, so you might want to select a value of READ-ONLY or ALL.

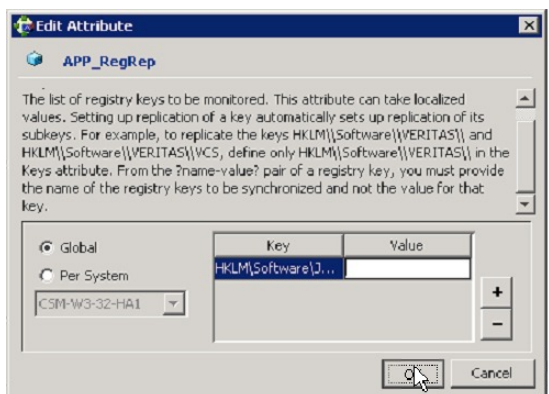
Creating the Application Service Group



Step 8 Right-click the **APP** service group and select **Add Resource**.

Add the RegRep resource and check the Critical and Enabled check boxes.

- Resource name = APP_RegRep
- Resource Type = RegRep
- MountResName = APP_MountV(defined as a Global attribute)
- ReplicationDirectory = \REGREP\DEFAULT(defined as a Global attribute)
- Keys (defined as a Global attribute)Key = HKLM\Software\JavaSoft\Prefs\vmsValue = <blank>




Note Security Manager stores client user preferences in the server registry under HKEY_LOCAL_MACHINE\SOFTWARE\JavaSoft\Prefs\vms. The registry replication agent (RegRep) monitors changes to the specified registry location on the active server and synchronizes these changes to a secondary server in the event of a failover.

Step 9 Right-click the **APP** service group and select **Add Resource**.

Add the CSManager resource and check the Critical and Enabled check boxes.

- Resource name = APP_CSManager
- Resource Type = CSManager
- PathName = <Selected Drive Letter>:\Program Files\CSCOPx\ (defined as a Global attribute)
- EventIPAddress = The same IP address as used in APP_IP(defined as a Global attribute)

- CertificateDir = See [Security Certificates for SSL](#), on page 56, for an explanation of this attribute.

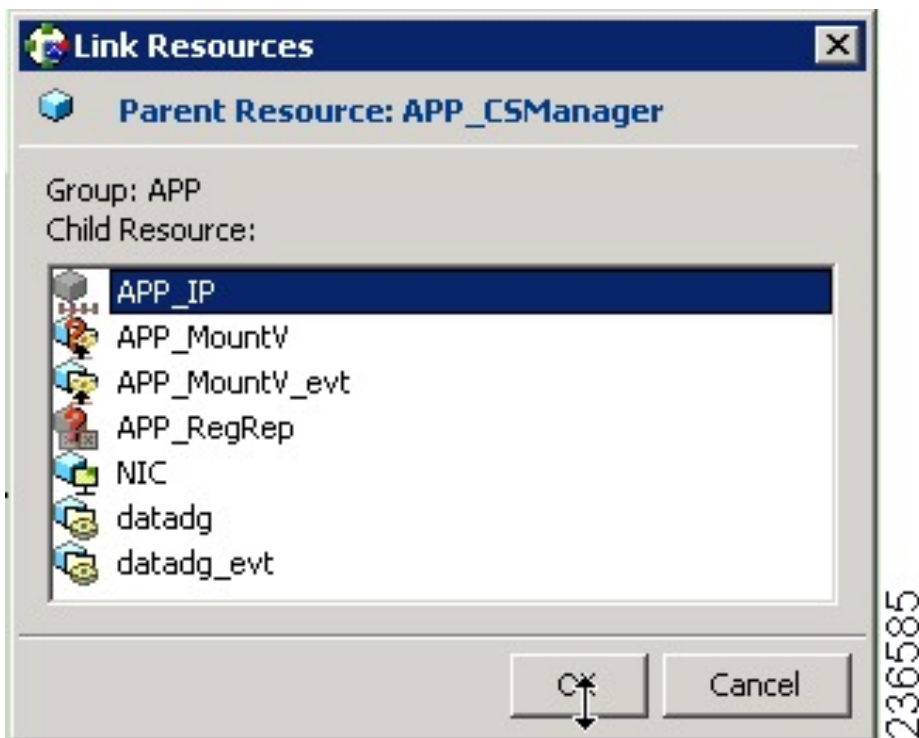
**Step 10**

Link the resources as defined in the following table (see [Figure 12: Resource View: APP Group \(Dual Cluster, Single Node\)](#), on page 77).

Parent Resource	Child Resource
APP_CSManager	APP_RegRep
APP_CSManager	APP_IP
APP_IP	NIC
APP_RegRep	APP_MountV
APP_RegRep	APP_MountV_evt
APP_MountV	datadg
APP_MountV_evt	datadg_evt

To link resources, follow these steps:


- Right-click the parent resource, and then select **Link**.
The Link Resources dialog box appears.



- b) Select the child resource, and then click **OK**.

The selected resources are linked.

When all links have been made, your resource view should look like the following:



Creating the Cluster Level Configuration

To create the cluster level configuration, follow these steps:

-
- Step 1** Link the APP service group as the parent of the APPrep service group with an online local firm dependency. Perform this step on both the primary and secondary clusters.
- Step 2** Under the cluster properties specify the cluster address, which is the same IP address that you used in the VCS_IP resource.
- Step 3** From the primary cluster, choose **Edit > Add/Delete Remote Cluster** to use the Remote Cluster Configuration wizard to add the secondary cluster.
- Step 4** From the primary cluster, choose **Edit > Configure Global Groups** to use the Global Group Configuration wizard to configure the APP service group as a global group.
- See [Figure 15: Service Group View \(Dual Cluster, Single Node\)](#), on page 79.
-



CHAPTER 5

Maintenance Activities

This chapter describes maintenance activities related to Security Manager when it is used in an HA/DR configuration. This chapter containing the following topics:

- [Customizing VCS Behavior, on page 55](#)
- [Security Certificates for SSL, on page 56](#)
- [Manually Starting, Stopping, or Failing Over Security Manager, on page 57](#)
- [Integrating Cisco Secure ACS with Security Manager, on page 59](#)
- [Upgrading Security Manager, on page 60](#)
- [Backing Up Security Manager, on page 60](#)
- [Uninstalling Security Manager, on page 61](#)
- [Migrating a Non-HA Security Manager to HA, on page 61](#)

Customizing VCS Behavior

VCS supports an extensive number of variables to control VCS behavior, such as responses to resource failures. If you followed the default installation as described in this document, here are some of the resulting failover behaviors. You should review these and other behavior controls as described in *Veritas Cluster Server User's Guide*.

- If Security Manager fails, VCS does not try to restart the application on the same server; instead, VCS fails it over to the standby server in the cluster. However, you can use a resource-level attribute, `RestartLimit`, to control the number of times the agent tries to restart the resource before declaring the resource as faulted.
- When first trying to bring the Security Manager application online on a given server, VCS will attempt to bring the resource online only once. The `OnlineRetryLimit` resource-level attribute specifies the number of times the online entry point is retried if the initial attempt fails.
- By default, VCS runs the Security Manager application monitor script every 60 seconds. This means that it can take up to 60 seconds to detect an application failure. The `MonitorInterval` is a resource-level attribute that can be adjusted.
- If you are using dual clusters, failover between the clusters is a manual operation by default. This avoids having both clusters running the application simultaneously. If communication between the clusters is lost (which can readily happen if no redundant paths exist between geographically separated data centers), VCS cannot determine whether the remote cluster failed or a communication problem exists. If you prefer

automatic failover between clusters, you can configure it with the ClusterFailOverPolicy attribute on the APP service group.

Security Certificates for SSL

Security Manager allows configuring the use of Secure Socket Layer (SSL) encryption between the server and the client browser or application. SSL encryption requires the creation and placement of a digital certificate on the server. Part of the identity information contained in the digital certificate is the Common Name (CN) or “Host Name” as shown on the Common Services web GUI. In a HA/DR configuration where there are multiple servers and corresponding hostnames you may want to take special steps to ensure that you maintain a certificate that matches the hostname or IP address used to access the application.

In the case of a single cluster, you access the application with a single virtual IP address or virtual hostname. In this case you should create a certificate with the CN equal to the virtual IP address or virtual hostname. Because the virtual IP or virtual hostname address is valid regardless of the server in the cluster running the application, you do not need to update the digital certificate files in the event of a failover.

However, in the case of a dual geographic cluster configuration, each cluster has its own IP address or hostname associated with the application. As a result, if the digital certificate file has been created to match one cluster, it will not match when the application fails over to the other cluster. In this case you might want to update the digital certificate files to match the other cluster in the event of an inter-cluster failover.



Note If you use a virtual hostname to access the application, you can avoid having to update the certificates for an inter-cluster failover by instead using DNS updating. In the event of an inter-cluster failover, DNS is updated with the new IP address associated with the virtual hostname. Because clients are using the same virtual hostname in all cases to access the application, there is no need to update the certificate files.

The Security Manager Agent for VCS can automatically copy digital certificate files stored on a non-shared, non-replicated local directory prior to starting the application. However, you need to place the appropriate files in this directory on each server in the clusters. The directory is specified to the agent using the CertificateDir parameter.

In the case of a geographic redundancy (DR) configuration where there is a single server at each site, a simpler option is available. You can configure the agent to regenerate the certificate files based on the hostname of the server. This works because there are no virtual IP addresses or virtual hostnames involved. To configure the agent for this behavior, specify the keyword regen for the value of the CertificateDir parameter.

When Security Manager is installed, it creates by default a self-signed certificate matching the local hostname of the server. If appropriate for your configuration, to generate a self-signed certificate matching a virtual IP address or virtual hostname, follow this procedure:

-
- Step 1** Log in to the web browser interface of the server (<http://<hostname or IP address>:1741>).
 - Step 2** Access the self-signed Certificate Setup screen as follows:
 - a) On the Cisco Security Management Suite homepage, click **Server Administration**.
 - b) From the menu on the Server Admin page, select **Server > Single Server Management > Certificate Setup**.
 - Step 3** Populate the fields of the certificate and specify either the virtual IP address or virtual hostname in the CN field, then click Apply.

The following certificate-related files are generated in the NMSROOT\MDC\Apache\conf\ssl directory:

- server.key
- server.crt
- server.pk8
- server.csr
- openssl.conf
- chain.cer

If you are using a single cluster, no further action is required. However, if you are using a dual geographic cluster configuration with multiple servers in each cluster, you should copy the certificate-related files listed above to a non-shared, non-replicated local directory on each server in the cluster. You should then do the same procedure for the secondary cluster, except this time specify the virtual IP address or virtual hostname of the secondary cluster. When you define the CSManager resource, specify the selected non-shared, non-replicated local directory for the **CertificateDir** attribute. The agent then automatically copies the certificate files to the appropriate working directory after a failover, prior to starting the application.

Manually Starting, Stopping, or Failing Over Security Manager

In a non-HA/DR configuration, you normally start and stop Security Manager with the Windows Services application or its command-line equivalents, net start and net stop. However, in an HA/DR configuration, you must not use this approach. Specific scripts are provided for starting and stopping Security Manager in an HA/DR configuration. These scripts perform additional procedures necessary if you start Security Manager on a different server. These scripts and others make up the Security Manager agent for VCS. The agent enables VCS to control and monitor Security Manager. If you are not using VCS, you can use these scripts to manually start and stop Security Manager.

The section contains the following topics:

- [VCS Case, on page 57](#)
- [Non-VCS Case, on page 58](#)

VCS Case

If you are using VCS, you should use the VCS controls to manually start, stop, or fail over the Security Manager service group (APP). In VCS terminology, start and stop are referred to as online and offline, respectively. You can bring online, bring offline, or fail over the Security Manager service group using the VCS GUI or the VCS command-line interface. Appendix B, [High Availability and Disaster Recovery Certification Test Plan, on page 81](#), has examples for performing such operations.



Caution

If you manually stop Security Manager outside of VCS (such as by using net stop), VCS views this as an application failure and tries to initiate recovery.

Non-VCS Case

If you are not using VCS, you can use the online and offline scripts provided with Security Manager to start and stop Security Manager. These scripts can be found at:

\$NMSROOT\MDC\athena\ha\agent\Veritas60 (for Veritas 6.0.1)

\$NMSROOT\MDC\athena\ha\agent\Veritas602 (for Veritas 6.0.2)

\$NMSROOT\MDC\athena\ha\agent\Veritas61 (for Veritas 6.1)

\$NMSROOT\MDC\athena\ha\agent\Veritas70 (for Veritas 7.0)

\$NMSROOT\MDC\athena\ha\agent\Veritas72 (for Veritas 7.2)

\$NMSROOT\MDC\athena\ha\agent\Veritas74 (for Veritas 7.4)

Windows Server 2012, 2012R2 Syntax for Veritas 6.0.1, Veritas 6.0.2, Veritas 6.1, Veritas 7.0, Veritas 7.2, Veritas 7.4, and Veritas 7.4.2:

```
perl online.pl CSManager <PathName> <EventIPAddress> [ <CertificateDir>|regen ]
```

For example:

```
perl online.pl CSManager F:\Progra~1\CSCOpX 192.0.2.1
```

Note You must select the Run as administrator option when opening the Command Prompt.

Syntax	Description
<PathName>	The Security Manager installation path (for example, “F:\Program Files\CSCOpX”). If the installation path contains spaces, enclose the argument in quotes.
<EventIPAddress>	The IP address that the Security Manager application uses for client/server and server/device communications.
<CertificateDir>	Optional. Allows you to specify a nonshared, non-replicated local directory where SSL certificate files are kept. If specified, the script copies these files to the appropriate directory under the installation directory for use by the application. If the keyword regen is used, the script regenerates the SSL certificate based on the local hostname of the server. Regardless of the value used for this parameter, if the hostname of the server matches that of the Security Manager application files, no actions are taken on the certificates. See also Security Certificates for SSL, on page 56 .

The offline script syntax for Windows Server 2012 , 2012R2 is shown below:

Windows Server 2012 , 2012R2 Syntax for Veritas 6.0.1, Veritas 6.0.2, Veritas 6.1, Veritas 7.4, and Veritas 7.4.2:

```
perl offline.pl CSManager <PathName> <EventIPAddress>
```

For example:

```
perl offline.pl CSManager F:\Progra~1\CSCOpX 192.0.2.1
```

Note You must select the Run as administrator option when opening the Command Prompt.

Syntax	Description
<PathName>	The Security Manager installation path (for example, "F:\Program Files\CSCOpX"). If the installation path contains spaces, enclose the argument in quotes.
<EventIPAddress>	The IP address that the Security Manager application uses for client/server and server/device communications.

For ease of use, you might want to create an online and offline batch file (for example, online.bat and offline.bat), which includes the appropriate attributes for your configuration.

To perform a manual failover, you can use VEA or the command line to transfer the primary role within your replicated volume group. If both the primary and secondary server are functioning, you can migrate the primary role to the secondary (effectively reversing the direction of replication) or, if the primary server has failed and is unavailable, you can have the secondary server take over the primary role (with or without fast-failback). Refer to the Veritas Volume Replicator administrator's guide for details.

The following is an outline of the manual failover procedure for a disaster recovery configuration using replication between two servers:

-
- Step 1** Stop Security Manager on the primary server using the offline.pl script.
 - Step 2** Unassign the drive letter from the volume used for Security Manager on the primary server.
 - Step 3** Migrate ownership from the primary server to the secondary server using the VEA GUI.
 - Step 4** Assign the drive letter for the volume used for Security Manager on the secondary server.
 - Step 5** Start Security Manager on the secondary server using the online.pl script.

Note If you are migrating/failing-over to the secondary server for the first time, you must upgrade the file permissions for the casusers group. This is a one-time activity. For more information, see [Updating Permissions on the Working Volume, on page 35](#).

Integrating Cisco Secure ACS with Security Manager

As described in the Installation Guide for Cisco Security Manager, you can integrate Cisco Secure ACS with Security Manager to provide enhanced authorization for Security Manager users. In an HA/DR configuration,

you need to add each Security Manager server involved in the configuration as a AAA client in ACS. When you specify the server in ACS, specify the fixed IP address associated with server's physical hostname.

If you are using an HA/DR configuration for Security Manager with ACS integration, you should also deploy multiple ACS servers to avoid having ACS become a single point of failure. If you only have one ACS server and it fails, you cannot log in to Security Manager without taking corrective action to either restore ACS or reset the Security Manager server to use local authentication. ACS supports the deployment of a primary ACS along with multiple secondary ACSs, where database replication is used to keep the secondary ACSs synchronized with the primary ACS. Security Manager supports specifying up to three ACSs, so if the first ACS is unavailable, it tries the second, and finally the third, if necessary.



Note Beginning with Cisco Security Manager 4.21, Cisco Identity Services Engine (ISE) can be used for authentication purposes, in the place of earlier ACS server.

Upgrading Security Manager

Security Manager upgrades come in various forms:

- Major releases (change in the first number of the release, for example, 3.x to 4.x)
- Minor releases (change in the second digit of the release, for example, 3.1 to 3.2)
- Maintenance releases (change in the third digit of the release, for example, 3.1 to 3.1.1)
- Service packs (identified by a service pack identifier, such as SP2 for Security Manager 3.1)

When you upgrade Security Manager in an HA/DR configuration, the main difference is whether it is necessary to upgrade just the primary server with the active instance of Security Manager or to also upgrade the secondary servers, which have only a spare copy of Security Manager for establishing the correct registry configuration necessary for Security Manager to run on the server. If an upgrade modifies the registry, you must perform the upgrade on all servers in the HA/DR configuration. Normally service packs do not affect the registry, so it is sufficient to install the service pack just on the primary server. For major, minor, or maintenance releases, normally you should upgrade all servers. However, check the readme file or release notes for exceptions to these guidelines.

When upgrading a secondary server, you must mount the spare copy of the Security Manager to the standard \$NMSROOT (such as F:\Program Files\CSCOpX) path used on all servers in the configuration and then install the regular upgrade. This ensures that the registry settings are correct for running the upgraded version of Security Manager on any secondary server.

Before you upgrade, stop VCS on all servers (using **hastop -all -force** on any server in the cluster stops VCS on all servers in the cluster and leaves the application and its resources operational). If you are upgrading on all servers and your configuration uses replication, you should pause or stop the replication during the upgrade and then synchronize the secondary servers after the upgrade is complete.

Backing Up Security Manager

An HA/DR deployment configuration of Security Manager does not replace the need for backing up Security Manager regularly. The HA/DR configuration protects you against loss of data or application downtime due

to hardware failures; however, it does not protect you against user actions such as accidentally or maliciously modifying or deleting important information maintained in Security Manager. Therefore, you should continue to back up the Security Manager database and information files; you can use the backup feature in Security Manager.

You should back up only the primary active instance of Security Manager, not the spare instances associated with secondary servers. Security Manager can be restored on any server in the HA/DR configuration or any server that has the compatible Security Manager application installed.

Uninstalling Security Manager

To uninstall Security Manager from all servers in the HA/DR configuration, follow these steps:

-
- Step 1** Make sure Security Manager is running on the primary server in the primary cluster.
 - Step 2** Using the Cluster Explorer, right-click the **APP_CSManager** resource and uncheck the **critical** check box. You are prompted to switch to read/write mode, so click **Yes** when this dialog box appears.
 - Step 3** Right-click the **APP_CSManager** resource and select **Offline** on the primary server. Wait for Security Manager to go offline.
 - Step 4** Perform required maintenance activities, as needed.
 - Step 5** Start the daemon manager manually on the server, using the **net start crmdmgtd** command.
 - Step 6** The **APP_CSManager** will come online; check the **critical** check box.
 - Step 7** Delete the **APP_CSManager** resource and then save the VCS configuration.
 - Step 8** If you are using replication, stop replication using the VEA GUI.
 - Step 9** To uninstall Security Manager on the primary server, choose **Start > All Programs > Cisco Security Manager > Uninstall Cisco Security Manager**.
 - Step 10** On the secondary server, import the disk group if not already done, which contains the **cscopx_spare** volume, using either the VEA GUI or the command line.
 - Step 11** Assign the selected drive letter to the **cscopx_spare** volume using either the VEA GUI.
 - Step 12** To uninstall Security Manager on the primary server, choose **Start > All Programs > Cisco Security Manager > Uninstall Cisco Security Manager**.
 - Step 13** Repeat Steps 10 through 12 on any other secondary servers or the primary server in a secondary cluster.
- Note** If you do not plan to re-install Security Manager, you should also delete service groups in VCS associated with Security Manager and the Replicated Volume Group if using replication. You should also delete any unneeded volumes and disk groups as well.
-

Migrating a Non-HA Security Manager to HA

If you have an existing Security Manager installed in a regular non-HA configuration, this section addresses how to migrate that instance to an HA configuration. Use the following steps to perform the migration:

-
- Step 1** Perform a backup of the existing Security Manager instance as described in the *User Guide for CiscoWorks Common Services 3.2*. See the section entitled *Backing Up Data* in the *Configuring the Server* Chapter available here: http://www.cisco.com/en/US/docs/net_mgmt/ciscoverks_common_services_software/3.2/user/guide/admin.html
- Step 2** Create the desired Security Manager HA or DR deployment environment as described in this document.
- Step 3** Restore the backup taken from the original Security Manager instance to the primary server in the HA or DR deployment environment as described in the *User Guide for CiscoWorks Common Services 3.2*. See the section entitled “Restoring Data” available at the link above.
- Step 4** Manually synchronize the database passwords in the registry on any secondary servers with the passwords on the primary server. On the primary server, use the registry editor (**Start > Run > regedit**) to find and note the values for the CWEPWD registry entries under folders cmf, vms, rmeng, and aus under HKEY_LOCAL_MACHINE\SOFTWARE\OBDC\OBDC.INI. Edit the CWEPWD registry values on any secondary machine to match those on the primary.
-



CHAPTER 6

High Availability and Disaster Recovery in Virtual Machines

This document explains how to install Cisco Security Management Suite (Security Manager) in a VMware based high availability (HA) or disaster recovery (DR) environment. Security Manager supports the following scenarios:

- [Host-based Failover \(Local HA\), on page 63](#)
- [Fault Tolerance, on page 65](#)
- [Disaster Recovery, on page 68](#)

The steps to configure Security Manager in the above scenarios are described as follows:

- [Host-based Failover \(Local HA\), on page 63](#)
- [Fault Tolerance, on page 65](#)
- [Disaster Recovery, on page 68](#)

Host-based Failover (Local HA)

In this configuration Security Manager is installed on a virtual machine on an ESXi host within a VMware cluster. In the event of a hardware failure on the existing ESXi host, the host-based failover configuration automatically starts up the same virtual machine (VM) on another host within the VMware cluster.

The VMware HA agent monitors the heartbeats, which are sent every second (by default), between the primary and the secondary hosts to detect host failure. It is recommended that you configure redundant heartbeat networks. This allows reliable detection of failures and helps to prevent isolation conditions from occurring.

The same primary VM, with the same Operating System and Application Volume, is started on a different ESXi host. The hostname and IP address remain the same in host-based failover configuration. This configuration works with shared SAN infrastructure between the physical hosts. This process of failover to another host may take few minutes.



Note The following configuration is meant for reference only. You must refer to the VMware documentation for the specific steps to set up the VMware infrastructure. The steps described in this chapter are not Security Manager specific steps.

Prerequisites for Creating VMware HA Clusters

The following prerequisites must be met for creating VMware clusters:

- All virtual machines and their configuration files must reside on shared storage, such as a Storage Area Network (SAN).
- The ESXi hosts must be configured to have access to the same virtual machine network.
- Each host in the VMware HA cluster must have a host name assigned to it and a static IP address.
- There must be CPU compatibility between the hosts. An ideal cluster is a cluster with exactly the same hardware and memory size.
- It is recommended that you use redundant Service Console and VMkernel networking configuration.

Configuring Security Manager for Host-based Failover

Follow these steps to configure Security Manager for host-based failover:

Step 1 Configure two physical hosts that meet the requirements described in the *Deployment Planning Guide for Cisco Security Manager 4.27*.

Note The CPUs on each of the hosts must be compatible.

Step 2 Install VMware ESXi on each of the hosts that you created in Step 1.

Step 3 Create a VMware cluster and add the hosts to the cluster.

Step 4 Configure vSphere HA settings on the ESXi hosts. See VMware documentation for more information.

Step 5 Create a VM on one of the ESXi hosts. See *Deployment Planning Guide for Cisco Security Manager 4.27* for more information.

Step 6 Install Security Manager on the VM you created in Step 5. See *Installation Guide for Cisco Security Manager 4.27* for more information.

Step 7 Start Security Manager.

In the event of a hardware failure on the ESXi host on which Security manager is installed on a VM, the VM is moved to the other ESXi host within the cluster and the VM is started. This movement takes a few minutes to complete and hence there is a downtime.

Limitations

The following limitations exist in the host-based failover configuration:

- You need to manually restart the virtual machine on the failed host.
- If an application stops running on the VM on a failed host and the application data becomes corrupt, then even though the VM is manually restarted after the failover, the application may still remain unusable.
- If a host in the VMware cluster loses its connection to the heartbeat network but the host itself is running, it is isolated from the cluster. In this event, VMware High Availability solution waits for 12 seconds before it decides that the host is isolated from the cluster.



Note Application-based monitoring is not supported in Security Manager. This means that if a Security Manager process stops running, it will not be restarted automatically. You must manually resolve the problem and restart the process, and then manually restart Security Manager.

Fault Tolerance

In the VMware Fault Tolerance configuration, when a hardware failure is detected on a host, a second VM is created on a different host and Security Manager starts running on the second VM without an interruption of service. VMware Fault Tolerance enables a new level of guest redundancy. VMware Fault Tolerance implies that two copies of the VM are maintained, each on separate hosts. This feature can be enabled by turning on Fault Tolerance on the VM on which Security Manager has been installed.

The key difference between VMware's Fault Tolerance and Host-based Failover (HA) solutions is in the interruption to the VM operation in the event of an ESX/ESXi host failure. Fault tolerant systems instantly transition to a new host, whereas high-availability systems see the VMs fail on the host before restarting on another host. The VM on the host that has failed is called the Primary VM and the VM that takes over is the Secondary VM. The failover from the Primary to Secondary VM is dynamic with the Secondary VM continuing to run from the exact point where the Primary VM left. This process happens automatically with no data loss, downtime, or interruption of services. After the dynamic failover, the Secondary VM becomes the new Primary VM and a new Secondary VM is spawned automatically.



Note The following configuration is meant for reference only. You must refer to the VMware documentation for the specific steps to set up the VMware infrastructure. The steps described in this chapter are not Security Manager specific steps.

Creating Fault Tolerant Systems

Prerequisites

The following prerequisites must be met to be able to create Fault Tolerance systems:


- Make sure that all versions of the VMware software used in a Fault Tolerant environment are compatible as per the list in the vSphere Compatibility Matrix. See [VMware Compatibility Guide](#) for more information.
- On a hardware and guest Operating System level, only certain processors and Operating Systems are supported. For information about the support, see the VMware Guest OS Compatibility Guide at <http://www.vmware.com>. Further, to check whether your existing VMware setup is suitable for Fault Tolerance, run the site survey at https://www.vmware.com/support/shared_utilities.
- Enable Hardware Virtualization in the BIOS for each host in the cluster. Since the process for enabling Hardware Virtualization differs for each BIOS, contact your hardware vendor for specific instructions on enabling Hardware Virtualization.
- Make sure that all ESX hosts used by VMware Fault Tolerance are members of a VMware High Availability (HA) cluster. You must enable VMware HA for VMware Fault Tolerance to function. For information about enabling VMware HA see the vSphere Availability Guide at <http://www.vmware.com>.

- Make sure that the ESX hosts that run the primary and secondary Fault Tolerance nodes are running the same build of ESX. Additionally, make sure to apply the patches that have been released as the patches contain improvements to the VMware Fault Tolerance features.
- Make sure you have configured the following for your environment to enable VMware Fault Tolerance:
 - The virtual machine must reside on shared storage, that is, storage that is visible to all ESX hosts in the cluster.
 - Storage must be FC SAN, iSCSI or NFS, and not local storage.
 - Virtual machines must not have snapshots. If there are snapshots, you must commit them before proceeding.
- Make sure to perform the following tasks to configure networking in your Fault Tolerance environment:
 - Define a separate VMkernel port group for Fault Tolerance logging. See the ESX Configuration Guide for instructions to create the port group.
 - Define the Fault Tolerance logging and VMotion port groups and assign a physical network card for uplink. This network card must be of at least 1GB size. It is recommended that you use a 10GB network card.
 - Enable use of Jumbo Frames for the Fault Tolerance logging. For detailed steps, see the Advanced Networking section of the ESX Configuration Guide.
 - VMware recommends enabling fully redundant NICs to ensure availability, although Fault Tolerance can function without it.

After you have configured your environment as per the list of prerequisites, make sure you turn ON Fault Tolerance as shown in the following figure.

Figure 423485

Figure 5: Turning on Fault Tolerance



Note Security Manager must have a minimum of six virtual CPUs for Small Deployment, with VMware ESXi version 5.102 up to ESXi version 6.0. See *Cisco Security Manager Deployment Planning Guide* for more information.



Note Fault tolerant virtual machine on vCenter Server version 5.x supports one virtual CPU per protected virtual machine. vCenter Server version 6.0 supports up to four virtual CPUs depending on the licensing.

Disaster Recovery

Security Manager uses the VMware vCenter Site Recovery Manager tool with VMware vSphere Replication for disaster recovery and management.

Site Recovery Manager integrates natively with VMware vSphere Replication and supports a broad set of high-performance array-based replication products to reliably copy virtual machines across sites according to business requirements. Site Recovery Manager is an extension to VMware vCenter Server that delivers a disaster recovery solution that helps to plan, test, and run the recovery of virtual machines. Site Recovery Manager can discover and manage replicated datastores, and automate migration of inventory between vCenter Server instances.

System Requirements

Hardware Requirements

For hardware requirements, see the *VMware Site Recovery Manager 6.1 Documentation Center* at <http://www.vmware.com>

Software Requirements

Following are the high level software requirements for setting up the VMware Site Recovery Manager solution for Disaster Recovery:

- Virtual Center 6.0 license applied on both the primary (protected) and recovery sites.
- ESXi Server 6.0 licenses applied on both primary and recovery sites.
- VSphere SRM 6.0 license applied on both primary and recovery sites.
- SQL Server Database for Site Recovery Manager installed on both primary and recovery sites.



Note All VMware tools must be on version 6.0.



Note VMware Disaster Recovery solution has been tested with VMware Site Recovery Manager. However, other VMware solutions might also work with Security Manager.

Configuring VMware Site Recovery Manager

Follow these steps to install Site Recovery Manager on the vCenter server:

Step 1 Start the installation of Site Recovery Manager by clicking **install.exe**.

Step 2 Accept the VMware End User License Agreement.

Step 3 On the VMware vCenter Site Recovery Manager—vSphere Replication window, select **Install vSphere Replication** and then click **Next**.

Step 4 Enter the vCenter Server Address, Port (81, by default), Username, and Password.

Step 5 Accept the security warning.

Step 6 On the VMware vCenter Site Recovery Manager—Certificate Type Selection window, select the Certificate Source as **Automatically Generate a Certificate**.

Step 7 Enter the vCenter Server information:

On the VMware vCenter Site Recovery Manager Extension window, enter the following:

- Local Site name—The VCenter site Fully Qualified Domain Name.
- Administrator Email—The Administrator’s email as per your organizational requirements.
- Additional Email—Any additional email ID that you may wish to enter.
- Local Host—Current Host IP Address; this is automatically populated.
- Listener Ports
 - SOAP Port—Default is 8095.
 - HTTP Port—Default is 9085.
- API Listener Port—Default is 9007.

Step 8 The Site Recovery Manager Server requires its own database, which it uses to store data such as recovery plans and inventory information. The Site Recovery Manager database is a critical part of a Site Recovery Manager installation. You must create the Site Recovery Manager database and establish a database connection before you can install Site Recovery Manager.

On the VMware vCenter Site Recovery Manager—Database Configuration window, enter the following, and then click **Next**:

- Select **Database Client type** from the drop-down list.
- Enter or select **Data Source Name**. Click **ODBC DSN Setup...** to set up a System DSN.
- Enter the Database Username and Password.
- Enter the Connection Count and Maximum Connections.

Figure 6: Database Configuration

VMware vCenter Site Recovery Manager

Database Configuration

Enter VMware vCenter Site Recovery Manager database information.

Database Information

Select database client type.

Database Client:

Enter or select Data Source Name (DSN). Click ODBC DSN Setup button to set up a System DSN.

Data Source Name:

Enter database user credentials.

Username:

Password:

Connection information.

Connection Count: Max. Connections:

InstallShield

Figure 7: ODBC DSN Setup

ODBC Data Source Administrator

User DSN | System DSN | File DSN | Drivers | Tracing | Connection Pooling | About

System Data Sources:

Name	Driver
PCSMHA	SQL Server Native Client 10.0

An ODBC System data source stores information about how to connect to the indicated data provider. A System data source is visible to all users on this machine, including NT services.

Step 9 Click Next to complete the installation of Site Recovery Manager.

Note Each Site Recovery Manager site requires its own instance of the Site Recovery Manager database. Use a different database server instance to run the individual Site Recovery Manager databases for each site. If you use the same database server instance to run the databases for both the sites, and if the database server experiences a problem, neither Site Recovery Manager site will work and you will not be able to perform a recovery.

Configuring vCenter


Follow these steps to configure vCenter on the Protected Site.

The vCenter server details are those that you provided in [Step 7](#). The Site Recovery Manager 5.5 Plug-in Manager will be enabled after you provide the vCenter server information.

Step 1 From the Plug-in Manager, download and install the VMware vCenter Site Recovery Manager extension.

Note The above step is applicable only for Site Recovery Manager version 5.5. For Site Recovery Manager version 6.0 you no longer need to install the Plug-in as it supports vSphere web client.

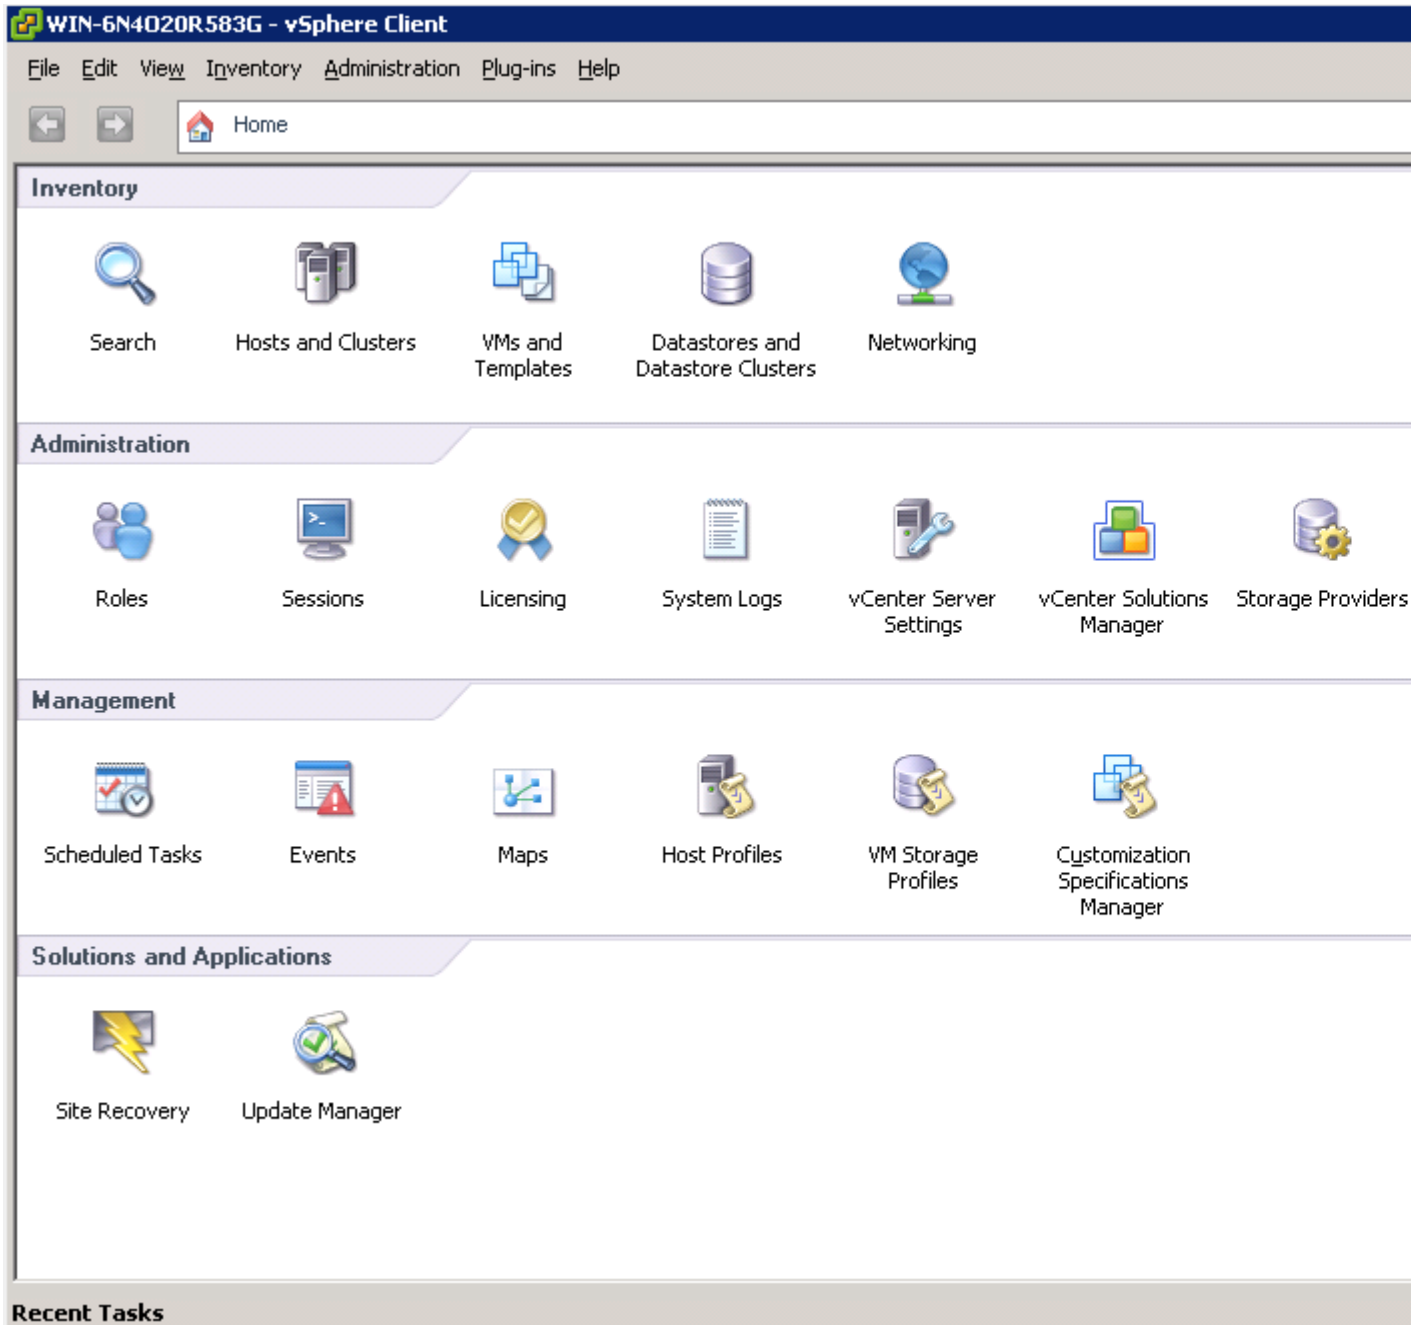
Figure 8: Plug-in Manager



Step 2 To configure and manage Site Recovery Manager, click Site Recovery on the Home page.

Note The above step is applicable only for Site Recovery Manager version 5.5. For Site Recovery Manager version 6.0 you no longer need to install the Plug-in as it supports vSphere web client.

Figure 9: Site Recovery Manager 5.5



The following image shows the vSphere web client that you can use to access Site Recovery Manager version 6.0:

Figure 10: VMware vSphere Web Client



Configuring the Recovery Site

Perform the tasks described in [Configuring VMware Site Recovery Manager, on page 68](#) and [Configuring vCenter, on page 71](#) on the Recovery Site.

Configuring Replication

Perform the following tasks on the Protected Site and Recovery Site. These steps are applicable for site Recovery Manager versions 5.5 and 6.0:

- Configure the Protection and Replication sites
- Configure the Inventory
- Install and configure the vSphere Replication
- Configure Protection Groups
- Configure Recovery Plans

See the VMware User Guide at the *VMware vCenter Site Recovery Manager versions 5.5 and 6.0 Documentation Center* for more information.

Installing Security Manager in Disaster Recovery Environment

After configuring Site Recovery Manager on both the Protected and Recovery sites, you must install Security Manager on the VM of the Protected site. To install Security Manager, see *Installation Guide for Cisco Security Manager 4.27*.

The VMware Site Recovery Manager tool replicates the installation onto the Recovery site. The synchronization between the Protected and recovery sites is performed based on the bandwidth and data size of the Protected site. After enabling Site Recovery Manager on the Protected site host, vSphere Replication performs an initial full synchronization of the source VM and its replica to the Recovery site.

After a full synchronization is completed for the first time, vSphere Replication is performed based on the Recovery Point Objective (RPO) time interval configured in the Site Recovery Manager.



Note You must configure the RPO time interval based on the environment, that is, the bandwidth and data size.



Caution It is recommended that you allow at least 15 minutes of grace time for the Security Manager services to come up on the Recovery site before stating to work with the application. If the grace time is not allowed, Security Manager may not start properly and this might lead to reinstalling the application.



Note If Security Manager is integrated with Access Control Server (ACS), for authentication purposes, you must provide the IP addresses of both the Protected and Recovery sites (where Security Manager has been installed) as AAA client to the ACS server.



Note Beginning with Cisco Security Manager 4.21, Cisco Identity Services Engine (ISE) can be used for authentication purposes, in the place of earlier ACS server.



APPENDIX **A**

VCS Resource Views for the Reference Configurations

This section provides information about Veritas Cluster Server (VCS) resource and service group views for the HA/DR Security Manager configurations described in this document. Figure A-1 through Figure A-5 illustrate the dependencies among resources in a service group or the dependencies among service groups. In the figures, the line between two resources represents a dependency, or parent-child relationship. Resource dependencies specify the order in which resources are brought online and taken offline. During failover, the resources closest to the top of the figure must be taken offline before the resources linked to them are taken offline. Similarly, the resources that appear closest to the bottom of the figure must be brought online before the resources linked to them can come online. A resource that depends on other resources is a parent resource. The figure links a parent resource icon to a child resource icon below it.

This appendix covers the following topics:

- [VCS Resource Views for the Reference Configurations, on page 75](#)

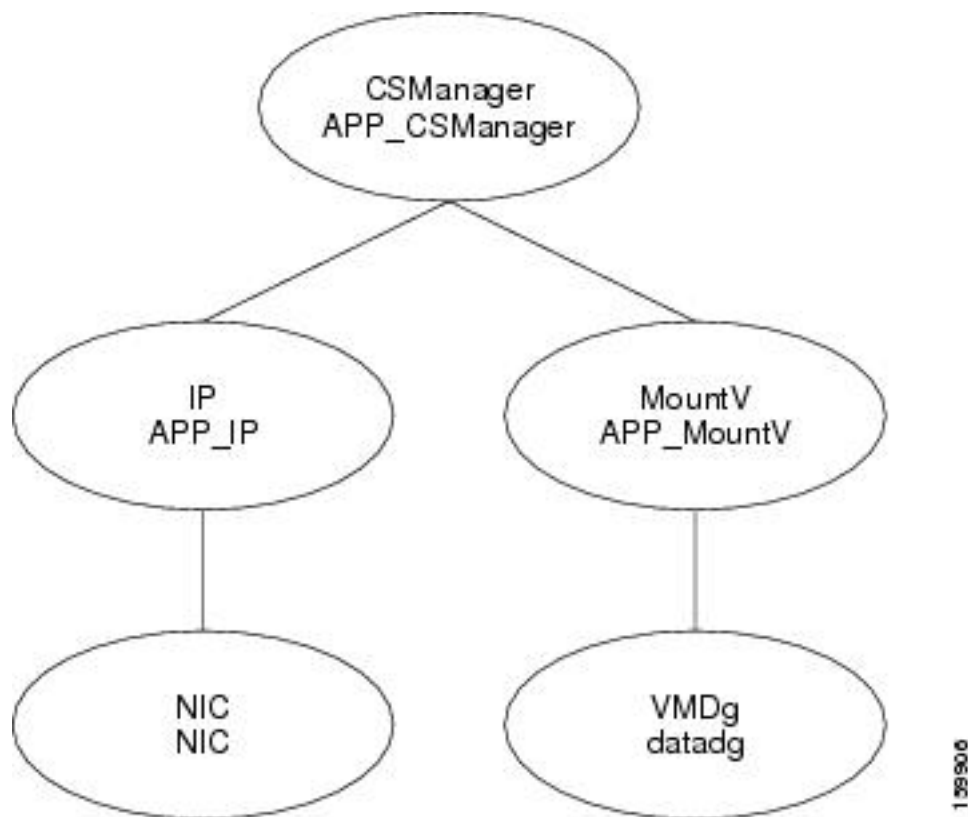
VCS Resource Views for the Reference Configurations

This section provides information about Veritas Cluster Server (VCS) resource and service group views for the HA/DR Security Manager configurations described in this document. Figure A-1 through Figure A-5 illustrate the dependencies among resources in a service group or the dependencies among service groups. In the figures, the line between two resources represents a dependency, or parent-child relationship. Resource dependencies specify the order in which resources are brought online and taken offline. During failover, the resources closest to the top of the figure must be taken offline before the resources linked to them are taken offline. Similarly, the resources that appear closest to the bottom of the figure must be brought online before the resources linked to them can come online. A resource that depends on other resources is a parent resource. The figure links a parent resource icon to a child resource icon below it.

Single Local Cluster (Dual-Node) Configuration

Figure A-1 shows the Veritas Cluster Server (VCS) resource views for a single cluster with two servers in the cluster.

Figure 11: Resource View: APP Group (Single Cluster, Dual Node)



Dual Geographic Cluster (Single-Node) Configuration

Figure 12: Resource View: APP Group (Dual Cluster, Single Node) through Figure 13: Resource View: APPrep Group (Dual Cluster, Single Node) show the Veritas Cluster Server (VCS) resource views for a dual cluster configuration with one server in the cluster.

Figure 12: Resource View: APP Group (Dual Cluster, Single Node)

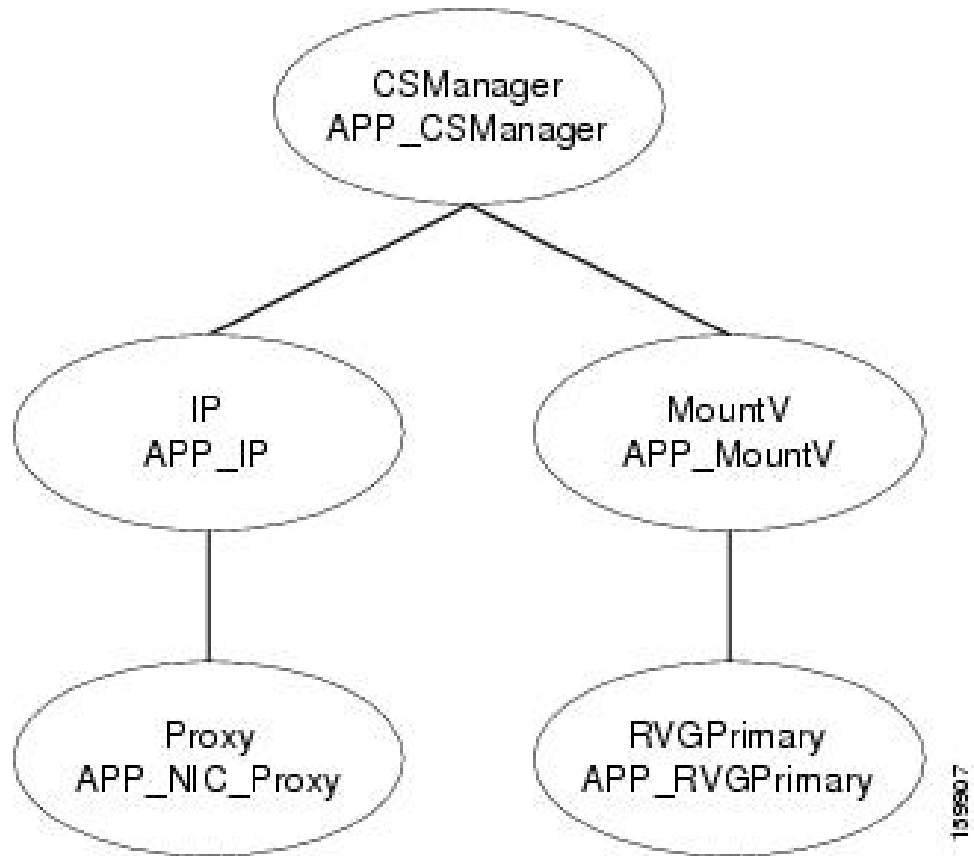


Figure 13: Resource View: APPrep Group (Dual Cluster, Single Node)




Figure 14: Resource View: ClusterService Group (Dual Cluster, Single Node)





Figure 15: Service Group View (Dual Cluster, Single Node)





APPENDIX B

High Availability and Disaster Recovery Certification Test Plan

The HA/DR certification test plan validates that the Security Manager application is highly available and can survive various hardware and software failures. The test plan also covers maintenance activities, such as manually switching the application between servers.



Note Security Manager client sessions require active users to log in again after an application failover. This behavior is equivalent to stopping and starting Security Manager services running on the server.

The following test case categories are contained in this appendix:

- [Manual Switches, on page 81](#)
- [Ethernet/Network Failures, on page 83](#)
- [Server Failures, on page 90](#)
- [Application Failures, on page 96](#)

Manual Switches

This section covers two different types of manual switches. In a single cluster with two servers, you can switch between the two servers in the cluster (intracluster switch); in a dual cluster configuration with a single server in each cluster, you can switch between clusters (intercluster switch).

This section contains the following topics:

- [IntraCluster Switch, on page 81](#)
- [InterCluster Switch, on page 82](#)

IntraCluster Switch

Test Case Title: Manual application switch within a cluster.

Description: The application is manually switched to a different server in the same cluster using VCS.

Test Setup: A dual node cluster ([Figure](#)) in a single cluster configuration.

-
- Step 1** Ensure that the APP service group is running on the primary server. Using the VCS Cluster Explorer, select the APP service group. From the shortcut menu, select Switch To, and choose the secondary server. Alternatively, issue the following command:
- Example:**
- ```
C:\> hagr -switch APP -to secondary_server_name
```
- Step 2** From the Resource view of the APP service group, observe that the resources in the service group go offline on the primary server and then come online on the secondary server. Or issue the following command to observe the status of the APP service group.
- Example:**
- ```
C:\> hagr -state APP
```
- Step 3** From a client machine, launch the Security Manager client, using the virtual hostname or IP address in the Server Name field of the login dialog box. Verify that you can log in to the application successfully.
-

InterCluster Switch

Test Case Title: Manual application switch between clusters.

Description: The application is manually switched to a server in a different cluster using VCS.

Test Setup: A dual cluster configuration as shown in [Figure](#) with a single server in each cluster.

- Step 1** Using the VCS Cluster Explorer, select the **APP** service group. From the shortcut menu, select **Switch To**, then **Remote Switch(...)**, to open the Switch global dialog box. In the dialog box, specify the remote cluster and, if desired, a specific server in the remote cluster. Alternatively, issue the following command:
- Example:**
- ```
C:\> hagr -switch APP -any -clus secondary_cluster_name
```
- Step 2** From the Resource view of the APP service group, observe that the resources in the service group go offline in the primary cluster. Select the root cluster node in the tree and use the Remote Cluster Status view to see that the APP service group goes online on the remote cluster. Or issue the following command to observe the status of the APP service group.
- Example:**
- ```
C:\> hagr -state APP
```
- | #Group | Attribute | System | Value |
|--------|-----------|------------------------------|---------|
| APP | State | csm_primary:<Primary Server> | OFFLINE |
| APP | State | localclus:<Secondary Server> | ONLINE |
- Step 3** From a client machine, launch the Security Manager client by entering the appropriate hostname or application IP address used in the secondary cluster in the Server Name field of the Login dialog box. Verify that you can successfully log in to the application.
- Step 4** Log out of the Security Manager client, and then switch the APP service group to the primary cluster using either the VCS Cluster Explorer or the following command:

Example:

```
C:\> hagr -switch APP -any -clus primary_cluster_name
```

Ethernet/Network Failures

HA/DR configurations have two types of server Ethernet connections. The first are the Ethernet connections used for network communications (public interfaces); the second are Ethernet interfaces dedicated for intracluster communications (private interfaces). This section covers failure test cases for each type of Ethernet interface.

- [Network Communication Failures, on page 83](#)
- [Cluster Communication Failure, on page 88](#)

Network Communication Failures

This section describes the tests used to verify that VCS can detect failure of the network Ethernet ports used for network communications. This section contains the following topics:

- [Network Ethernet Failure on Secondary Server, Single Cluster, on page 83](#)
- [Network Ethernet Failure on Primary Server, Single Cluster, on page 84](#)
- [Network Ethernet Failure on Secondary Server, Dual Cluster, on page 85](#)
- [Network Ethernet Failure on Primary Server, Dual Cluster, on page 87](#)

Network Ethernet Failure on Secondary Server, Single Cluster

Test Case Title: A failure occurs in the network Ethernet connection on the secondary server in a single cluster configuration.

Description: This test case verifies that VCS can detect a failure on the network Ethernet port on the secondary server and then recover after the failure is repaired.

Test Setup: A dual node cluster ([Figure](#)) in a single cluster configuration with a single network connection per server.

Step 1 Verify that the application is running on the primary server.

Step 2 Log in to the application from a client machine.

Step 3 Remove the Ethernet cable from the network port on the secondary server to isolate the server from communicating with the switch/router network. Wait for at least 60 seconds for VCS to detect the network port failure. Verify that VCS detects a failure of the NIC resource on the secondary server by running the following command:

Example:

```
C:\> hastatus -sum
-- SYSTEM STATE
```

Network Ethernet Failure on Primary Server, Single Cluster

```

-- System          State          Frozen
A <PrimaryServer>  RUNNING        0
A <SecondaryServer> RUNNING        0
-- GROUP STATE
-- Group          System          Probed    AutoDisabled  State
B APP            <PrimaryServer> Y         N            ONLINE
B APP            <SecondaryServer> Y         N            OFFLINE | FAULTED
-- RESOURCES FAILED
-- Group          Type          Resource          System
C APP            NIC          NIC              <SecondaryServer>

```

Step 4 Restore the Ethernet cable to the network port on the secondary server. Verify that VCS detects that the failure was cleared by running the following command:

Example:

```

C:\> hastatus -sum
-- SYSTEM STATE
-- System          State          Frozen
A <PrimaryServer>  RUNNING        0
A <SecondaryServer> RUNNING        0
-- GROUP STATE
-- Group          System          Probed    AutoDisabled  State
B APP            <PrimaryServer> Y         N            ONLINE
B APP            <SecondaryServer> Y         N            OFFLINE

```

Network Ethernet Failure on Primary Server, Single Cluster

Test Case Title: A failure occurs in the network Ethernet connection on the primary server in a single cluster configuration.

Description: This test case verifies that VCS can detect a failure on the network Ethernet port of the primary server and automatically switch the application to the secondary server. After the problem is fixed, you can switch the application back to the primary server manually.

Test Setup: A dual node cluster ([Figure 4: Ethernet and Storage Connections for a Dual-Node Site, on page 15](#)) with a single network connection per server.

Step 1 Verify that the application is running on the primary server.

Step 2 Remove the Ethernet cable from the network port on the primary server to isolate the server from communicating with the switch/router network. Verify that VCS detects a failure of the NIC resource and automatically switches the APP service group to the secondary server:

Example:

```

C:\> hastatus -sum
-- SYSTEM STATE
-- System          State          Frozen
A <PrimaryServer>  RUNNING        0
A <SecondaryServer> RUNNING        0
-- GROUP STATE
-- Group          System          Probed    AutoDisabled  State
B APP            <PrimaryServer> Y         N            OFFLINE | FAULTED
B APP            <SecondaryServer> Y         N            ONLINE
-- RESOURCES FAILED
-- Group          Type          Resource          System

```

```
C APP          NIC          NIC          <PrimaryServer>
C APP          IP           APP_IP       <PrimaryServer>
```

Step 3 Verify that you can log in to the application while it is running on the secondary server.

Step 4 Replace the Ethernet cable on the network port of the primary server and manually clear the faulted IP resource on the primary server:

Example:

```
C:\> hares -clear APP_IP -sys primary_server_name
```

Step 5 Manually switch the APP service group back to the primary server.

Example:

```
C:\> hagrps -switch APP -to primary_server_name
```

Network Ethernet Failure on Secondary Server, Dual Cluster

Test Case Title: A failure occurs in the network Ethernet connection on the secondary server in a dual cluster configuration.

Description: This test case verifies that VCS can detect a failure on the network Ethernet port and then recover after the failure is repaired.

Test Setup: A dual cluster configuration (Figure) with a single node in each cluster and a single Ethernet network connection for each server.

Step 1 Verify that the APP service group is running on the primary cluster/server.

Step 2 Log in to the Security Manager from a client machine.

Step 3 Remove the Ethernet cable from the network port on the server in the secondary cluster. This isolates the server from communicating with the switch/router network and interrupts replication. From the primary server, verify that replication was interrupted (disconnected) by running the following command:

Example:

```
C:\> vxprint -P1
Diskgroup = datadg
Rlink      : rlk_172_6037
info       : timeout=500 packet_size=1400
            latency_high_mark=10000 latency_low_mark=9950
            bandwidth_limit=none
state      : state=ACTIVE
            synchronous=off latencyprot=off srlprot=off
assoc      : rvg=CSM_RVG
            remote_host=172.25.84.34
            remote_dg=datadg
            remote_rlink=rlk_172_32481
            local_host=172.25.84.33
protocol   : UDP/IP
flags      : write attached consistent disconnected
```

Step 4 Run the following command from the primary server to verify that communication with the secondary cluster was lost:

Example:

```

C:\> hastatus -sum
-- SYSTEM STATE
-- System                State                Frozen
A <PrimaryServer>       RUNNING                0
-- GROUP STATE
-- Group                System                Probed    AutoDisabled    State
B APP                   <PrimaryServer>       Y         N                ONLINE
B APPrep                <PrimaryServer>       Y         N                ONLINE
B ClusterService       <PrimaryServer>       Y         N                ONLINE
-- WAN HEARTBEAT STATE
-- Heartbeat            To                    State
L Icmp                  csm_secondary         ALIVE
-- REMOTE CLUSTER STATE
-- Cluster              State
M csm_secondary         LOST_CONN
-- REMOTE SYSTEM STATE
-- cluster:system       State                Frozen
N csm_secondary:<SecondaryServer> RUNNING                0
-- REMOTE GROUP STATE
-- Group                cluster:system        Probed    AutoDisabled    State
O APP csm_secondary:<SecondaryServer> Y         N                OFFLINE

```

Step 5 Reattach the network Ethernet cable to the secondary server and verify that replication resumed.

Example:

```

C:\> vxprint -Pl
Diskgroup = datadg
Rlink      : rlk_172_6037
info       : timeout=29 packet_size=1400
            latency_high_mark=10000 latency_low_mark=9950
            bandwidth_limit=none
state      : state=ACTIVE
            synchronous=off latencyprot=off srlprot=off
assoc      : rvg=CSM_RVG
            remote_host=172.25.84.34
            remote_dg=datadg
            remote_rlink=rlk_172_32481
            local_host=172.25.84.33
protocol   : UDP/IP
flags      : write attached consistent connected

```

Step 6 Verify that communications to the secondary cluster has been restored.

Example:

```

C:\> hastatus -sum
-- SYSTEM STATE
-- System                State                Frozen
A <PrimaryServer>       RUNNING                0
-- GROUP STATE
-- Group                System                Probed    AutoDisabled    State
B APP                   <PrimaryServer>       Y         N                ONLINE
B APPrep                <PrimaryServer>       Y         N                ONLINE
B ClusterService       <PrimaryServer>       Y         N                ONLINE
-- WAN HEARTBEAT STATE
-- Heartbeat            To                    State
L Icmp                  csm_secondary         ALIVE
-- REMOTE CLUSTER STATE
-- Cluster              State
M csm_secondary         RUNNING
-- REMOTE SYSTEM STATE

```

```

-- cluster:system                State                Frozen
N csm_secondary:<SecondaryServer> RUNNING            0
-- REMOTE GROUP STATE
-- Group                cluster:system        Probed        AutoDisabled    State
O APP csm_secondary:<SecondaryServer> Y              N              OFFLINE

```

Step 7 If replication has not recovered you may need to manually clear the IP resource if it has faulted and then start the APPrep service group on the secondary as follows:

Example:

```

C:\> hares -clear APP_IP
C:\> hagr -online APPrep -sys secondary_server_name

```

Network Ethernet Failure on Primary Server, Dual Cluster

Test Case Title: A failure occurs in the network Ethernet connection on the primary server.

Description: This test case verifies that VCS can detect a failure on the primary server network Ethernet port and can recover by starting the application on the secondary server. After the Ethernet connection is restored, you can manually fail over back to the original primary server, retaining any data changes that were made while running on the secondary.

Test Setup: A dual cluster configuration ([Figure](#)) with a single node in each cluster.

Step 1 Verify that the APP service group is running on the primary cluster.

Step 2 Remove the network Ethernet cable from the port on the server in the primary cluster to isolate the server from communicating with the switch/router network. VCS should detect this as a failure of the IP and NIC resources. Verify that VCS detected the failure and brought down the APP service group.

Example:

```

C:\> hastatus -sum
-- SYSTEM STATE
-- System                State                Frozen
A <PrimaryServer>        RUNNING            0
-- GROUP STATE
-- Group                System                Probed        AutoDisabled    State
B APP                    <PrimaryServer>      Y              N              OFFLINE
B APPrep                <PrimaryServer>      Y              N              OFFLINE|FAULTED
B ClusterService        <PrimaryServer>      Y              N              ONLINE
-- RESOURCES FAILED
-- Group                Type                Resource                System
C APPrep                IP                  APP_IP                <PrimaryServer>
C APPrep                NIC                NIC                  <PrimaryServer>
-- WAN HEARTBEAT STATE
-- Heartbeat            To                State
L Icmp                  csm_secondary      DOWN
-- REMOTE CLUSTER STATE
-- Cluster                State
M csm_secondary        FAULTED
-- REMOTE SYSTEM STATE
-- cluster:system        State                Frozen
N csm_secondary:<SecondaryServer> FAULTED            0
-- REMOTE GROUP STATE
-- Group                cluster:system        Probed        AutoDisabled    State
O APP                    csm_secondary:<SecondaryServer> Y              N              OFFLINE

```

Step 3 Start the APP service group on the secondary cluster using the following command on the secondary server:

Example:

```
C:\> hagrps -online -force APP -sys secondary_server_name
```

Step 4 From your client machine, log in to Security Manager to verify that it is operational. Change some data so that you can verify that changes are retained when you switch back to the primary server.

Step 5 Reconnect the network Ethernet cable to the primary cluster server.

Step 6 Clear any faults on the IP resource and turn on the APPrep service from the primary server:

Example:

```
C:\> hares -clear APP_IP
C:\> hagrps -online APPrep -sys primary_server_name
```

Step 7 Convert the original primary RVG to secondary and synchronize the data volumes in the original primary RVG with the data volumes on the new primary RVG using the fast failback feature. Using the Cluster Explorer for the secondary cluster, right-click the RVGPrimary resource (APP_RVGPrimary), select actions, then select fbsync from the Actions dialog box, and then click OK. Alternatively, you can issue the following command:

Example:

```
C:\> hares -action APP_RVGPrimary fbsync 0 -sys secondary_server_name
```

Step 8 Using the VCS Cluster Explorer on the secondary cluster, select the APP service group. From the short-cut menu, select Switch To, then Remote Switch(...), to open the Switch global dialog box. In the dialog box, specify the primary cluster and the primary server. Alternatively, issue the following command:

Example:

```
C:\> hagrps -switch APP -any -clus primarycluster
```

Step 9 Log in to the application to verify that the changes you made on the secondary server were retained.

Cluster Communication Failure

Test Case Title: Failures occur in the Ethernet used for cluster communication.

Description: The dedicated Ethernet connections used between servers in the cluster for intracluster communication fail. The test verifies that the cluster communications continue to function when up to two of the three redundant communication paths are lost.

Test Setup: A dual-node cluster ([Figure](#)) in a single cluster configuration, with two dedicated cluster communication Ethernet connections and a low-priority cluster communication connection configured on the network Ethernet connection.



Note In addition to the commands given in this test case, you can monitor the status of the cluster communications from the Cluster Explorer by selecting the root node in the tree and selecting the System Connectivity tab.

Step 1 Issue the following command to verify that all systems are communicating through GAB.

Note Group Membership Services/Atomic Broadcast (GAB) is a VCS protocol responsible for cluster membership and cluster communications.

Example:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen e8cc02 membership 01
Port h gen e8cc01 membership 01
```

Step 2 Remove the Ethernet cable from the first dedicated Ethernet port used for cluster communication on the primary server.

Step 3 Issue the following command to view the detailed status of the links used for cluster communication and verify that the first dedicated cluster communication port is down.

Note The asterisk (*) in the output indicates the server on which the command is run. The server where the command is run always shows its links up, even if one or more of those ports are the ones that are physically disconnected.

Example:

```
# llstat -nvv
LLT node information:
  Node          State   Link  Status  Address
  * 0 <PrimaryServer> OPEN
                        Adapter0  UP    00:14:5E:28:52:9C
                        Adapter1  UP    00:14:5E:28:52:9D
                        Adapter2  UP    00:0E:0C:9C:20:FE
  1 <SecondaryServer> OPEN
                        Adapter0  DOWN
                        Adapter1  UP    00:14:5E:28:27:17
                        Adapter2  UP    00:0E:0C:9C:21:C2
...

```

Step 4 If you configured a low-priority heartbeat link on the network interface, remove the Ethernet cable from the second dedicated Ethernet port used for cluster communication on the primary server.

Step 5 Issue the following command to verify that all systems are communicating through GAB. Also confirm that both servers in the cluster are now in a Jeopardy state, since each server has only one heartbeat working.

Example:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen e8cc02 membership 01
Port a gen e8cc02 jeopardy ;1
Port h gen e8cc01 membership 01
Port h gen e8cc01 jeopardy ;1
```

Step 6 Issue the following command to view the detailed status of the links used for cluster communication and verify that the second dedicated Ethernet port for cluster communications on the primary server is down.

Example:

```
# llstat -nvv
LLT node information:
```

```

Node           State   Link  Status  Address
* 0 <PrimaryServer>  OPEN
                Adapter0  UP    00:14:5E:28:52:9C
                Adapter1  UP    00:14:5E:28:52:9D
                Adapter2  UP    00:0E:0C:9C:20:FE
1 <SecondaryServer> OPEN
                Adapter0  DOWN
                Adapter1  UP    00:14:5E:28:27:17
                Adapter2  DOWN

```

Step 7 Replace the Ethernet cable on the second dedicated Ethernet port for cluster communications on the primary server.

Step 8 Verify that the Jeopardy condition was removed by issuing the following command:

Example:

```

# gabconfig -a
GAB Port Memberships
=====
Port a gen  e8cc02 membership 01
Port h gen  e8cc01 membership 01

```

Step 9 Replace the Ethernet cable on the first dedicated Ethernet port for cluster communications on the primary server.

Server Failures

This section covers causing server failures by removing the power from the server to cause a failure. Four cases are covered:

- [Standby Server Failure, Single Cluster, on page 90](#)
- [Primary Server Failure, Single Cluster, on page 91](#)
- [Standby Server Failure, Dual Cluster, on page 92](#)
- [Primary Server Failure, Dual Cluster, on page 94](#)

Standby Server Failure, Single Cluster

Test Case Title: The standby server in a single cluster configuration fails.

Description: This test case verifies that the application running in the primary server is unaffected and that after the standby server is repaired, the application can successfully rejoin the cluster configuration.

Test Setup: A dual node cluster ([Figure 4: Ethernet and Storage Connections for a Dual-Node Site, on page 15](#)) with two dedicated cluster communication Ethernet connections and a low-priority cluster communication connection on the network Ethernet connection.

Step 1 Verify that the application is running on the primary server in the cluster.

Example:

```

C:\> hastatus -sum
-- SYSTEM STATE

```



```

-- System                State                Frozen
A <PrimaryServer>       RUNNING                0
A <SecondaryServer>     RUNNING                0
-- GROUP STATE
-- Group                System                Probed    AutoDisabled    State
B APP                   <PrimaryServer>       Y         N                ONLINE
B APP                   <SecondaryServer>     Y         N                OFFLINE

```

Step 2 Remove the power for the secondary server and verify that VCS detected the failure and that the application continues to operate on the primary server.

Example:

```

C:\> hastatus -sum
-- SYSTEM STATE
-- System                State                Frozen
A <PrimaryServer>       RUNNING                0
A <SecondaryServer>     FAULTED                0
-- GROUP STATE
-- Group                System                Probed    AutoDisabled    State
B APP                   <PrimaryServer>       Y         N                ONLINE

```

Step 3 Reapply power and boot the secondary server. After the server recovers, verify that it rejoined the cluster in a healthy state by running the following command. The output should be identical to the output in Step 1.

Example:

```

C:\> hastatus -sum

```

Primary Server Failure, Single Cluster

Test Case Title: The primary server in a single cluster fails.

Description: This test case verifies that if a primary server fails, the application starts running on the secondary server and that after the primary server is restored, the application can be reestablished on the primary server.

Test Setup: A dual node cluster ([Figure](#)).

Step 1 Verify that the APP service group is running on the primary server in the cluster by examining the output of the following command:

Example:

```

C:\> hastatus -sum
-- SYSTEM STATE
-- System                State                Frozen
A <PrimaryServer>       RUNNING                0
A <SecondaryServer>     RUNNING                0
-- GROUP STATE
-- Group                System                Probed    AutoDisabled    State
B APP                   <PrimaryServer>       Y         N                ONLINE
B APP                   <SecondaryServer>     Y         N                OFFLINE

```

Step 2 Remove the power from the primary server and verify that VCS detected the failure and that the APP service group automatically moved to the secondary server.

Example:

```
C:\> hastatus -sum
-- SYSTEM STATE
-- System          State          Frozen
A <PrimaryServer>  FAULTED          0
A <SecondaryServer> RUNNING          0
-- GROUP STATE
-- Group           System          Probed    AutoDisabled  State
B APP              <SecondaryServer> Y         N             ONLINE
```

Step 3 Verify that you can successfully log in to Security Manager from a client machine.

Step 4 Restore the power to the primary server and verify that the server can rejoin the cluster in a healthy condition. Run the following command. The output should be identical to the output in Step 1.

Example:

```
C:\> hastatus -sum
```

Step 5 Manually switch the APP service group back to the primary server.

Example:

```
C:\> hagrpswitch APP -to primary_server_name
```

Standby Server Failure, Dual Cluster

Test Case Title: The standby server in a dual cluster configuration fails.

Description: This test case verifies that an application running in the primary cluster is unaffected by a standby server failure and that after the standby server is repaired, the application can successfully rejoin the dual cluster configuration.

Test Setup: A dual cluster configuration, with replication ([Figure](#)), with a single node in each cluster.

Step 1 Verify that the APP and ClusterService service groups are running in the primary cluster by running the following command on the primary server:

Example:

```
C:\> hastatus -sum
-- SYSTEM STATE
-- System          State          Frozen
A <PrimaryServer>  RUNNING          0
-- GROUP STATE
-- Group           System          Probed    AutoDisabled  State
B APP              <PrimaryServer> Y         N             ONLINE
B APPPrep         <PrimaryServer> Y         N             ONLINE
B ClusterService <PrimaryServer> Y         N             ONLINE
-- WAN HEARTBEAT STATE
-- Heartbeat      To              State
L Icmp            csm_secondary  ALIVE
-- REMOTE CLUSTER STATE
-- Cluster        State
M csm_secondary  RUNNING
```

```

-- REMOTE SYSTEM STATE
-- cluster:system          State          Frozen
N csm_secondary:<SecondaryServer> RUNNING 0
-- REMOTE GROUP STATE
-- Group                   cluster:system          Probed    AutoDisabled    State
O APP                      csm_secondary:<SecondaryServer> Y         N          OFFLINE

```

Step 2 Remove the power from the secondary server and verify that the primary cluster detects a loss of communication to the secondary cluster:

Example:

```

C:\> hastatus -sum
-- SYSTEM STATE
-- System                  State          Frozen
A <PrimaryServer>        RUNNING        0
-- GROUP STATE
-- Group                   System          Probed    AutoDisabled    State
B APP                    <PrimaryServer> Y         N          ONLINE
B APPrep                 <PrimaryServer> Y         N          ONLINE
B ClusterService        <PrimaryServer> Y         N          ONLINE
-- WAN HEARTBEAT STATE
-- Heartbeat              To              State
L Icmp                   csm_secondary  ALIVE
-- REMOTE CLUSTER STATE
-- Cluster                State
M csm_secondary         LOST_CONN
-- REMOTE SYSTEM STATE
-- cluster:system          State          Frozen
N csm_secondary:<SecondaryServer> RUNNING 0
-- REMOTE GROUP STATE
-- Group                   cluster:system          Probed    AutoDisabled    State
O APP                      csm_secondary:<SecondaryServer> Y         N          OFFLINE

```

Step 3 Restore the power to the secondary server. After the server restarts, verify that the primary cluster reestablished communications with the secondary cluster by running the following command. The output should be identical to the output in Step 1.

Example:

```
C:\> hastatus -sum
```

Step 4 Verify that the replication is operational and consistent by running the following command:

Example:

```

C:\> vxprint -Pl
Diskgroup = BasicGroup
Diskgroup = datadg
Rlink      : rlk_172_6037
info       : timeout=16 packet_size=1400
            latency_high_mark=10000 latency_low_mark=9950
            bandwidth_limit=none
state      : state=ACTIVE
            synchronous=off latencyprot=off srlprot=off
assoc     : rvg=CSM_RVG
            remote_host=172.25.84.34
            remote_dg=datadg
            remote_rlink=rlk_172_32481
            local_host=172.25.84.33

```

```
protocol : UDP/IP
flags : write attached consistent connected
```

Primary Server Failure, Dual Cluster

Test Case Title: The primary server in a dual cluster configuration fails.

Description: This test case verifies that if a primary server fails, the application starts running on the secondary server and that after the primary server is restored, the application can be reestablished on the primary server.

Test Setup: A dual cluster configuration, with replication (Figure), with a single node in each cluster.

Step 1 Verify that the APP and ClusterService service groups are running in the primary cluster by running the following command from the secondary server:

Example:

```
C:\> hastatus -sum
-- SYSTEM STATE
-- System          State          Frozen
A <SecondaryServer>  RUNNING      0
-- GROUP STATE
-- Group           System          Probed    AutoDisabled  State
B APP              <SecondaryServer> Y         N           OFFLINE
B APPPrep          <SecondaryServer> Y         N           ONLINE
B ClusterService  <SecondaryServer> Y         N           ONLINE
-- WAN HEARTBEAT STATE
-- Heartbeat       To              State
L Icmp             csm_primary    ALIVE
-- REMOTE CLUSTER STATE
-- Cluster         State
M csm_primary      RUNNING
-- REMOTE SYSTEM STATE
-- cluster:system  State          Frozen
N csm_primary:<PrimaryServer> RUNNING      0
-- REMOTE GROUP STATE
-- Group           cluster:system  Probed    AutoDisabled  State
O APP              csm_primary:<PrimaryServer> Y         N           ONLINE
```

Step 2 Remove the power from the primary server to cause a server failure. Verify that the secondary cluster reported a loss of connectivity to the primary cluster.

Example:

```
C:\> hastatus -sum
-- SYSTEM STATE
-- System          State          Frozen
A <SecondaryServer>  RUNNING      0
-- GROUP STATE
-- Group           System          Probed    AutoDisabled  State
B APP              <SecondaryServer> Y         N           OFFLINE
B APPPrep          <SecondaryServer> Y         N           ONLINE
B ClusterService  <SecondaryServer> Y         N           ONLINE
-- WAN HEARTBEAT STATE
-- Heartbeat       To              State
L Icmp             csm_primary    ALIVE
-- REMOTE CLUSTER STATE
```

```

-- Cluster          State
M csm_primary      LOST_CONN
-- REMOTE SYSTEM STATE
-- cluster:system          State          Frozen
N csm_primary:<PrimaryServer> RUNNING      0
-- REMOTE GROUP STATE
-- Group              cluster:system          Probed      AutoDisabled      State
O APP                  csm_primary:<PrimaryServer> Y              N              ONLINE

```

Step 3 Confirm that the state of the replication is disconnected. You can see this state from the flags parameter in the output of the following command:

Example:

```

C:\> vxprint -Pl
Diskgroup = BasicGroup
Diskgroup = datadg
Rlink      : rlk_172_32481
info       : timeout=500 packet_size=1400
            latency_high_mark=10000 latency_low_mark=9950
            bandwidth_limit=none
state      : state=ACTIVE
            synchronous=off latencyprot=off srlprot=off
assoc     : rvg=CSM_RVG
            remote_host=172.25.84.33
            remote_dg=datadg
            remote_rlink=rlk_172_6037
            local_host=172.25.84.34
protocol   : UDP/IP
flags     : write attached consistent disconnected

```

Step 4 Start the application on the secondary server by using the following command.

Example:

```

C:\> hagr -online -force APP -sys secondary_server_name

```

Step 5 Log in to the application and change some data so that you can verify later that changes made while the application operating on the secondary server can be retained when you revert to the primary server.

Step 6 Restore power to the primary server and allow the server to fully start up.

Step 7 Verify the status of the replication to show that the replication is connected; however, the two sides are not synchronized.

Example:

```

C:\> vxprint -Pl
Diskgroup = BasicGroup
Diskgroup = datadg
Rlink      : rlk_172_32481
info       : timeout=500 packet_size=1400
            latency_high_mark=10000 latency_low_mark=9950
            bandwidth_limit=none
state      : state=ACTIVE
            synchronous=off latencyprot=off srlprot=off
assoc     : rvg=CSM_RVG
            remote_host=172.25.84.33
            remote_dg=datadg
            remote_rlink=rlk_172_6037
            local_host=172.25.84.34
protocol   : UDP/IP
flags     : write attached consistent connected dcm_logging failback_logging

```

Step 8 Convert the original primary RVG to secondary and synchronize the data volumes in the original primary RVG with the data volumes on the new primary RVG using the fast failback feature. Using the Cluster Explorer for the secondary cluster, right-click the RVGPrimary resource (APP_RVGPrimary), select actions, then select fbsync from the Actions dialog box, and then click OK. Alternatively you can issue the following command:

Example:

```
C:\> hares -action APP_RVGPrimary fbsync 0 -sys secondary_server_name
```

Step 9 Verify that the current secondary (former primary) is synchronized with the current primary (former secondary) by looking for the keyword consistent in the flags parameter of the output of the following command:

Example:

```
C:\> vxprint -Pl
Diskgroup = BasicGroup
Diskgroup = datadg
Rlink      : rlk_172_32481
info       : timeout=29 packet_size=1400
            latency_high_mark=10000 latency_low_mark=9950
            bandwidth_limit=none
state      : state=ACTIVE
            synchronous=off latencyprot=off srlprot=off
assoc     : rvg=CSM_RVG
            remote_host=172.25.84.33
            remote_dg=datadg
            remote_rlink=rlk_172_6037
            local_host=172.25.84.34
protocol   : UDP/IP
flags      : write attached consistent connected
```

Step 10 Using the VCS Cluster Explorer on the secondary cluster, select the APP service group. From the shortcut menu, select Switch To, then Remote Switch(...) to open the Switch global dialog box. In the dialog box specify the primary cluster and the primary server. Alternately issue the following command, where primarycluster is the name of the primary cluster:

Example:

```
C:\> hagrps -switch APP -any -clus primarycluster
```

Step 11 Log in to the application to verify that the changes you made on the secondary server were retained.

Application Failures

This section covers test cases where the Security Manager application fails. Two cases are covered: a single cluster configuration and a dual cluster configuration. This section contains the following topics:

- [Application Failure, Single Cluster, on page 96](#)
- [Application Failure, Dual Cluster, on page 98](#)

Application Failure, Single Cluster

Test Case Title: The application fails on the primary server in a single cluster configuration.

Description: This test case verifies that VCS detects an application failure and that VCS automatically moves the application to the secondary server.

Test Setup: A dual node cluster (Figure) using the default application failover behavior.

Step 1 Verify that the APP service group is running on the primary server in the cluster by running the following command:

Example:

```
C:\> hastatus -sum
-- SYSTEM STATE
-- System          State          Frozen
A <PrimaryServer>  RUNNING      0
A <SecondaryServer> RUNNING      0
-- GROUP STATE
-- Group           System          Probed    AutoDisabled  State
B APP              <PrimaryServer> Y         N             ONLINE
B APP              <SecondaryServer> Y         N             OFFLINE
```

Step 2 On the server where Security Manager is running, stop the application by issuing the following command:

Example:

```
C:\> net stop crmdmgtd
```

Step 3 Verify that VCS detects that Security Manager failed on the primary server and starts the application on the secondary server.

Example:

```
# hastatus -sum
-- SYSTEM STATE
-- System          State          Frozen
A <PrimaryServer>  RUNNING      0
A <SecondaryServer> RUNNING      0
-- GROUP STATE
-- Group           System          Probed    AutoDisabled  State
B APP              <PrimaryServer> Y         N             OFFLINE|FAULTED
B APP              <SecondaryServer> Y         N             ONLINE
-- RESOURCES FAILED
-- Group           Type            Resource   System
C APP              CSManager      APP_CSManager <PrimaryServer>
```

Step 4 Manually clear the fault on the APP service group.

Example:

```
C:\> hagrps -clear APP -sys primary_server_name
```

Step 5 Manually switch the APP service group back to the primary server.

Example:

```
C:\> hagrps -switch APP -to primary_server_name
```

Application Failure, Dual Cluster

Test Case Title: The application fails on the primary server in a dual cluster configuration.

Description: This test case verifies that VCS detects an application failure.

Test Setup: A dual cluster configuration, with replication (Figure), with a single node in each cluster. Likewise, the assumption is that the default application failover behavior has not been modified (that is, failover between clusters requires manual intervention).

Step 1 Verify that the APP and ClusterService service groups are running in the primary cluster by running the following command from the primary server:

Example:

```
C:\> hastatus -sum
-- SYSTEM STATE
-- System                State                Frozen
A <SecondaryServer>     RUNNING             0
-- GROUP STATE
-- Group                System                Probed    AutoDisabled    State
B APP                  <SecondaryServer>    Y         N                OFFLINE
B APPrep              <SecondaryServer>    Y         N                ONLINE
B ClusterService     <SecondaryServer>    Y         N                ONLINE
-- WAN HEARTBEAT STATE
-- Heartbeat            To                    State
L Icmp                csm_primary          ALIVE
-- REMOTE CLUSTER STATE
-- Cluster              State
M csm_primary         RUNNING
-- REMOTE SYSTEM STATE
-- cluster:system      State                Frozen
N csm_primary:<PrimaryServer> RUNNING             0
-- REMOTE GROUP STATE
-- Group                cluster:system        Probed    AutoDisabled    State
O APP                  csm_primary:<PrimaryServer> Y         N                ONLINE
```

Step 2 On the server where Security Manager is running, stop the application by issuing the following command:

Example:

```
C:\> net stop crmdmgt
```

Step 3 Verify that VCS detects that the application failed and stops the APP service group. Issue the following command and observe the output.

Example:

```
# hastatus -sum
-- SYSTEM STATE
-- System                State                Frozen
A <PrimaryServer>      RUNNING             0
-- GROUP STATE
-- Group                System                Probed    AutoDisabled    State
B APP                  <PrimaryServer>    Y         N                OFFLINE | FAULTED
B APPrep              <PrimaryServer>    Y         N                ONLINE
B ClusterService     <PrimaryServer>    Y         N                ONLINE
-- RESOURCES FAILED
-- Group                Type                Resource                System
C APP                  CSManager          APP_CSManager          <PrimaryServer>
```



```
-- WAN HEARTBEAT STATE
-- Heartbeat      To          State
L  Icmp           csm_secondary  ALIVE
-- REMOTE CLUSTER STATE
-- Cluster        State
M  csm_secondary  RUNNING
-- REMOTE SYSTEM STATE
-- cluster:system State          Frozen
N  csm_secondary:<SecondaryServer> RUNNING  0
-- REMOTE GROUP STATE
-- Group          cluster:system Probed   AutoDisabled State
O  APP           csm_secondary:<SecondaryServer> Y       N       OFFLINE
```

Step 4 Manually clear the fault on the APP service group.

Example:

```
C:\> hagr -clear APP
```

Step 5 Put the APP service group online on the primary server to restart the application.

Example:

```
C:\> hagr -online APP -sys primary_server_name
```
