

Administrator Guide

VMware vCenter Server Heartbeat 6.3 Update 1

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see <http://www.vmware.com/support/pubs>.

EN-000562-01

vmware[®]

You can find the most up-to-date technical documentation on the VMware Web site at:

<http://www.vmware.com/support/>

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

Copyright © 2010 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

About This Book 7

Getting Started

- 1 Introduction 11
 - Overview 11
 - vCenter Server Heartbeat Concepts 11
 - Architecture 11
 - Protection Levels 13
 - Communications 16
 - vCenter Server Heartbeat Switchover and Failover Processes 17
- 2 Configuring vCenter Server Heartbeat 21
 - Server Configuration Wizard 22
 - Configuring the Machine Identity 22
 - Configuring the Server Role 23
 - Configuring the Client Connection Port 23
 - Configuring Channel Routing 23
 - Configuring the Default Channel Port 24
 - Configuring Low Bandwidth Module 24
 - Configuring Public IP Addressing 24
 - Configuring for Identical Nodes 25
 - Configuring for Non-Identical Nodes 25
 - Managing vCenter Server Heartbeat License Keys 26
 - Configuring Message Queue Logs 27
 - Configuring the Maximum Disk Usage 27

System Administration and Management

- 3 Server Protection 31
 - Server Protection Overview 31
 - Checking the Server Pair Status 31
 - Monitoring the Status of Servers 33
 - Configuring Heartbeat Settings 33
 - Configure Pings 33
 - Configure Failover 33
 - Configuring Response Times 34
 - Common Administrative Tasks in vCenter Server Heartbeat 34
 - Forcing a Switchover 34
 - Recovering From a Failover 35
 - Configuring Split-Brain Avoidance 36
- 4 Network Protection 39
 - Communication Status 39
 - Reviewing the VMware Channel Status 39

Configuring Public Network Connection Checks	39
Setting Max Server Time Difference	40

5 Application Protection 41

Application Protection Overview	41
Applications Tab	41
Reset the Application Health Status	41
View Application Status	42
Editing Individual Applications	42
Remove an Application	42
Manually Start and Stop Applications	42
Configuring Applications	42
Application Maintenance Mode	42
Reviewing the State of an Application	43
Reviewing the Applications Log	43
Filtering Application Log Entries	43
Services Tab	43
Adding a Service	43
Editing a Service	44
Checking the Status of Services	44
Unprotecting User Defined Services and Stopping Monitoring	44
Change the Order of Services	44
Removing a Service	44
Applications: Tasks	45
Adding a Task	45
Editing a Task	45
Remove a Task	45
Change the Order of Tasks	46
Starting a Task Manually	46
View, Add, and Remove User Accounts	46
Applications: Plugins	46
Install a Plug-In	46
Editing a Plug-in	46
Uninstalling a Plug-in	47

6 Status and Control 49

vCenter Server Heartbeat Console	49
About vCenter Server Heartbeat Console	49
Navigate vCenter Server Heartbeat Console	50
Change the Font for vCenter Server Heartbeat Console	50
Work with Groups and Pairs	50
Add or Remove a vCenter Server Group	50
Remove a vCenter Server Heartbeat Group	51
Add, Edit, Move, and Remove Pairs in VCenter Server Heartbeat Groups	51
Add a New Connection	51
Edit a Connection	51
Move a Connection	52
Remove a Connection	52
Edit Username and Password Settings	52
Review the Status of vCenter Server Heartbeat Groups and Pairs	52
Exit vCenter Server Heartbeat Console	53
Shut Down Windows Without Stopping vCenter Server Heartbeat	53

7 Performance Protection 55

Applications: Rules	55
---------------------	----

	Rules	55
	Checking a Rule Condition	55
	Edit a Rule	55
	Rules Installed by vCenter Server Heartbeat Plug-Ins	55
8	Data Protection	57
	Data Protection Overview	58
	Replication	58
	Registry and File Synchronization Status	58
	Initiate a Full Registry Check	58
	Initiate a Full System Check	58
	Initiate File Synchronization Manually	59
	Initiate Verify and Synchronize Manually	59
	File Filters	59
	Determine Effective Filters	59
	Add a User-Defined Inclusion Filter	60
	Add a User-Defined Exclusion Filter	60
	Edit User Defined Inclusion/Exclusion Filters	60
	Remove User-Defined Filters	60
	Automatic Filter Discovery	60
9	Other Administrative Tasks	61
	Configure Alerts	61
	Configure Alert Reporting	61
	Test Alert Reporting	62
	Configure Event Log Files	62
	Configure Log File Email Recipients	62
	Review Event Logs	63
	Event Log Filters	63
10	Troubleshooting	65
	Troubleshooting Unexpected Behaviors	65
	Two Active Servers	65
	Symptoms	65
	Causes	66
	Resolution	66
	Two Passive Servers	67
	Symptom	67
	Causes	67
	Resolution	67
	Synchronization Failures	68
	Services Running on the Passive Server	68
	VMware Channel Incorrectly Configured	68
	Incorrect or Mismatched Disk Configuration	69
	Passive Server Has Less Available Space than Active Server	69
	Registry Status is Out of Sync	70
	Resource Issues	70
	Registry Security Issues	70
	Channel Drops	70
	Performance Issues	70
	Passive Server Does Not Meet Minimum Hardware Requirements	71
	Hardware or Driver Issues on VMware Channel NICs	71
	Firewall Connection	72
	Incorrect VMware Channel Configuration	72
	VMware vCenter Server Heartbeat Packet Filter Is Enabled on the Channel NIC(s)	73

- Subnet or Routing Issues 74
 - LAN Deployment 74
 - WAN Deployment 74
- MaxDiskUsage Errors 74
 - Send Queue 75
 - Receive Queue 75
- MaxDiskUsage Error Messages 75
 - [L9]Exceeded the Maximum Disk Usage (VCChannelExceededMaxDiskUsageException) 75
 - [L9]Exceeded the Maximum Disk Usage on the ACTIVE Server 76
 - [L9]Exceeded the Maximum Disk Usage on the PASSIVE Server 76
 - [L20]Out of Disk Space (VCChannelOutOfDiskSpaceException) 77
 - Application Slowdown 78
 - Poor Application Performance 78
 - Both Servers Can Accommodate the Initial Load but the Load Has Increased 78
 - One Server Can Provide Adequate Resource Support, but the Other Cannot 78
 - Scheduled Resource Intensive Tasks 79

- Glossary 81

About This Book

The *Administrator Guide* provides information configuring VMware vCenter Server Heartbeat network protection, application protection, data protection, Split-brain Avoidance, and more. To help you protect your VMware vCenter Server, the book provides an overview of protection offered by vCenter Server Heartbeat and the actions that vCenter Server Heartbeat can take in the event of a network, hardware, or application failure.

Intended Audience

This guide assumes the reader has working knowledge of networks including the configuration of TCP/IP protocols and domain administration on the Windows™ 2003 and 2008 platforms, notably in Active Directory and DNS.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation go to <http://www.vmware.com/support/pubs>.

Overview of Content

This guide is designed to give guidance on the configuration and administration of vCenter Server Heartbeat, and is organized into the following sections:

- Preface — *About This Book* (this chapter) provides an overview of this guide and the conventions used throughout.
- Chapter 1 — *Introduction* presents an overview of vCenter Server Heartbeat concepts including the Switchover and Failover processes.
- Chapter 2 — *Configuring vCenter Server Heartbeat* shows you how to use the Server Configuration Wizard to configure your new installation of vCenter Server Heartbeat.
- Chapter 3 — *Server Protection* gives an overview of how vCenter Server Heartbeat provides protection against server system crash or server hardware failure, shows you how to check the server pair status, and explains how to configure settings, shutdown options, and Split-Brain Avoidance.
- Chapter 4 — *Network Protection* describes how vCenter Server Heartbeat protects against network failure and provides a way to monitor communication status. It also explains how to configure public network connection checks and maximum server time difference.
- Chapter 5 — *Application Protection* discusses how vCenter Server Heartbeat maintains the protected application environment ensuring that applications and services stay alive on the network.
- Chapter 6 — *Status and Control* introduces you to the vCenter Server Heartbeat Console and shows you how to configure its look and feel.

- Chapter 7 — *Performance Protection* describes how vCenter Server Heartbeat monitors system and application attributes to prevent an unexpected system or application failure.
- Chapter 8 — *Data Protection* discusses how vCenter Server Heartbeat intercepts all data written by users and protected applications and maintains a copy of this data for use in case of failure.
- Chapter 9 — *Other Administrative Tasks* discusses additional tasks for the administrator to configure system logging and alerting functions.
- Chapter 10 — *Troubleshooting* provides techniques to troubleshoot common issues and unexpected behaviors.

Document Feedback

VMware welcomes your suggestions for improving our documentation and invites you to send your feedback to docfeedback@vmware.com.

Abbreviations Used in Figures

The figures in this book use the abbreviations listed in [Table 1](#).

Table 1. Abbreviations

Abbreviation	Description
Channel	VMware Channel
NIC	Network Interface Card
P2P	Physical to Physical
P2V	Physical to Virtual
V2V	Virtual to Virtual

Technical Support and Education Resources

The following sections describe the technical support resources available to you. To access the current version of this book and other books, go to www.vmware.com/support/pubs.

Online and Telephone Support

Go to www.vmware.com/support to use online support to submit technical support requests, view your product and contract information, and register your products.

Go to www.vmware.com/support/phone_support.html to find out how to use telephone support for the fastest response on priority 1 issues (applies to customers with appropriate support contracts).

Support Offerings

Go to www.vmware.com/support/services to find out how VMware support offerings can help meet your business needs.

VMware Professional Services

Go to www.vmware.com/services to access information about education classes, certification programs, and consulting services. VMware Education Services courses offer extensive hands-on labs, case study examples, and course materials designed for use as on-the-job reference tools. Courses are available onsite, in the classroom, and live online. For onsite pilot programs and implementation best practices, VMware Consulting Services provides offerings to help you assess, plan, build, and manage your virtual environment.

Getting Started

Introduction

This chapter includes the following topics:

- [“vCenter Server Heartbeat Concepts”](#) on page 11
- [“vCenter Server Heartbeat Switchover and Failover Processes”](#) on page 17

Overview

vCenter Server Heartbeat is a Windows based service specifically designed to provide high availability protection for vCenter Server configurations without requiring any specialized hardware.

vCenter Server Heartbeat Concepts

Architecture

vCenter Server Heartbeat software is installed on a Primary (production) server and a Secondary (ready-standby) server. These names refer to the physical hardware (identity) of the servers.

Depending on the network environment, vCenter Server Heartbeat can be deployed in a Local Area Network (LAN) or Wide Area Network (WAN). Additionally, vCenter Server Heartbeat allows for installation using identical nodes (LAN and WAN) or non-identical nodes (LAN only). These features provide flexibility necessary to address most network environments.

Depending on the network environment and architecture selected, vCenter Server Heartbeat can be configured where the Primary and Secondary server have the same domain name and Principal (Public) IP address (identical nodes) or where the Primary and Secondary servers have different domain names but share the Principal (Public) IP address (non-identical nodes). In either architecture, the Secondary server in the will have the same file and data structure and can run all of the same applications and services as the Primary server. The vCenter Server Heartbeat software is symmetrical in almost all respects, and either the Primary server or the Secondary server can take the active role and provide protected applications to the user.

When deployed, vCenter Server Heartbeat uses two servers with either identical Principal (Public) IP addresses for identical nodes or a shared Principal (Public) IP address for non-identical nodes. One of the servers performs the role of the active server that is visible on the Public network while the other is the passive server that is hidden from the Public network but remains as a ready-standby server. Only one server can display the Principal (Public) IP address and be visible on the Public network at any given time.

Identical Nodes Configuration

When configured for identical nodes, vCenter Server Heartbeat uses two servers with the same domain name, Principal (Public) network IP address, file and data structure, applications, and services to ensure that vCenter Server is running and available. With a packet filter installed on both servers Principal (Public) NICs, vCenter Server Heartbeat applies the packet filter to the Principal (Public) NIC on the passive server thereby blocking communications and hiding the passive server from the Public network. When a switchover occurs, the packet

filter on the currently active server is enabled preventing communications on the Public network while the packet filter on the currently passive server is disabled thereby allowing the two servers to switch roles. From this point on, the previously passive server is now active and servicing the clients. The previously active server is now passive and hidden from the Public network.

In the event of a failover, the previously active server is deemed to have failed and no further actions are necessary. The packet filter on the previously passive server is disabled allowing the server to become visible on the Public network and service clients as the new active server.

Advantages of Identical Nodes Configuration

Identical node configuration provides for a simplified installation and supports Virtual to Virtual, Virtual to Physical, and Physical to Physical architectures. Additionally, identical node configuration supports use of both the Pre-clone and Install clone, installation techniques. When configured with identical nodes, vCenter Server Heartbeat provides two identical servers that can function as the active server servicing clients. Additionally, with identical nodes, organizations can deploy vCenter Server Heartbeat in WAN environments thereby extending protection over wide geographic distances.

Non-Identical Nodes Configuration

Non-identical node configuration provides enhanced management capabilities to the vCenter Heartbeat Server cluster by providing continuous access to the passive server simultaneously while the active server continues to service clients. When configured with non-identical nodes, vCenter Server Heartbeat uses unique server names for each server and enhanced IP address management capabilities to accommodate Microsoft (updatable) DNS servers and Locked Down mode or non-Microsoft DNS servers while providing the same levels of protection as when configured with identical nodes.

The Secondary server has a different Fully Qualified Domain Name (FQDN) than the Primary server but uses the same file and data structure, same Principal (Public) network address, and can run all the same applications and services as the Primary server. With vCenter Server Heartbeat software symmetrical in almost all respects, either the Primary server or the Secondary server can assume the active role and provide protected applications to the user.

Additionally, this configuration allows the passive server to be easily accessed for maintenance purposes, updating anti-virus definition files, receiving operating system hot-fixes, updates and patches from third-party management software, and allows use of third-party monitoring tools.

Switchover/Failover Without Updating DNS Servers

Select the *Non-Identical* option when you want to deploy vCenter Server Heartbeat with non-identical nodes in an environment that does not permit updating DNS (a non-updatable Microsoft DNS server (or Locked Down Mode) or a non-Microsoft DNS server). This configuration requires vCenter Server vCenter Server 4.0 Update 1 or later, if SQL Server is deployed, it must be remote, and since DNS will not be updated automatically, you must prepopulate the DNS server with entries for the new management names and IP addresses of the Primary and Secondary servers. vCenter Server Heartbeat uses permanently assigned unique Management IP addresses on both the Primary and Secondary servers. Adjacent IP addresses should be reserved and used for the Principal (Public) IP address and the Management IP addresses for the Primary and Secondary Servers when installing vCenter Server Heartbeat on non-identical nodes on servers running Windows 2008. When vCenter Server Heartbeat is started, the shared Principal (Public) IP address is provided to the active server. When a switchover occurs, the shared Principal (Public) IP address is removed from the previously active server (making it passive) and provided to the previously passive server (making it active) allowing the servers to switch roles and the new active server to begin servicing clients.

Microsoft DNS Server Updated During Switchover/Failover

Select the *Non-identical (Updates MS DNS)* option when you want to deploy vCenter Server Heartbeat with non-identical nodes in an environment where vCenter Server 4.0 Update 1 or later is installed, if SQL Server is deployed it must be remote, a Microsoft DNS server that permits updates is used, and you want to take advantage of the enhanced passive server management capabilities. vCenter Server Heartbeat interacts with the DNS server and following a switchover, will ensure that servers are correctly identified on the network by updating DNS information. vCenter Server Heartbeat uses unique server names and Management IP

addresses for the Primary and Secondary servers separate from the shared Public (Principal) IP address. Clients connect to vCenter Server using a unique virtual service name configured in vCenter Server Heartbeat that resolves to the Public (Principal) IP address. When started, vCenter Server Heartbeat provides the passive server its unique Management IP address and provides the active server with the shared Principal (Public) IP address.

When a switchover occurs, vCenter Server Heartbeat removes the shared Principal (Public) IP address from the currently active server and replaces it with its unique Management IP address. On the currently passive server, vCenter Server Heartbeat removes the Management IP address and replaces it with the shared Principal (Public) IP address thereby allowing the servers to switch roles.

Advantages of Non-Identical Nodes Configuration

Deployment of non-identical nodes allows for easy access to the passive (hidden/ready standby) server at any time regardless of which server is active and servicing clients. Additionally, this configuration allows the passive server to be easily accessed for maintenance purposes, updating anti-virus definition files, receiving operating system hot-fixes, updates and patches from third-party management software, and allows use of third-party monitoring tools.

Protection Levels

vCenter Server Heartbeat provides the following protection levels:

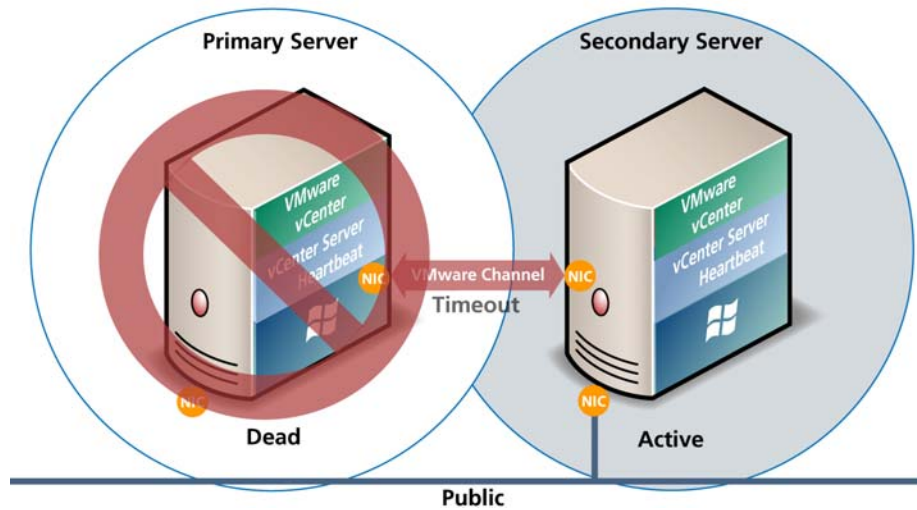
- **Server Protection** – vCenter Server Heartbeat provides continuous availability to end users through a hardware failure scenario or operating system crash. Additionally, vCenter Server Heartbeat protects the network identity of the production server, ensuring users are provided with a replica server on the failure of the production server.
- **Network Protection** – vCenter Server Heartbeat proactively monitors the network by polling up to three nodes to ensure that the active server is visible on the network.
- **Application Protection** – vCenter Server Heartbeat maintains the application environment ensuring that applications and services stay alive on the network.
- **Performance Protection** – vCenter Server Heartbeat proactively monitors system performance attributes to ensure that the system administrator is notified of problems and can take pre-emptive action to prevent an outage.
- **Data Protection** – vCenter Server Heartbeat intercepts all data written by users and applications, and maintains a copy of this data on the passive server that can be used in the event of a failure.

vCenter Server Heartbeat provides all five protection levels continuously, ensuring all facets of the user environment are maintained at all times, and that the network (Principal (Public) network) continues to operate through as many failure scenarios as possible.

Server Protection

vCenter Server Heartbeat provides continuous availability to end users through a hardware failure scenario or operating system crash. Additionally, vCenter Server Heartbeat protects the network identity of the production server, ensuring users are provided with a replica server including server name (single identity only), IP address on the failure of the production server.

Two instances of vCenter Server Heartbeat regularly send “I’m alive” messages and message acknowledgments to one another over a network connection referred to as the VMware Channel to detect interruptions in responsiveness. If the passive server detects that this monitoring process (referred to as the heartbeat) has failed, it initiates a failover as illustrated in [Figure 1-1](#).

Figure 1-1. Failover

A failover is similar to a switchover but is used in more urgent situations, such as when the passive server detects that the active server is no longer responding. This can occur when the active server hardware fails, loses its network connections, or otherwise becomes unavailable. Rather than the active server gracefully closing, the passive server determines that the active server has failed and requires no further operations. In a failover, the passive server immediately assumes the active server role. The failover process is discussed later in this guide.

Network Protection

vCenter Server Heartbeat proactively monitors the network by polling up to three nodes to ensure that the active server is visible on the network. vCenter Server Heartbeat polls defined nodes around the network, including the default gateway, the primary DNS server, and the global catalog server at regular intervals. If all three nodes fail to respond, for example, in the case of a network card failure or a local switch failure, vCenter Server Heartbeat can initiate a switchover, allowing the Secondary server to assume an identical network identity as the Primary server.


Application Protection

vCenter Server Heartbeat running on the active server locally monitors the applications and services it has been configured to protect (through the use of plug-ins) to verify that protected applications are operational and not in an unresponsive or stopped state. This level of monitoring is fundamental in ensuring that applications remain available to users.

If a protected application fails, vCenter Server Heartbeat first tries to restart the application on the active server (1) in [Figure 1-2](#).

If the application does not successfully restart, vCenter Server Heartbeat initiates a switchover (2) in [Figure 1-2](#). Refer to “[vCenter Server Heartbeat Switchover and Failover Processes](#)” on page 17 for further information about the switchover process.

Figure 1-2. Switchover



A switchover gracefully closes any protected applications that are running on the active server and restarts them on the passive server, including the application or service that caused the failure. In the example where the Primary server is active and the Secondary server is passive, the Primary server is demoted to a passive role and is hidden from the network when the Secondary server is promoted to an active role and is made visible to the network. The mechanics of switchovers are discussed in more detail later in this guide.

Performance Protection

Ensuring that your protected applications are operational and providing service at a level of performance adequate for users to remain productive is important. The vCenter Server Heartbeat plug-in provides these monitoring and pre-emptive repair capabilities.

vCenter Server Heartbeat proactively monitors system performance attributes to ensure that the system administrator is notified of problems and can take pre-emptive action to prevent an outage.

In addition to monitoring application services, vCenter Server Heartbeat can monitor specific application attributes to ensure that they remain within normal operating ranges. Similar to application monitoring, various rules can be configured to trigger specific corrective actions whenever these attributes fall outside of their respective ranges.

vCenter Server Heartbeat provides the same level of flexibility to define and perform multiple corrective actions in the event of problems on a service by service or even attribute by attribute basis.

Data Protection

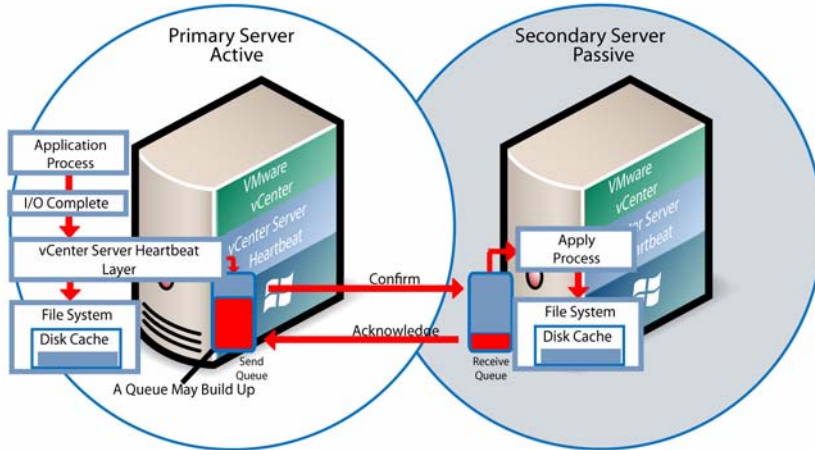
You can configure vCenter Server Heartbeat to protect the application environment. All data files that users or the applications require in the application environment are made available should a failure occur. After installation, vCenter Server Heartbeat configures itself to protect files, folders, and registry settings for vCenter Server on the active server by mirroring them in real time to the passive server. If a failover occurs, all files protected on the failed server are available to users after the failover, hosted on the Secondary server.

vCenter Server Heartbeat intercepts all file system I/O operations on the active server. If the intercepted write and update operations are within the protected set, these are placed in a queue on the active server referred to as the active server's send queue, pending transmission to the passive server. Each request is numbered to maintain its order in the queue.

With the request in the active server's send queue, vCenter Server Heartbeat allows the disk I/O to continue with the requested disk operation.

If the channel is connected, the active server's send queue is transferred to the passive server, which places all the requests in the passive server's receive queue. The passive server confirms the changes were logged by sending the active server an acknowledgement. The active server clears the data from its queue.

Figure 1-3. Apply Process



The apply process running on the passive server’s receive queue applies all updates in strict sequence, duplicating an identical set of file operations on the passive server as illustrated in [Figure 1-3](#).

Communications

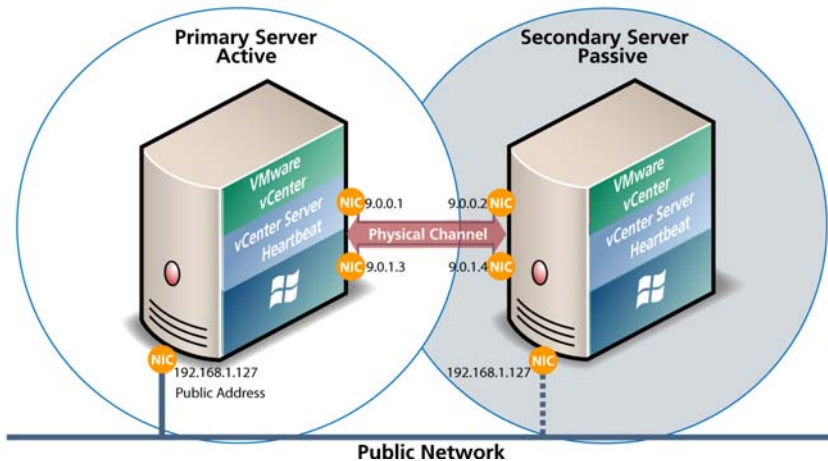
The VMware Channel is a crucial component of the setup and can be configured in a number of ways.

Both the Primary and Secondary servers must have two or more network interface connections (NICs).

The Principal (Public) network requires one NIC. The VMware Channel uses a separate NIC for the private connection between the servers used for control and data transfer between the pair of servers.

A second pair of NICs can be used to provide a degree of redundancy for the VMware Channel. In this configuration, the VMware Channel has a dual channel if more than one dedicated NIC is provided for the VMware Channel on each server. To provide added resilience, the communications for the second channel must be completely independent from the first channel. They must not share any switches, virtual switches, routers or the same WAN connection.

Figure 1-4. Communication Between Primary and Secondary Servers



The IP address a client uses to connect to the active server (the Principal (Public) IP address) must be configured as a static IP address, that is, not DHCP (Dynamic Host Configuration Protocol) enabled. In the figure above, the IP address is configured as 192.168.1.127.

NOTE Obtain the IP address: type **ipconfig** at the prompt in a DOS shell. For additional information about the IP configuration, add the switch **/All** to the **ipconfig** command.

When deployed with identical nodes, the Principal (Public) NICs on the passive server are configured to use the same IP address as that of the active server but are prevented from communicating with the live network through an IP packet filtering system installed with vCenter Server Heartbeat. This packet filter prevents traffic using the Principal (Public) address from being committed to the wire. It also prevents NetBIOS traffic utilizing other IP addresses on the NIC from being sent to prevent NetBIOS name resolution conflicts.

The NICs on the active and passive servers used for the VMware Channel are configured so that their IP addresses are outside of the subnet range of the Principal (Public) network. These addresses are referred to as VMware Channel addresses.

During installation, setup will switch off NetBIOS for the VMware Channel(s) on the active and passive servers as this connection remains live and both the passive and active machines have the same NetBIOS name. Following restore and after the vCenter Server Heartbeat installation completes (runtime), NetBIOS is disabled across the channel(s). This occurs during installation to prevent a name conflict, which occurs when both servers have the same name.

The NICs that support connectivity across the VMware Channel can be standard 100BaseT Ethernet cards providing a throughput of 100 Mbits per second across standard Cat-5 cabling. In its most basic form, a dedicated channel requires no hubs or routers, but the direct connection requires crossover cabling.

When configured for a WAN deployment, configure the VMware Channel to use static routes over switches and routers to maintain continuous communications independent from corporate or public traffic.

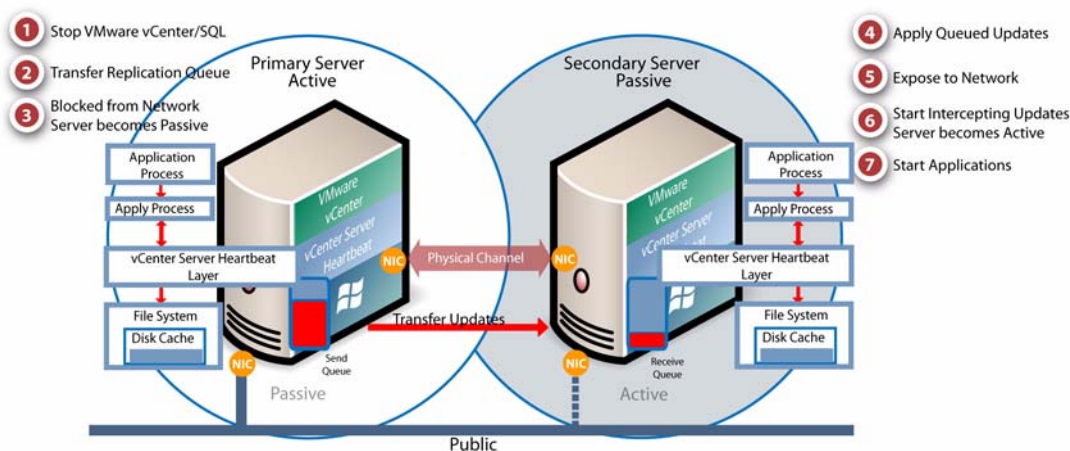
vCenter Server Heartbeat Switchover and Failover Processes

vCenter Server Heartbeat uses four different procedures — managed switchover, automatic switchover, automatic failover, and managed failover — to change the role of the active and passive servers depending on the status of the active server.

Managed Switchover

You can click **Make Active** on the **vCenter Server Heartbeat Console Server: Summary** page to manually initiate a managed switchover. When a managed switchover is triggered, the running of protected applications is transferred from the active machine to the passive machine in the server pair. The server roles are reversed.

Figure 1-5. Switchover



A managed switchover performs the following steps:

- 1 Stop the protected applications on the active server. After the protected applications stop, no more disk updates are generated.
- 2 Send all updates that are still queued on the active server to the passive server. After this step, all updates are available on the passive server.

- 3 Re-designate the Secondary server as the new active server. After this step, vCenter Server Heartbeat:
 - Hides the previously active server from the network.
 - Makes the newly active server visible on the network. The newly active server has the same identity as the previously active server, and begins to intercept and queue disk I/O operations for the newly passive server.
- 4 vCenter Server Heartbeat causes the newly passive server to begin accepting updates from the active server.
- 5 vCenter Server Heartbeat starts the same protected applications on the new active server. The protected applications become accessible to users. The managed switchover is complete

Automatic Switchover

Automatic switchover (auto-switchover) is similar to failover (discussed in the next section) but is triggered automatically when system monitoring detects failure of a protected application.

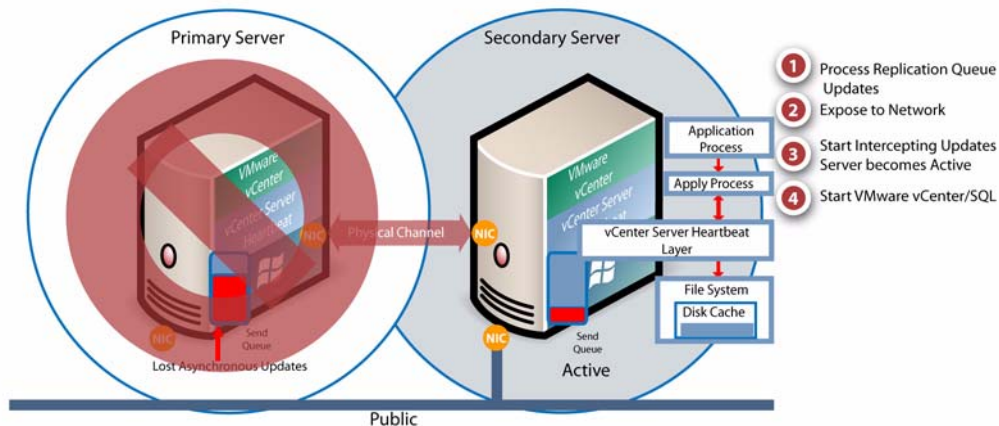
Like managed switchover, auto-switchover changes the server roles but then stops vCenter Server Heartbeat on the previously active server to allow the administrator to investigate the cause of the auto-switchover and verify the integrity of the data.

After the cause for the auto-switchover is determined and corrected, the administrator can use vCenter Server Heartbeat Console to return the server roles to their original state.

Automatic Failover

Automatic failover is similar to automatic switchover (discussed above) but is triggered when the passive server detects that the active server is no longer running properly and assumes the role of the active server.

Figure 1-6. Failover



During the automatic failover, the passive server performs the following steps:

- 1 Apply any intercepted updates currently in the passive server's receive queue as identified by the log of update records that are saved on the passive server but not yet applied to the replicated files.

The amount of data in the passive server's receive queue affects the time required to complete the failover process. If the passive server's receive queue is long, the system must wait for all updates to the passive server to complete before the rest of the process can take place. An update record can be applied only if all earlier update records are applied, and the completion status for the update is in the passive server's receive queue. When no more update records can be applied, any update records that cannot be applied are discarded.

- 2 Switch mode of operation from passive to active.

This enables the public identity of the server. The active and passive servers both use the same system name and same Principal (Public) IP address. This Principal (Public) IP address can be enabled only on one system at anytime. When the public identity is enabled, any clients previously connected to the server before the automatic failover are able to reconnect.

- 3 Start intercepting updates to protected data. Any updates to the protected data are saved in the send queue on the local server.
- 4 Start all protected applications. The applications use the replicated application data to recover, and then accept re-connections from any clients. Any updates that the applications make to the protected data are intercepted and logged.

At this point, the originally active server is offline and the originally passive server is filling the active role and is running the protected applications. Any updates that completed before the failover are retained. Application clients can reconnect to the application and continue running as before.

Managed Failover

Managed failover is similar to automatic failover in that the passive server automatically determines that the active server has failed and can warn the system administrator about the failure; but no failover actually occurs until the system administrator manually triggers this operation.

Automatic Switchover and Failover in a WAN Environment

Automatic switchover and failover in a WAN environment differ from a automatic switchover and failover in a LAN environment due to the nature of the WAN connection. In a WAN environment, automatic switchover and failover are disabled by default in the event that the WAN connection is lost.

Should a condition arise that would normally trigger an automatic switchover or failover, the administrator will receive vCenter Server Heartbeat alerts. The administrator must manually click the **Make Active** button on the **Server: Summary** page of the vCenter Server Heartbeat Console to allow the roles of the servers to switch over the WAN.

To enable Automatic Switchover in a WAN

- 1 In the vCenter Server Heartbeat Console, click the **Network** tab to display the **Network Monitoring** page.
- 2 Click **Configure Auto-switchover**.
- 3 Select the **Auto-switchover if client network connectivity lost for** check box.
- 4 Configure the number of pings to wait before performing the auto-switchover.
- 5 Click **OK**.

Recovery from a Failover

Assuming that the Primary server was active and the Secondary server was passive before the failover, after the failover the Secondary server is active and the Primary server is passive.

NOTE When failover conditions such as a power failure cause both active and passive servers to fail, the condition may result that causes both servers to restart in passive mode. In this situation, manual intervention is required.

After rectifying the problem that initiated the failover, it is a simple process to reinstate the Primary server as the active server and the Secondary server as the passive server.

When vCenter Server Heartbeat starts on the failed Primary server, it detects that it did not stop cleanly the previous time. It disables the public identity by deploying the IP packet filter and halts vCenter Server Heartbeat so that the issues that caused the failure can be resolved.

To restore the previously failed server to the active role

- 1 Correct the conditions that caused the failover.
- 2 Verify the integrity of the disk data on the failed server.
- 3 Restart the failed, now passive, server after all issues are resolved.
- 4 Start vCenter Server Heartbeat on the passive server.

At this point, the instances of vCenter Server Heartbeat running on the servers connect and begin to re-synchronize the data on the Primary server.

- 5 Wait until vCenter Server Heartbeat is fully synchronized.

When the re-synchronization is complete, you can continue operating with this configuration (for example, the Secondary server is the active server and the Primary server is the passive server), or initiate a managed switchover.

- 6 Optionally, perform a managed switchover to return the Primary and Secondary servers to the same roles they had before the failover.

Configuring vCenter Server Heartbeat

This chapter includes the following topics:

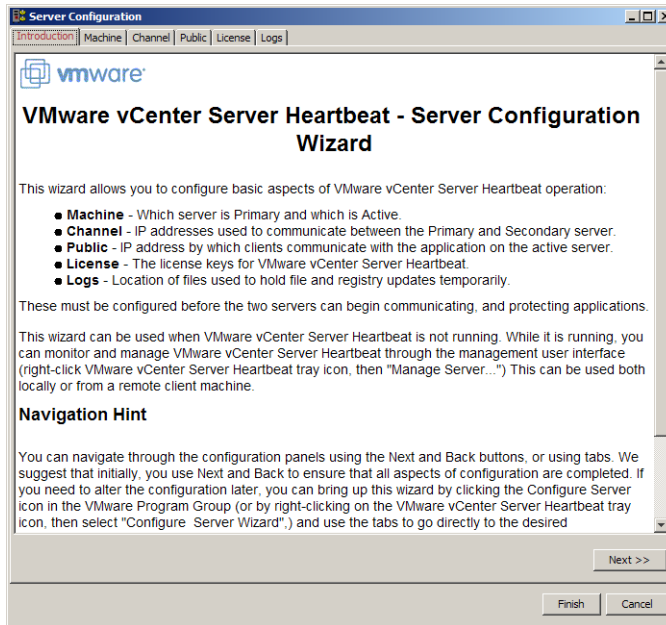
- [“Server Configuration Wizard”](#) on page 22
- [“Configuring the Machine Identity”](#) on page 22
- [“Configuring the Server Role”](#) on page 23
- [“Configuring the Client Connection Port”](#) on page 23
- [“Configuring Channel Routing”](#) on page 23
- [“Configuring the Default Channel Port”](#) on page 24
- [“Configuring Low Bandwidth Module”](#) on page 24
- [“Configuring Public IP Addressing”](#) on page 24
- [“Managing vCenter Server Heartbeat License Keys”](#) on page 26
- [“Configuring Message Queue Logs”](#) on page 27
- [“Configuring the Maximum Disk Usage”](#) on page 27

Server Configuration Wizard

The **VMware vCenter Server Heartbeat – Server Configuration Wizard** (Configure Server wizard) sets up and maintains communications between the vCenter Server Heartbeat servers. Configuration information includes the IP address for the VMware Channel(s) and Principal (Public) addresses on both servers.

After the system is set up and is functioning correctly, you do not need to reconfigure the system. The Configure Server wizard becomes redundant during daily operations of the software.

- 1 Before launching the Configure Server wizard to change the vCenter Server Heartbeat configuration, you must stop vCenter Server Heartbeat.
- 2 Click the Configure Server icon on the desktop or **Start > All Programs > VMware > VMware vCenter Server Heartbeat > Configure Server** to launch the Configure Server wizard.



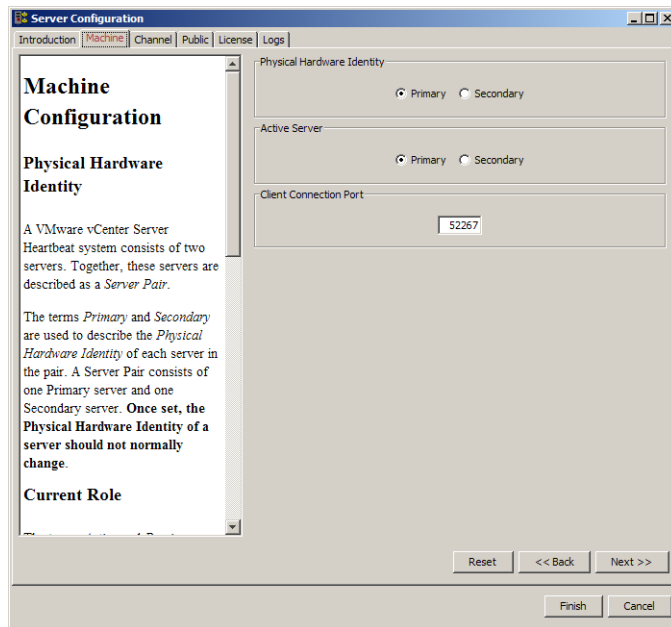
Configuring the Machine Identity



CAUTION The machine Identity should only be changed when directed to do so by VMware Support or when instructed to by a knowledge base article. vCenter Server Heartbeat is designed not to allow two passive or two active servers to connect.

To change the machine identity

- 1 Click the **Machine** tab and select a **Physical Hardware Identity** for the local machine.
- 2 Click **Next** or **Finish**.



Configuring the Server Role

To change the server role, click the **Machine** tab, select the **Current Role** of the local machine, and click **Next** or **Finish**.

NOTE Before changing the role of the local server, verify that the other (remote) server in the pair is not already performing the same role.

Configuring the Client Connection Port

Clients such as the vCenter Server Heartbeat Console use the Client Connection Port to connect to vCenter Server Heartbeat. Do not change this port unless another application is using it. To change the Client Connection Port, click the **Machine** tab, edit the default entry (52267) and click **Next** or **Finish**.

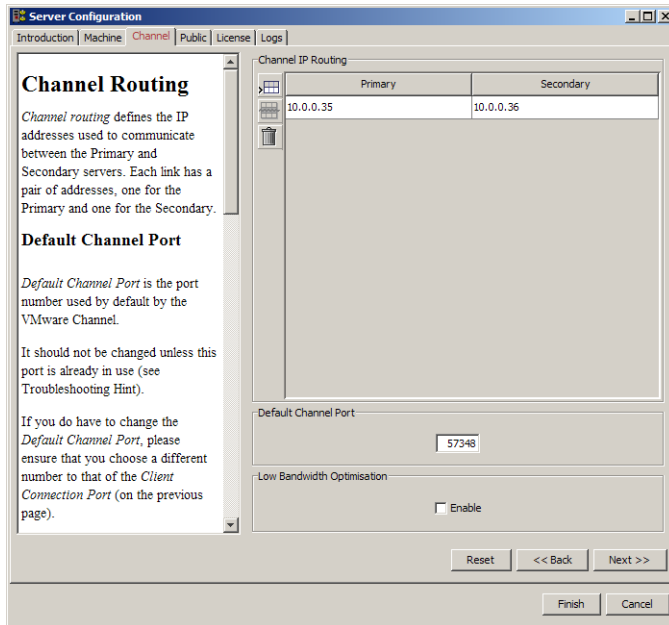
Configuring Channel Routing

Channel IP routing defines the IP addresses used to communicate between the Primary and Secondary servers. Each link has a pair of addresses, one for the Primary, and one for the Secondary.

To add an additional VMware Channel after installing the NICs and configuring them

- 1 Click the **Channel** tab. Click **Add Row** to add the new IP addresses for both the Primary and Secondary server to the VMware Channel IP Routing table.
- 2 Use the drop-down menu to view a list of available IP addresses on the local server.

- 3 Type the remote server IP address.



- 4 To change the VMware Channel IP addresses, select and edit the entry in the table.
- 5 Click **Next** or **Finish**.

Configuring the Default Channel Port

VMware Channel uses the Default Channel Port to communicate between the Primary and Secondary server. Do not change this port unless another application is using it. To change the Default Channel Port, click the **Channel** tab, edit the default entry (57348), and click **Next** or **Finish**.

Configuring Low Bandwidth Module

This feature is automatically enabled during installation when configured for a WAN. To disable this feature, click the **Channel** tab and clear the Low Bandwidth Optimization check box. When enabled, the VMware Channel optimizes communications for low bandwidth connections. Low Bandwidth Optimization (LBO) stores data on disk rather than in memory and is essential for WAN installations or when bandwidth is limited.

This setting should not be changed unless directed to do so by VMware Support.

NOTE This feature is designed for implementations where the available throughput on the VMware Channel is slower than 10 Mbit/s. Do not enable the Low Bandwidth Module in a LAN, this feature is not designed to work in a LAN where the throughput is much faster.

Configuring Public IP Addressing

vCenter Server Heartbeat servers are configured with one or more Principal (Public) IP addresses. These are the addresses used by clients of the protected application to connect to the application. Typically, there is one Principal (Public) IP address.

You must configure all of the Principal (Public) IP addresses on both Primary and Secondary servers. On the passive server, all traffic to and from these Principal (Public) IP addresses is blocked. On the active server, all traffic passes through. When the server roles switch, the IP filtering mode also switches, so client systems always connect to the Principal (Public) IP addresses on the server that is currently active. When the vCenter Server Heartbeat service shuts down, the filtering remains in place so no IP address conflicts occur between the two servers.

Configuring for Identical Nodes

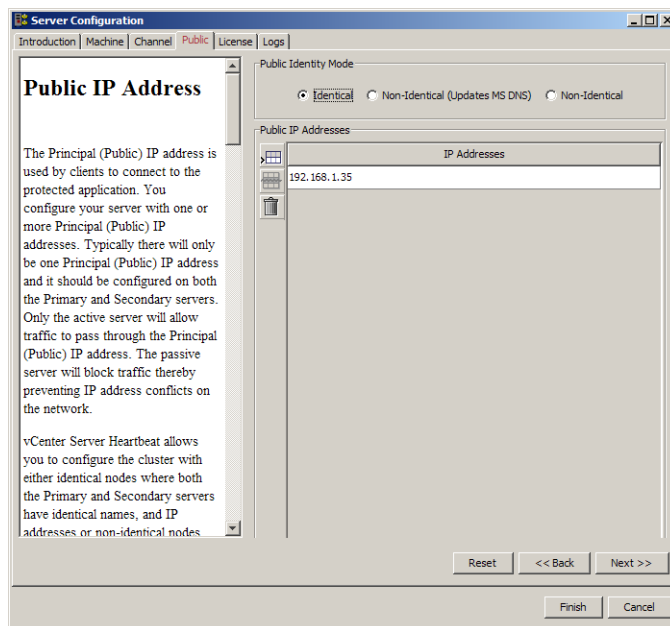
If the **Identical** option was selected during installation, you must configure the server cluster to operate with identical nodes.

To configure the Principal (Public) IP address for Identical Nodes

- 1 Click the **Public** tab and list all the addresses used as Principal (Public) IP addresses.

An address must not appear more than once. The Principal (Public) IP addresses listed cannot be the same as any of the IP addresses listed in the **Channel** tab.

- 2 To add an address, double-click a row and type in the address or select it from a list of currently defined addresses.
- 3 Click **Next** or **Finish**.



Configuring for Non-Identical Nodes

If the **Non-Identical (Updates MS DNS or Non-Identical)** option was selected during installation, you must configure the server cluster to operate with non-identical nodes.

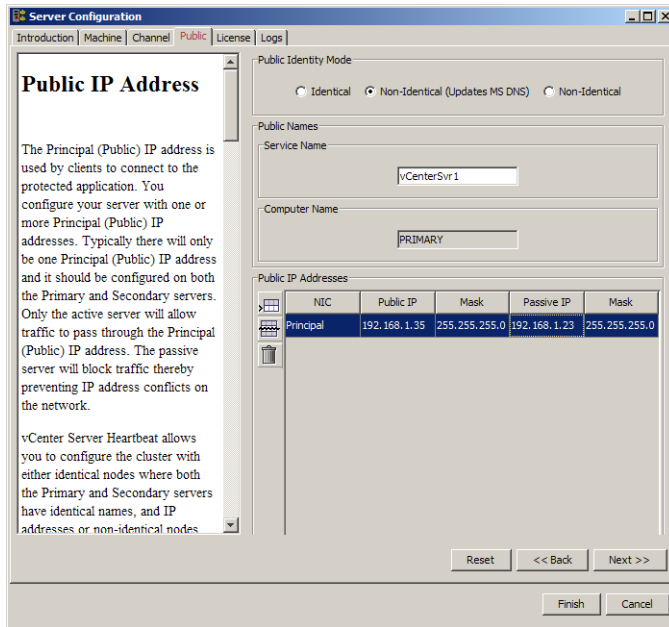
To configure for Non-Identical Nodes

- 1 Select the **Public** tab of the Configure Server wizard and verify that **Non-Identical (Updates MS DNS or Non-Identical)** is selected in the **Public Identity Mode** pane.
- 2 Enter the vCenter Server or SQL Server name in the **Service Name** field.
- 3 In the **NIC** field, select the Principal (Public) network connection in the drop-down.

NOTE If **Non-Identical** was selected for the *Public Identity Mode*, adjacent IP addresses should be reserved and used for the Principal (Public) IP address and the Management IP addresses for the Primary and Secondary Servers when installing vCenter Server Heartbeat on non-identical nodes on servers running Windows 2008.

- 4 Enter the Principal (Public) IP address in the **Public IP** field.
- 5 Enter the Principal (Public) IP address Subnet Mask in the first **Mask** field.
- 6 Enter the reserved Management IP address in the **Passive IP** field.

- 7 Enter the reserved Management IP address Subnet Mask in the second **Mask** field.
- 8 Click **Next** or **Finish**.
- 9 Click **Finish**.

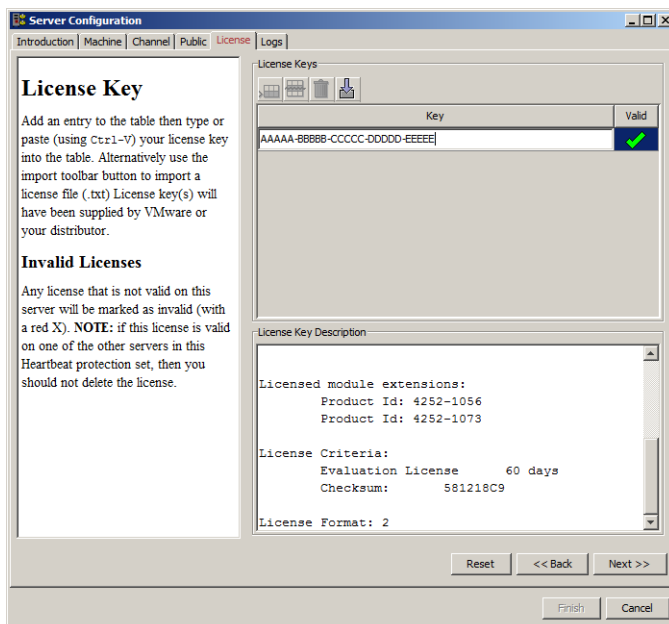


Managing vCenter Server Heartbeat License Keys

To manage vCenter Server Heartbeat license keys, select the **License** tab of the **Configure Server** wizard.

To add an entry to the License Keys table

- 1 Click the **Add Row** icon and enter your VMware vCenter Server Heartbeat serial number.
- 2 Manually type or paste (**using Ctrl-V**) your license key into the table.
- 3 Click **Next** or **Finish**.



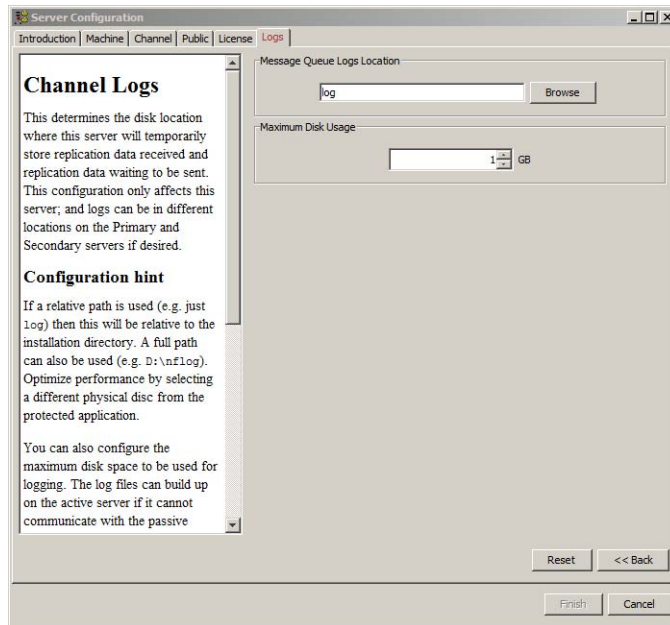
Configuring Message Queue Logs

The server temporarily stores replication data received in the passive server's receive queue and the replication data waiting in the active server's send queue in message queue logs.

This configuration only affects the local server. Logs can be in different locations on the Primary and Secondary servers.

To configure the location of the message queue logs

- 1 Click the **Logs** tab.
- 2 Click **Browse** to navigate to the folder to store the message queue logs.
- 3 Select the folder and click **Next** or **Finish**.



Configuring the Maximum Disk Usage

You can configure the maximum disk space allocated for logging. Log files increase in size on the active server under the following conditions:

- If the active server cannot communicate with the passive server
- Certain operations on the passive server
- If the server is under heavy load

When the disk reaches quota, replication stops and the system is no longer protected.

If using a dedicated disk for log files, consider disabling the quota. To do this, set the quota to zero. If vCenter Server Heartbeat runs out of physical disk space, it must be shut down before it can resume replication. Set the quota with sufficient overflow space so vCenter Server Heartbeat can stop replicating gracefully.

To configure Maximum Disk Usage, click the **Logs** tab, type the maximum dedicated disk space allocated for message queue log files, and click **Finish**.

System Administration and Management

Server Protection

This chapter includes the following topics:

- [“Server Protection Overview”](#) on page 31
- [“Checking the Server Pair Status”](#) on page 31
- [“Monitoring the Status of Servers”](#) on page 33
- [“Configuring Heartbeat Settings”](#) on page 33
- [“Common Administrative Tasks in vCenter Server Heartbeat”](#) on page 34
- [“Configuring Split-Brain Avoidance”](#) on page 36

Server Protection Overview

Protection against operating system or hardware failure affecting the active server is facilitated by two instances of the vCenter Server Heartbeat that monitor one another by sending “I’m alive” messages and reciprocating with acknowledgments over the VMware Channel. If the passive server detects that this process (the heartbeat) has failed, an automatic switchover is initiated.

Additionally, vCenter Server Heartbeat proactively monitors the capability of the active server to communicate with the rest of the network by polling defined nodes around the network, including by default, the default gateway, the primary DNS server, and the Global Catalog server at regular intervals. If all three nodes fail to respond, for example, due to network card failure or a local switch failure, vCenter Server Heartbeat can initiate an automatic switchover, allowing the passive server to assume an identical network identity as the active server.
















Checking the Server Pair Status

The **Server: Summary** page is the default page that opens when administering a pair of servers. The **Server: Summary** page allows you to view the roles that the servers are performing (active or passive), the actions that the servers are currently performing, and summary information of the status of communications and data replication between servers. The lower pane displays status information for each server in the pair.

NOTE To change the currently displayed server, click the server in the graphical representation in the upper pane, or select the server Identity tab (Primary or Secondary Server) in the bottom pane.

The following table lists the possible system statuses and their meanings.

Table 3-1. System Status

Status	Icon	Description
Heartbeat service shutdown		The Heartbeat service is shut down
Initializing		
Replicating		(Normal status.) File and registry changes on the active server are intercepted and replicated to the passive server. The protected applications are monitored.
Not replicating		File and registry replication is in the process of stopping and all protected applications may be closing down.
Switching active server		The system is in the process of performing a switchover.
Connecting to peer server		VMware Channel connections have been established between the two servers.
Disconnecting from peer server		VMware Channel connections have been lost between the two servers.
Stopping replication		File replication is in the process of being stopped and, optionally, all protected applications may be closing down.
Starting replication		The replication process is starting and protected applications are optionally starting.
Starting as active server		The Heartbeat service is initializing on the active server and starting protected applications.
Heartbeat service shutting down		The Heartbeat service is stopping. The Heartbeat service is shutting down, and will no longer participate in replication. Optionally, protected applications may be stopped.
Lost active server		The passive server has lost connection to the active server. If this condition persists for the failover timeout, and failover is permitted between the pair of servers, then a failover will occur.
Active following failover		
Not participating		
Server not responding		The Heartbeat service cannot be contacted on the server.

When viewing the status of the passive server, the status of the file system and registry are displayed graphically. The following tables list possible synchronization statuses and their meanings.

Table 3-2. File Synchronization Status









Status	Icon	Description
Synchronized		Fully synchronized
Unchecked		There are files that are currently unchecked. A full system check did not complete
Out of Sync		Not synchronized
Uninitialized		

Table 3-3. Registry Synchronization Status

Status	Icon	Description
Checking		The registry is currently in the process of synchronization.
Synchronized		Fully synchronized
Error		Not synchronized
Uninitialized		

When the vCenter Server Heartbeat pair establishes a connection, it triggers a file synchronization and verification process to ensure all protected files on both server are identical. The process checks each 64K block of each protected file and performs a checksum to determine whether the blocks differ. If the blocks are the same, the block is marked as synchronized. If the blocks differ, then the block is replicated to the passive server and then marked as synchronized. The file verification and synchronization process is finished after all blocks of all stipulated files are marked as synchronized.

Monitoring the Status of Servers

The **Server: Monitoring** page provides additional information about the status of communications between the servers within the pair. The graphical representation provides an overview of the status of communications between the servers. A green channel icon indicates that the channel is connected and healthy while a yellow dashed channel icon indicates that communications are not operational between the indicated servers. In addition to the heartbeat sent between the servers, VMware vCenter Server Heartbeat also sends a ping to ensure that the servers remain visible to one another.

Configuring Heartbeat Settings

The **Server: Monitoring** page provides three configuration features: **Configure Pings**, **Configure Failover**, and **Configure Response Times**.

Configure Pings

IP addresses of all NICs used for the VMware Channel are identical during installation. vCenter Server Heartbeat, by default, added those IP addresses to the **Server: Monitoring Ping Configuration** dialog. You can add additional targets to the list for each server's channel connection in the event of redundant NICs. The settings in the **Server: Monitoring Ping Configuration** dialog allow vCenter Server Heartbeat to send pings across the VMware Channel in addition to the heartbeat ("I'm alive" messages) to confirm that the server is still operational and providing service.

To configure pings

- 1 Click **Configure Pings** to open the **Server Monitoring: Ping Configuration** dialog.
- 2 Click on the **Ping Settings** tab to configure the Ping Interval.
- 3 Click on the **Ping Routing** tab to add additional IP address for redundant NICs.

Configure Failover

The **Failover timeout** dictates how long vCenter Server Heartbeat waits for a missed heartbeat before it takes a pre-configured action. This value is set to 60 seconds by default.

To configure failover

- 1 Click **Configure Failover** to open the **Server Monitoring: Failover Configuration** dialog.
- 2 Type a new numeric value (seconds) in the **Failover timeout** text box or use the arrow buttons to set a new value.
- 3 Mark or clear the check boxes to select the actions to take if the specified **Failover timeout** is exceeded.
- 4 Click **OK**.

Configuring Response Times

vCenter Server Heartbeat also allows you to configure the following timeouts:

- Time to wait following channel connection before starting replication
- Time to wait following channel disconnection before stopping replication

To configure response times

- 1 Click **Configure Response Times** to open the **Server Monitoring: Response Times** dialog.
- 2 Type new numeric values (seconds) into the text boxes or use the arrow buttons to select new values.
- 3 Click **OK**.

Common Administrative Tasks in vCenter Server Heartbeat

The Server Summary page provides the following buttons that allow you to quickly perform common administrative tasks:

- **Make Active** — Prompts to verify that you want to make the passive server in the pair active. Click **Yes**.
- **Shutdown** — Prompts you to select the server(s) to shut down. If you select the active server, additional options to stop or not stop protected applications appear in the dialog. Click **OK**.
- **Start Replication** — Opens the **Start Replication Options** dialog. Select to start or not start the protected applications and click **OK**. By default, all protection modes start when vCenter Server Heartbeat starts and a manual start is not required unless the system stopped in response to an automated stop.
- **Stop Replication** — Opens the **Stop Replication Options** dialog. Use this method to stop replication, such as to contain a virus infection or to upgrade a protected application. Select whether to stop or not stop protected applications and click **OK**. Replication of data files stops and, if selected, protected applications also stop.

NOTE The VMware vCenter Server Heartbeat service continues to run on the servers, providing heartbeats and protecting the system and network facets of the active server.

- **Start Applications** — Click to start the protected applications on the active server.
- **Stop Applications** — Click to stop the protected applications on the active server.
- **Configure** — Click to open the **Configure** dialog. Select the radio button corresponding to whether you want to stop or leave running the protected applications when vCenter Server Heartbeat is shut down. You can select whether to leave protected applications running upon shutdown when a `net stop` command is issued, and to start protected applications upon startup when a `net start` command is issued. Type a number (seconds) or use the arrow buttons to select an alert threshold value for time difference between servers, which is checked at handshake following startup. Click **OK**.

Forcing a Switchover

After configuring vCenter Server Heartbeat to protect all the required applications and data, the Secondary server can take over from the Primary server in a managed and seamless way called a managed switchover.

This is particularly useful when maintenance work performed on the Primary server requires rebooting the server.

Since a managed switchover cannot be performed during synchronization, it is important to review the queue information prior to attempting a managed switchover. If the queues are large, file operations on the active server are high and for this reason it may be prudent to delay a managed switchover due to the length of time required to completely clear the queue. Queue lengths can be viewed in the **Data: Traffic/Queues** page of the vCenter Server Heartbeat Console.

Prior to performing work on the Primary server, a managed switchover can be triggered by selecting the Secondary server and clicking **Make Active** on the **Server: Summary** page. This changes the server roles such that the active server becomes passive and the passive server becomes active. This means users are able to work continuously while the Primary server is off line.

When the Primary server is back up and running, the managed switchover can be triggered again so that the Primary server becomes active and the previously active server becomes passive.

IMPORTANT The managed switchover process may be performed at any time as long as the systems are fully synchronized with respect to data files and registry replication. *Switchovers cannot be performed if either server is in an unsynchronized or unknown state.*

Recovering From a Failover

A failover differs from a switchover. A switchover is a controlled switch (initiated manually from the vCenter Server Heartbeat Console, or initiated by vCenter Server Heartbeat when preconfigured) between the active and passive servers. A failover happens when any of the following fail on the active server: power, hardware, or VMware Channel communications. The passive server waits a preconfigured time after the first missed heartbeat before initiating a failover. When this period expires, the passive server automatically assumes the active role and starts the protected applications.

The following recovery scenario is based on vCenter Server Heartbeat configuration with the Primary server as active and the Secondary server as passive.

NOTE When failover conditions, such as a power failure, cause failures in both active and passive servers, a condition may result that causes both servers to restart in passive mode. In this situation, manual intervention is required. See [Appendix 10, “Two Passive Servers,”](#) on page 67.

A failover has occurred and the Secondary server is now running as the active server.

- 1 Review event logs on both servers to determine the cause of the failover. For assistance, use the Log Collector tool to collect information and send the output to VMware Support.
- 2 If any of the following issues exist on the Primary server, performing a switchback to the Primary server may not be possible until other important actions are carried out. Do not restart vCenter Server Heartbeat until the following issues have been resolved:
 - **Hard Disk Failure** – Replace the defective hard disk.
 - **Power Failure** – Restore power to the Primary server.
 - **Virus** – Clean the server of all viruses.
 - **Communications** – Replace or repair the physical network hardware.
 - **Blue Screen** – Determine cause and resolve. As required, submit the dump file to VMware Support (www.vmware.com/support) for analysis.
- 3 Run the Configure Server wizard and verify the server identity is set to Primary and the role is set to passive. Click **Finish** to accept the changes.
- 4 Disconnect the VMware Channel network cables or disable the network card.
- 5 Resolve the list of possible failures.
- 6 Restart this server and reconnect or enable the network card again.

- 7 After restart, check that the **Taskbar** icon now reflects the changes by showing **P / -** (Primary and passive).
- 8 On the Secondary active server or from a remote client, launch vCenter Server Heartbeat Console and confirm that the Secondary server is reporting as active.

If the Secondary server is not displaying as active, perform the following steps:

- 1 If the vCenter Server Heartbeat Console is unable to connect remotely, try running it locally. If you are still unable to connect locally, use the Service Control Manager to verify that the service is running. If the service is not running, review the event logs for a cause.
- 2 Run the Configure Server wizard and confirm that the server is set to Secondary and active.

NOTE If vCenter Server Heartbeat is running, you can run the Configure Server wizard but any changes made will not be saved.

- 3 Verify that the protected application is accessible from clients. If accessible, start vCenter Server Heartbeat on the Secondary server. If the application is not accessible, review the application logs to determine why the application is not running.
- 4 Start vCenter Server Heartbeat on the Secondary active server.

NOTE At this point, the data on the Secondary (active) server should be the most up to date and this server should also be the live server on your network. When vCenter Server Heartbeat starts, it overwrites all the protected data (configured in the File Filter list) on the Primary passive server. If you are not sure that the data on the active server is the most current and up to date, contact VMware Support (www.vmware.com/support). Go on to the next step only if you are sure that you want to overwrite the protected data on the passive server.

- 5 Start vCenter Server Heartbeat on the Secondary active server and check that the **Taskbar** icon now reflects the correct status by showing **S / A** (Secondary and active).

Configuring Split-Brain Avoidance

Split-brain Avoidance ensures that only one server becomes active if the VMware Channel connection is lost, but both servers remain connected to the Principal (Public) network. Split-brain Avoidance works by pinging from the passive server to the active server across the Principal (Public) network. If the active server responds, the passive server does not failover, even if the VMware Channel connection is lost. WAN installations require different IP addresses on the Principal (Public) network for the local and remote servers.

To enable Split-brain Avoidance, open the **Server: Monitoring** page in the vCenter Server Heartbeat Console, click **Configure Failover**, and select **Prevent failover if channel heartbeat is lost but Active server is still visible to other servers (recommended)**.

You must configure Management IP addresses on the Principal (Public) network cards of each server to allow the passive server to send a ping. Management IP addresses are additional IP addresses assigned to the network card connected to the Principal (Public) network. They are used to allow the passive server to communicate, because unlike the Principal (Public) IP address, they are not filtered. This allows the passive server to send pings, and is also required to allow the passive server to send email alerts. To configure a Management IP address on the Principal (Public) network card, follow the procedure below.

To configure a Management IP address:

- 1 Open the network properties for the Principal (Public) network connection.
- 2 Double-click **TCP/IP** to display the properties.
- 3 Click **Advanced**.
- 4 Enter an additional (currently unused) IP address in the table.
- 5 Reposition the IP addresses in the list so that the additional (Management) IP address appears first, and the Principal (Public) network address (by which clients connect to the server) appears second.

- 6 Click **OK** on all three dialogs to accept the configuration changes to the network connection.
- 7 After completing all of the steps click **Next** or **Finish**.

The active server must respond within the time period value specified in the **Failover timeout** to prevent a failover from occurring. If the active server responds in a timely manner, the failover process ceases. If the active server does not respond, the failover proceeds.

Network Protection

This chapter includes the following topics:

- [“Communication Status”](#) on page 39
- [“Reviewing the VMware Channel Status”](#) on page 39
- [“Configuring Public Network Connection Checks”](#) on page 39
- [“Setting Max Server Time Difference”](#) on page 40

Communication Status

Use the **Data: Traffic/Queues** page to check the status of the VMware Channel, the active server’s send, and passive server’s receive queues.

Reviewing the VMware Channel Status

The **Data: Traffic/Queues** page displays the VMware Channel status as connected (Green solid icon) or not connected (red broken icon), the statistics of the connection with regards to the data sent by either server, and the size and age of the oldest entry in the active server’s send queue and passive server’s receive queue. The **Channel Connection** tab in the lower pane displays the IP addresses used by the VMware Channel for the Primary to Secondary connections and the port that the communications are using.

Configuring Public Network Connection Checks

The **Network Monitoring** page allows you to view the status of the network and make adjustments to the IP addresses used to ping multiple servers within the network.

The Principal (Public) network monitoring feature, previously discussed, is enabled by default during the installation of VMware vCenter Server Heartbeat. This feature integrates the polling of designated waypoints around the network through the active server’s Principal (Public) connection to ensure connectivity with the Principal (Public) network is operational. By default, the IP addresses of the default gateway, the primary DNS server, and the Global Catalog server are all selected. When one or more of the automatically discovered waypoints are co-located on a physical machine (leading to duplication of IP addresses), the ability to specify additional waypoints manually becomes an advantage. To specify a manual target for Principal (Public) network checking, click **Configure Pings** to invoke the **Ping Configuration** dialog. Select the **Ping Routing** tab to add to or modify the existing target IP addresses for each server to ping.

In a WAN environment, the target addresses for Principal (Public) network monitoring on the Secondary server may be different to those automatically selected on the Primary server. Again, the ability to override automatically discovered selections is provided by manually specifying the target address.

Principal (Public) Network Monitoring is carried out by the active server effectively pinging the target addresses at regular time intervals. The time interval is set by default to every 10 seconds but the frequency may be increased or decreased as required.

Each target is allowed 5 seconds (default) to respond. On slower networks where latency and network collisions are high, increase this interval by changing the **Ping echo timeout** value.

The failure of all three targets to respond is allowed up to the **Auto-switchover if client network is lost for** threshold value. If the failure count of all three targets exceeds this value, vCenter Server Heartbeat initiates an auto-switchover.

Setting Max Server Time Difference

vCenter Server Heartbeat generates a warning if the Primary and Secondary server system clocks are not synchronized.

To override the warning

- 1 Open the **Server: Summary Configure** dialog by clicking the **Configure** button.
- 2 Type a number (seconds) or use the arrow buttons to select an alert threshold value for time difference between servers, which is checked at handshake following startup.
- 3 Click **OK**.

Application Protection

This chapter includes the following topics:

- [“Application Protection Overview”](#) on page 41
- [“Applications Tab”](#) on page 41
- [“Services Tab”](#) on page 43
- [“Applications: Tasks”](#) on page 45
- [“Applications: Plugins”](#) on page 46

Application Protection Overview

vCenter Server Heartbeat incorporates an Application Management Framework (AMFx) to manage vCenter Server Heartbeat plug-ins.

The AMFx provides additional functions while maintaining the traditional stability of VMware software. Use the AMFx to install and remove plug-ins on the fly while vCenter Server Heartbeat continues to provide protection to currently installed applications.

The AMFx also employs sponsorship for protected application files and services. With sponsorship, multiple plug-ins can share files or services. When removing a plug-in, sponsorship prevents removal of a shared file or service that is still required by a remaining plug-in.

vCenter Server Heartbeat uses the **System** plug-in to monitor the server performance. With the **System** plug-in, you can configure a variety of counters and assign actions when associated rules are exceeded.

Applications Tab

The **Applications: Summary** page displays the identity of the active server, the application state and health, details of application types and their corresponding running status and health. From this page, you can start, stop, and configure all protected applications. The lower portion of the pane provides an Applications Log that allows viewing of application events as they occur. This page also provides controls to edit, remove, start, and stop applications, and to configure and edit the configuration of all protected applications.

Reset the Application Health Status

To **Clear** (reset) the Application Health status, click **Clear** in the Application Health pane of the **Applications: Summary** page.

If a problem occurs (for example, a failed service or rule), the Application Health status becomes **Degraded**. Even if vCenter Server Heartbeat corrects the problem (for example, restarts the failed service) or the user corrects the problem, the Degraded status remains until manually cleared. In this state, the Service Discovery Task will not run.

After acknowledging the problem and solving it, click **Clear** to reset the Application Health status. The status updates to provide the actual current Application Health status.

View Application Status

After an application starts and is running, you can view its status in the Applications pane of the **Applications: Summary** page.

Editing Individual Applications

Use vCenter Server Heartbeat to configure the amount of time to wait for an application to start or stop before taking action or reporting the failure.

To configure the timeout settings

- 1 Right-click on the application and select **Edit** from the menu or select the application and click **Edit** at the top of the pane to invoke the **Edit Application** dialog.
- 2 Enter new values into the **Stop Timeout** and **Start Timeout** text boxes or use the arrow buttons to adjust the values (seconds). Click **OK**.

Remove an Application

To remove an application, select the application (in the Applications pane) and do one of the following:

- Right-click on the application and select **Remove** from the menu.

Alternatively, click **Remove** at the top of the pane.

Manually Start and Stop Applications

To stop all protected applications, click **Stop Applications** (at the top of the **Applications: Summary** page). The protected applications set stops. You can view the progress of the stopping in the Applications Log pane.

To start the stopped applications, click **Start Applications** (at the top of the **Applications: Summary** page). The applications start. You can view the progress of the starting in the Applications Log pane.

Configuring Applications

Use the **Applications** page to configure protected applications, enable and disable protection and monitoring. You can maintain applications without stopping vCenter Server Heartbeat or taking the full server offline. During installation, vCenter Server Heartbeat sets default settings for application configurations but accepts modifications to the configurations settings.

To configure applications

- 1 Click **Configure** on the **Applications** page.
You can protect services and start monitoring applications or unprotect services and stop monitoring applications. You can also enable **Verbose Plugin logging**, **Discover protected data at startup**, **Discover protected services at startup**, and set the rule trigger count.
- 2 After making modifications to the configuration, click **OK**.

Application Maintenance Mode

Use the Applications page to disable application protection, service monitoring, and recover for maintenance purposes.

To perform manual maintenance

- 1 On the **Applications** page, select the **Summary** tab.
- 2 Click **Configure**.

- 3 Select **Unprotect services and stop monitoring all applications (for manual application maintenance)**.
- 4 Perform the required maintenance.
- 5 When maintenance is complete, from the **Applications** page of vCenter Server Heartbeat Console, select the **Summary** tab.
- 6 Click **Configure**.
- 7 Select **Protect services and monitor all applications (recommended)**.

Reviewing the State of an Application

After an application successfully starts and is running, you can view the application state in the **Applications: Summary** page. If an application fails, right-click the event in the **Applications Log** and click on **Properties** to invoke the **Event Properties** dialog and investigate the failure.

Reviewing the Applications Log

The **Applications Log** helps troubleshoot the protected application environment. The **Applications Log** provides information about the behavior of all protected applications and includes events such as task status changes, rule triggering, task outputs, and application warnings. Use this log to troubleshoot application errors. The order that entries are displayed can be sorted either ascending or descending by clicking the column title.

Filtering Application Log Entries

vCenter Server Heartbeat can filter Applications Log files to limit the events displayed. By default, all events display in the Applications Log file.

To filter the events to display

- 1 Right-click an event in the **Applications Log** and click **Filter** in the Applications Log pane on the **Applications** page.
- 2 In the upper section, clear the event types you do not want to view.
- 3 To limit the date and time range, select **Only show events from** and edit the date or time range.
- 4 Click **OK**.

Services Tab

The **Applications: Services** page displays both services that you or plug-ins specify and the services related to them by dependency (either as dependents or depends-on). The target states of protected services for the Primary and Secondary server can be specified and are typically **Running** on the active and **Stopped** on the passive. Services are protected if they are **Running** or **Automatic**, and are otherwise logged as unprotected. vCenter Server Heartbeat manages services that depend on protected services (started and stopped) but not monitored (not restarted if stopped by some external agency). vCenter Server Heartbeat monitors protected services (restarted if stopped) but not managed (not stopped if protected applications are stopped).

Adding a Service

To protect a service

- 1 Right-click on a service and select **Add** from the menu or click **Add** on **Applications: Services** page to invoke the **Add Service** dialog. The **Name** drop-down list contains a list of all currently running services.
- 2 Select the service and set the values for **Target State on Active** and **Target State on Passive**. Normally the **Target State on Active** is set to **Running** and the **Target State on Passive** is set to **Stopped**.

NOTE Setting the target state for both the active and passive server to **Running** can cause the service to place a lock on some files preventing synchronization from completing.

- 3 If vCenter Server Heartbeat is to manage the start and stop of the service, select **Manage Starting and Stopping**. If vCenter Server Heartbeat is to monitor the state of the service, select **Monitor State**.

vCenter Server Heartbeat also assigns three sequential tasks to perform in the event of failure. Task options include **Recover Service**, **Application Restart**, **Log Warning**, **Switchover**, and any additional user-defined tasks previously created.

- 4 Assign a task to each of the three failure options and click **OK**.

NOTE If an application with the failure option set to **Application Restart** fails, only the services that have failed are restarted. Dependent services do not stop and restart as a result of the failure.

Editing a Service

To change the options of a protected service

- 1 Select the service and click **Edit**.

The **Edit Service** dialog opens to provide a subset of same options available when adding a new service.

- 2 Make the modifications and click **OK**.

Checking the Status of Services

The **Applications: Services** page displays the status of all protected services. The status shows both the target and actual state for both the Primary and Secondary servers and the Failure Counts for both servers.

Unprotecting User Defined Services and Stopping Monitoring

To unprotect and stop monitoring user defined services

- 1 Click **Applications: Services**, select the user defined service, and click **Edit**.
- 2 Clear **Manage Starting and Stopping** and **Monitor State**. Click **OK**.

Change the Order of Services

The order of services can be modified using **Up** and **Down** arrows. The exact order in which services start and stop is influenced by a number of key factors:

- The order of applications specified by plug-ins determines which services are started first.
- Services can have dependencies, and these must be respected. For example, if service B is listed after service A under the User-Defined group, and service A depends on Service B, Service B starts first.
- Multiple applications can use the same service (the same service can appear under more than one sponsor). The service starts when the first application to reference it starts.
- The order of stopping services is the reverse for starting services.

Removing a Service

To remove a service

- 1 Select the service in the **Applications: Services** page.
- 2 Click **Remove**.

The service is removed from the protected list.

Applications: Tasks

Tasks are a generalization and extension of start, stop, and monitor scripts. Task types are determined by when the tasks run, and include the following:

- **Network Configuration** – This is the first type of task that runs when applications start and is intended to launch `dnscmd` or `DNSUpdate`. The task can launch a batch script containing multiple `dnscmd` commands. Network Configuration tasks are the only types of task that can vary between Primary and Secondary servers.
- **Periodic** – These tasks are run a specific configurable intervals.
- **Pre/Post Start** – These tasks run before and after services start on the active server.
- **Pre/Post Stop** – These tasks run before and after services stop on the active server.
- **Pre/Post Shadow** – These tasks run before and after Data Rollback Module creates a shadow copy on the active server (Not available in this release).
- **Rule Action** – Configure these tasks to run in response to a triggered rule or when a service fails its check.

You can define and implement tasks at the command line, such as launching a batch-script. Examples of built-in tasks include monitoring a protected service state on the active and passive servers.

vCenter Server Heartbeat can use plug-ins to define and implement tasks. An example of a plug-in-defined task is the discovery of protected data and services for a particular application.

Click **Applications: Tasks** to open the Tasks page.

Adding a Task

To add a task

- 1 Click **Add** to invoke the **Add Task** dialog. Assign a name to the task.
- 2 Select the task type from the drop-down list.
- 3 Select the identity of the server the task runs on (Primary or Secondary).
- 4 In the **Command** text box, type in the path or browse to the script, `.bat` file, or command for the task to perform.

NOTE When the **Command** entry requires specific user credentials, you must select that user from the **Run As** drop-down list. To add a user account, click **User Accounts** (near the top of the pane). See [“View, Add, and Remove User Accounts”](#) on page 46.

- 5 Click **OK**.

Editing a Task

You can edit the interval of a task or disable a task.

To edit a task

- 1 Right-click on an existing task and select **Edit** from the menu or select the task and click **Edit** at the top of the pane to invoke the **Edit Task** dialog.
- 2 Edit the parameters of the task.
- 3 Click **OK**.

Remove a Task

To remove a task, select the task in the **Applications: Task** page and click **Remove**.

Change the Order of Tasks

To change the order of tasks, use the **Up** and **Down** arrows (near the top of the pane) or on the right-click menu to change the order in which the tasks appear in the tasks list.

Starting a Task Manually

vCenter Server Heartbeat provides options to launch a task immediately, or to launch a task after a designated time period elapses, or following the occurrence of a specified event.

To launch the task immediately

- 1 Select the task from the task list.
- 2 Right-click on the existing task and select **Run Now** from the menu or click **Run Now** at the top of the pane.

vCenter Server Heartbeat immediately launches the task.

View, Add, and Remove User Accounts

vCenter Server Heartbeat Console allows you to view, add, and remove user accounts used to run tasks.

Add a User Account

To add a user account, click **User Accounts** to invoke the **User Accounts** dialog.

- 1 Click **Add** to invoke the **Add User** dialog.
- 2 Type the name of the User, the associated Domain, and a Password into the corresponding text boxes.
- 3 Click **OK**.

Remove a User Account

- 1 To remove a user, select the user account from the list in the **User Accounts** dialog and click **Remove**. A confirmation message appears.
- 2 Click **Yes**.

Applications: Plugins

Plug-ins support specific applications and contain all of the components to protect the designated application. Plug-ins start and stop the application, monitor the application, and provide all rules necessary to ensure that application is available in the event of a failure by initiating a auto-switchover when configured.

Install a Plug-In

vCenter Server Heartbeat allows you to install and upgrade plug-ins as needed to support applications.

IMPORTANT Plug-ins should be installed only on the active server. Installation of a plug-in on a passive server may cause an Exception to occur.

To install a new plug-in

- 1 Click **Applications: Plugin** to open the **Plugins** page.
- 2 Right-click an existing plug-in and select **Install** from the menu or click **Install** at the top of the pane to invoke the **Install Plugins** dialog.
- 3 Type a path to the plug-in location or click **Browse** to navigate to the plug-in location. The path statement is case-sensitive.
- 4 Click **OK**.

Editing a Plug-in

vCenter Server Heartbeat allows you to edit the configuration of user installed plug-ins.

To edit the plug-in configuration

- 1 Right-click on an existing plug-in from the Plugins list and select **Edit** from the menu or select the plug-in and click **Edit** at the top of the pane to invoke the **Edit Plugin** dialog.
- 2 Review the configuration options before making modifications as they are specific to each plug-in.
- 3 Click **OK**.

Uninstalling a Plug-in

You can uninstall a plug-in when you upgrade or remove the application the plug-in protects, or when directed by VMware Support.

To uninstall a plug-in

- 1 Right-click an existing plug-in and select **Uninstall** or select the plug-in and click **Uninstall** at the top of the pane.
- 2 Click **Yes**.

Status and Control

This chapter includes the following topics:

- [“vCenter Server Heartbeat Console”](#) on page 49
- [“About vCenter Server Heartbeat Console”](#) on page 49
- [“Navigate vCenter Server Heartbeat Console”](#) on page 50
- [“Change the Font for vCenter Server Heartbeat Console”](#) on page 50
- [“Work with Groups and Pairs”](#) on page 50
- [“Add or Remove a vCenter Server Group”](#) on page 50
- [“Add, Edit, Move, and Remove Pairs in VCenter Server Heartbeat Groups”](#) on page 51

vCenter Server Heartbeat Console

vCenter Server Heartbeat operates over a Pair of vCenter Server Heartbeat servers and is administered in these Pairs.

The vCenter Server Heartbeat Console is used to carry out the day-to-day administration of one or more Pairs of servers.

The vCenter Server Heartbeat Console runs from either of the two servers in the Pair or remotely from another machine in the same subnet that has vCenter Server Heartbeat or the vCenter Server Heartbeat Client Tools installed.

NOTE You can install vCenter Server Heartbeat on a Windows XP and Windows Vista SP1 or later workstation to act as a client to the server Pair or Windows Server 2003. Run `setup.exe` from the `setup` CD folder on the workstation and select **Install Client Tools** under the installation set.

About vCenter Server Heartbeat Console

You can start vCenter Server Heartbeat Console from any server in the vCenter Server Heartbeat Pair.

To start vCenter Server Heartbeat Console

- 1 Right-click the VMware vCenter Server Heartbeat interactive status icon on the Windows tool tray (located on the right side of the Windows tool bar). The vCenter Server Heartbeat quick access menu opens.
- 2 Select **Manage Server**. The vCenter Server Heartbeat Console opens in a window and shows the Heartbeat Servers (overview) pane.

Alternatively you can start vCenter Server Heartbeat Console from the VMware program group on the Windows Start menu. This is the only method supported if vCenter Server Heartbeat Console has been installed on a workstation that is not part of the Pair.

Navigate vCenter Server Heartbeat Console

After vCenter Server Heartbeat Console is running, use the navigation panel on the left of the vCenter Server Heartbeat Console window to view and select Groups and Pair connections you can manage with vCenter Server Heartbeat Console.

NOTE A Group is an arbitrary collection of vCenter Server Heartbeat Pairs used for organization.

A Connection, or Pair Connection allows vCenter Server Heartbeat Console to communicate with a vCenter Server Heartbeat Pair either on the same machine or remotely.

See [“Add or Remove a vCenter Server Group”](#) on page 50 and [“Add, Edit, Move, and Remove Pairs in vCenter Server Heartbeat Groups”](#) on page 51 for information on how to add Groups and Pair Connections to vCenter Server Heartbeat Console.

The selection of Group or Pair you make in the navigation panel “points” the vCenter Server Heartbeat Console to that Group or Pair and vCenter Server Heartbeat Console provides information related to only the selected Group or Pair. To avoid confusion, pay particular attention to the selection in the navigation panel when you are managing more than one Group or Pair.

NOTE Groups and Pairs are not automatically detected by vCenter Server Heartbeat Console. Each Group or Pair you want to manage must be added to vCenter Server Heartbeat Console before you can use it to view status or change settings for that Group or Pair Connection.

Select a Pair in the navigation panel of vCenter Server Heartbeat to show a set of tabs and sub-tabs that offer detailed status and control of the associated vCenter Server Heartbeat server in the Pair.

Change the Font for vCenter Server Heartbeat Console

You can change the font used in the vCenter Server Heartbeat Console interface.

To change the font of the interface

- 1 Select **Font Selection** from the **Preferences** menu. The **Font Selection** dialog opens.
- 2 In the **Style** pane, scroll to and click to select a font.
- 3 In the **Size:** text box, type a new numeric (point) size or use the arrow buttons to change the font size.
- 4 Click **OK**. A confirmation message appears.
- 5 Click **Yes** to confirm the changes and restart vCenter Server Heartbeat to apply the new font settings. Click **No** to restart later; the changes will be applied the next time vCenter Server Heartbeat Console is started.

Work with Groups and Pairs

This section describes how to use vCenter Server Heartbeat Console to work with Groups (add or remove) and Pairs (add, edit, move, or remove), and to manage the Username and Password settings on the servers in a vCenter Server Heartbeat Pair.

Add or Remove a vCenter Server Group

The Add Group feature in vCenter Server Heartbeat Console allows you to add new vCenter Server Heartbeat Groups to manage.

To add a vCenter Server Heartbeat Group

- 1 Open vCenter Server Heartbeat Console and click **Add Group** in the tool bar, select **Add Group** from the **File** menu, or right-click an existing group in the navigation panel and select **Add Group** from the menu.
- 2 Type the name for the new group into the text box and click **OK**. The newly created group appears in the navigation panel on the left of the vCenter Server Center Heartbeat window.

Remove a vCenter Server Heartbeat Group

The Remove Group feature in vCenter Server Heartbeat allows you to remove existing vCenter Server Heartbeat Groups from management.

To remove a vCenter Server Heartbeat Group

- 1 Select the Group to be removed in the navigation panel of vCenter Server Heartbeat Console. Click **Remove Group** in the tool bar or select **Remove Group** from the **File** menu.
- 2 A confirmation message appears. Click **Yes**.

Add, Edit, Move, and Remove Pairs in VCenter Server Heartbeat Groups

When you created a vCenter Server Heartbeat Group using the instructions in [“Add or Remove a vCenter Server Group”](#) on page 50, you created an empty container. Next you must add the connections to the Pair or Pairs that make up your new vCenter Server Heartbeat Group.

Add a New Connection

The Add Connection feature in the vCenter Server Heartbeat Console allows you to add a new Pair Connection to an existing vCenter Server Heartbeat Group.

To Add a new connection

- 1 In the navigation panel, select the vCenter Server Heartbeat Group to receive the new connection. Click **Add Connection** in the tool bar, select **Add Connection** from the **File** menu, or right-click an existing group in the navigation panel and select **Add Connection** to invoke the **Add Connection** dialog.
- 2 Type the Host Name or IP address for the new connection into the text box, select the Port Number (if different from the default value of 52267), and select a group from the **Add to Group** drop-down list (to add the connection to a Group other than the one currently selected).

NOTE When a you attempt to connect to vCenter Server Heartbeat for the first time, you are presented the option to accept the SSL certificate from the server. To continue connecting to vCenter Server Heartbeat, you must accept the SSL certificate.

- 3 Click **OK**. The newly created connection appears in the navigation panel on the left of the vCenter Server Heartbeat Console window, and the **Server: Summary** page updates to represent any existing network relationships of the added server.

NOTE You may be prompted to login. If so, login using a valid administrator-level Username and Password for the server for which you are adding a connection, and click **OK**.

- 4 Enter the remaining connections necessary to define the new vCenter Server Heartbeat Group.

Edit a Connection

The Edit Connection feature in the vCenter Server Heartbeat Console allows you to change the Port Number for existing connections.

To edit a connection

- 1 In the navigation panel, select the connection you want to change and select **Edit Connection** from the **File** menu, or right-click an existing connection in the navigation panel and select **Edit Connection** from the menu to display the **Edit Connection** dialog.

NOTE When a configured connection is not found, an error message may be displayed. Click **Edit Connection** to reconfigure the connection.

- 2 Type the new value for the Port Number into the text box, or use the **Up** or **Down** arrow controls to the right of the text box to select a new value. Click **OK**.

Move a Connection

The Move Connection feature in vCenter Server Heartbeat Console allows you to reassign an existing Pair to a different Group.

To move a connection

- 1 Select the Pair in the navigation panel and click **Move Connection** in the tool bar, select **Move Connection** from the **File** menu, or right-click on the Connection in the navigation panel and select **Move Connection** from the menu to display the **Move Connection** dialog.
- 2 Select the destination Group to receive the Connection from the drop-down list and click **OK**.

Remove a Connection

The Remove Connection feature in vCenter Server Heartbeat allows you to remove an existing Connection.

To remove a connection

- 1 Select the Connection in the navigation panel and click **Remove Connection** in the tool bar, select **Remove Connection** from the **File** menu, or right-click on the connection in the navigation panel and select **Remove Connection** from the menu to.
- 2 A confirmation dialog appears. Click **Yes**.

Edit Username and Password Settings

The Edit Username and Password feature in vCenter Server Heartbeat Console allows you to change the Username and Password settings used to connect to a given Pair.

To edit Username and Password

- 1 Select a connection in the navigation panel and select **Edit User Name and Password** from the **File** menu or right-click on the Connection in the navigation panel and select **Edit User Name and Password** from the menu to display the **Edit User Name and Password** dialog.
- 2 Type new values for **User Name** and **Password** into the corresponding text boxes and click **OK**.

Review the Status of vCenter Server Heartbeat Groups and Pairs

Click on the top level of vCenter Server Heartbeat Console, Heartbeat Servers to view a list of all managed Pairs and a quick status of the protected applications, network, files system, and registry settings for each Group.

The status hyperlinks in this overview window links to pages that provide more specific related information and management controls.

Click:

- The **Server** connection name to view the **Server: Summary** page
- The **Applications** status to view the **Applications: Summary** page

- The **Network** status to view the **Network Monitoring** page
- The **File System** or **Registry** status to view the **Data: Replication** page

Exit vCenter Server Heartbeat Console

Click **Exit** on the **File** menu. The **Confirm Exit** message appears.

Click **Yes**.

Shut Down Windows Without Stopping vCenter Server Heartbeat

Always stop vCenter Server Heartbeat before attempting to shut down Microsoft Windows. If an attempt is made to shut down Windows without stopping vCenter Server Heartbeat, a confirmation message is displayed.

When the confirmation message is displayed, click **Cancel** and stop vCenter Server Heartbeat before attempting Windows shut down again.

Performance Protection

This chapter provides detailed information about the topic [“Applications: Rules”](#) on page 55.

Applications: Rules

Rules are implemented by plug-ins (there are no user-defined rules). Configure rule actions to trigger the rule that performs specific tasks. Rules have two trigger properties:

- **Timed** – They must evaluate as true continuously for the specified duration to trigger.
- **Latched** – They trigger as soon as they evaluate to true.

Rules

The **Applications: Rules** page provides a list of rules with their current status and two ways to edit and check rules.

Checking a Rule Condition

vCenter Server Heartbeat allows you to check the rule conditions of the current configuration against the attributes of the system or application.

To check a rule condition

Right-click on the rule and select **Check Now** from the menu or click **Check Now** at the top of the pane. The rule condition is displayed in the pane.

Edit a Rule

Rules are implemented by plug-ins and cannot be created by users. Each plug-in contains a default set of rules with options that may be modified by the user.

To edit a rule

- 1 Right-click on the rule and select **Edit** from the menu or click **Edit** at the top of the pane.
- 2 Edit the parameters of the rule and click **OK**.

Rules Installed by vCenter Server Heartbeat Plug-Ins

The following plug-ins implement the rules listed.

vCenter Server Plug-In

- Check health of Tomcat server
- Check vCenter License
- Check Connection to vCenter

vCenter SQL Server Plug-In

- Default Instance Buffer Cache Hit Ratio
- Default Free Pages
- Default Instance Free Pages
- Named Instance Working Set
- Named Instance Buffer Cache Hit Ratio
- Named Instance Free Pages
- Named Instance Total Server Memory

vCenter Server Heartbeat System Plug-In

- DiskAvgSecsPerRead
- DiskAvgSecsPerWrite
- DiskIO
- DiskQueueLength
- DiskReadsPerSec
- DiskWritesPerSec
- DiskWriteable
- FreeDiskSpace
- FreeDiskSpaceOnDrive
- MemoryCommittedBytes
- MemoryCommittedBytesPercent
- MemoryFreePTEs
- MemoryPageReadsPerSec
- MemoryPageWritesPerSec
- MemoryPagesPerSec
- MemoryPagingFileUseage
- PageFaultsPerSec
- ProcessorIntsPerSec
- ProcessorLoad
- ProcessorQueueLength
- RedirectorBytesTotalPerSec
- RedirectorNetworkErrorsPerSec
- ServerBytesTotalPerSec
- ServerWorkItemShortages ≥ 3 (if the rule for server work item shortages is triggered, consult Microsoft documentation on setting the registry values for InitWorkItems or MaxWorkItems accordingly).
- ServerWorkQueueLength
- SystemContextSwitches

Data Protection

This chapter includes the following topics:

- [“Data Protection Overview”](#) on page 58
- [“Replication”](#) on page 58
- [“Registry and File Synchronization Status”](#) on page 58
- [“Initiate a Full Registry Check”](#) on page 58
- [“Initiate a Full System Check”](#) on page 58
- [“Initiate File Synchronization Manually”](#) on page 59
- [“Initiate Verify and Synchronize Manually”](#) on page 59
- [“File Filters”](#) on page 59
- [“Determine Effective Filters”](#) on page 59
- [“Add a User-Defined Inclusion Filter”](#) on page 60
- [“Add a User-Defined Exclusion Filter”](#) on page 60
- [“Edit User Defined Inclusion/Exclusion Filters”](#) on page 60
- [“Remove User-Defined Filters”](#) on page 60
- [“Automatic Filter Discovery”](#) on page 60

Data Protection Overview

vCenter Server Heartbeat can protect many permutations or combinations of file structures on the active server by the use of custom inclusion and exclusion filters configured by the administrator. See [“File Filters”](#) on page 59 for more information.

The filter driver identifies files to protect and disk I/O operations to intercept and replicate to the passive server. Use this driver to filter files for inclusion in or exclusion from the replication process.

vCenter Server Heartbeat protects a folder called **Protected** on the system partition.

NOTE vCenter Server Heartbeat forbids replicating certain files and folders by using a veto. If an inclusion filter includes any of those files or folders, the entire filter is vetoed, even if an exclusion filter is used to prevent replication of those files and folders. Examples of folders are the vCenter Server Heartbeat installation directory or the `system32` folder.

The VMware application folder contains the active server’s send and passive server’s receive queues on the active and passive servers. This folder must be explicitly excluded from file protection.

Replication




You can view replication status and manage data replication using the **Data: Replication** page.

Registry and File Synchronization Status

Two panes near the top of the **Replication** page in vCenter Server Heartbeat Console, **File System Synchronization Status** and **Registry Synchronization Status**, provide graphical status information.

The synchronization status for each file or folder can read three different values depending on the verification and synchronization states as described in [Table 8-1](#).

Table 8-1. Synchronization Status

Icon	Description
	The file is verified and successfully synchronized.
	The file is not synchronized on the active and passive servers. This state often follows a failover and requires manual synchronization and verification.
	The file or folder has not been checked because a full system check has not been performed or the system check has not yet reached the file or folder.

Initiate a Full Registry Check

The registry check re-scans and synchronizes all registry keys specified in the built-in registry filters between the servers and the results are displayed in the **Registry Synchronization Status** pane.

To initiate a full registry check

Click **Full Registry Check** in the **Registry Synchronization** pane.

Initiate a Full System Check

You can verify and synchronize the entire protected file set with a Full System Check. A Full System Check performs the same block level check of all the files set by the file filters in the initial startup synchronization and verification.

When you click **Full System Check**, a dialog asks you to confirm the request and warns you that depending on the amount of data under protection, this task can take a long time to complete (for example, a number of hours). Click **Yes** to perform the check.

Switchover cannot occur until the full system check completes and the File System Status is **Synchronized**.

The File System Status is **Unchecked** when you cancel the task. Depending on the amount of data, resynchronization can take substantial time to complete.

Initiate File Synchronization Manually

The **Data: Replication File Hierarchy** pane displays files that were detected as out of synchronization.

To initiate file synchronization manually

- To manually synchronize the specified files, click **Synchronize**.
- You can resynchronize files at any time. Select multiple files with the **Shift** or **Ctrl** keys and click **Synchronize**.
- Select a folder and select **Including Subdirectories** to synchronize files within folders.

A progress graphic displays the status of the verification or synchronization operation. When complete, the status displays a green **Synchronized** icon.

Initiate Verify and Synchronize Manually

A manual or scheduled synchronization and verification request is defined as a task that is queued for processing after the running task completes. Tasks display in the **Pending Tasks** pane. You can cancel individual tasks. If you cancel a scheduled task, you risk an unchecked system. Possible consequences of canceling tasks display in a warning message.

To verify and synchronize folders

- 1 Use **Verify & Synchronize** and select **Include Subdirectories** to ensure all underlying files and subfolders are included in the verification and synchronization operation.
- 2 Right-click a folder to access a popup menu to perform quick synchronization and verification of folders and subfolders.

File Filters

File filters dictate which files are protected and the disk I/O operations to intercept and replicate to the passive server. File filters also allow you to customize the inclusion and exclusion of files from the replication process.

The File Filters pane of the **Data** page allows you to set up and manage inclusion and exclusion filters.

The File Filters pane contains three columns: Filter, State, and Detail.

- The Filter column lists the pattern for protecting files and folders on the active server.
- The State of the filter identifies the filter as Effective, Subset (contained within another filter), or Not Effective (not contained within another filter). An Effective filter is properly configured and functions to protect (replicate) the stipulated files to the passive server.
- The Detail describes the file filter details based upon the state of the file filter.

Determine Effective Filters

An Effective Filter is the result of the remainder of the files and folders stipulated in the Inclusion Filter after removing the files and folders in the Exclusion Filter.

Filters are compared with each other, and if one filter is a superset of another, the superset filter is used. You can configure a single, general filter to replace file servers with 1000s of individual shares requested by a plug-in.

Add a User-Defined Inclusion Filter

Inclusion Filters create a subset of files to specify items to include for protection.

To define filters that include files and folders for protection and replication

- 1 In the **Data: File Filters** pane, click **Add Inclusion Filter** to open the **Add Inclusion Filter** dialog.
- 2 Type the complete path and pattern, specify a pattern containing wildcards, or use **Browse** to locate the file or folder.
- 3 Click **OK**.

The two forms of wildcards available are *, which matches all files in the folder, and **, which matches all files, subfolders and the files in the subfolders of the folder.

After defining the filter, you can add additional Inclusion Filters.

Add a User-Defined Exclusion Filter

Exclusion Filters create a subset of an Inclusion Filter to specify items to exclude from protection.

To define filters that exclude files and folders from protection and replication

- 1 In the **Data: File Filters** pane, click **Add Exclusion Filter** to open the **Add Exclusion Filter** dialog.
- 2 Type the complete path and pattern, specify a pattern containing wildcards, or use **Browse** to locate the file or folder.
- 3 Click **OK**.

The two forms of wildcards available are *, which matches all files in the folder, and **, which matches all files, subfolders and the files in the subfolders of the folder.

Edit User Defined Inclusion/Exclusion Filters

Inclusion and exclusion filters can be edited by selecting the filter and clicking **Edit** at the top of the File Filters pane or right-clicking the filter and selecting **Edit** from the menu. Edit the value in the **Pattern:** text box by typing over the current file filter definition.

Remove User-Defined Filters

When necessary, user defined inclusion and exclusion filters can be removed.

To remove an Inclusion or Exclusion Filter

- 1 Select the filter in the File Filters list and click **Remove**, or right-click on the filter in the File Filters list and select **Remove** from the menu.
- 2 A confirmation message appears. Click **Yes**.

Automatic Filter Discovery

When Administrators make changes to the configuration, vCenter Server Heartbeat adjusts file filter protection for protected locations. Additionally, the SQL Server plug-in provides database protection including changes or additions to the database and log files.

Other Administrative Tasks

This chapter includes the following topics:

- [“Configure Alerts”](#) on page 61
- [“Configure Alert Reporting”](#) on page 61
- [“Test Alert Reporting”](#) on page 62
- [“Configure Event Log Files”](#) on page 62
- [“Review Event Logs”](#) on page 63

Configure Alerts

vCenter Server Heartbeat can send predefined alerts to remote administrators by email using **Logs > Configure Alerts**.

You can configure alerts in by clicking **Configure Alerts** on the **Logs** page

You can configure three alert states: Red alerts are critical, yellow alerts are not as serious, and green alerts are informational. These alerts are preconfigured with the recommended alerting levels.

To reconfigure each event to trigger red, yellow, green, or no alert, select the appropriate tab, select the check boxes, and click **OK**.

Configure Alert Reporting

vCenter Server Heartbeat can alert the administrators or other personnel and route logs by email when an Alert condition exists.

To configure email alerts

- 1 Click **Logs: Mail Settings** to open the **Mail Settings** dialog.
- 2 Type the outgoing SMTP server of both the Primary server (when active) and the Secondary server (when active) in the appropriate fields.
- 3 Type the FQDN of the mail server. Type an email address that is authorized to send mail through the SMTP server.
- 4 If the SMTP servers require authentication to accept and forward SMTP messages, select **Mail Server requires authentication** and specify the credentials for an appropriate authenticated user account.
- 5 Click **OK**.

You can configure email recipients in the **On Red Alert**, **On Yellow Alert**, and **On Green Alert** tabs of the **Configure Alerts** dialog after configuring the trigger levels and the email server.

Red, Yellow, or Green alert triggers email to the same or different recipients. The process to add recipients is the same for all trigger levels.

- 1 Click the **On Red Alert**, **On Yellow Alert**, or **On Green Alert** tab and select **Send mail**.
- 2 Select the frequency for the email to be sent.
- 3 Click **Add** and type a fully qualified email address for each recipient for the respective trigger level alert.
- 4 To delete a recipient, select the recipient's email address in the Mail Recipients pane and click **Remove**.

Use the preconfigured subject and content of the alert emails for Red, Yellow, or Green alerts. You can add content as required. VMware recommends leaving the preconfigured subject and content and if necessary, add additional information.

Another method to send an alert notification is:

- 1 Select **Run Command** under the pertinent alert state.
- 2 **Browse** to the script to run or use a command line argument to run on the alert trigger.

The preconfigured WScript command creates an event in the Application Event Log and can be customized to include vCenter Server Heartbeat specific informational variables as detailed in [Table 9-1](#).

Table 9-1. Script Variables

Variable	Value
\$EventId	Id of event as listed above
\$EventName	Human-readable name of event
\$EventDetail	Detail message for event
\$EventTime	Time at which event occurred

The following command line argument creates an event in the Application Event Log listing the machine that caused the alert, the time the alert occurred, the name, and details of the alert:

```
WScript //T:10 $(installdir)\bin>alert.vbs "VMware vCenter Server Heartbeat alert on
$EventHost at $EventTime because $EventName ($EventDetail). Event Id is $EventId"
```

- 3 Click **OK**.

Test Alert Reporting

Click **Test Alert Reporting** to run a test alert email. This way you can avoid triggering an actual alert during the operation of the active server.

Configure Event Log Files

vCenter Server Heartbeat allows you to configure Event Log files to direct where the log file is stored and the number of events to be recorded.

To configure default settings for log files

- 1 Click **Logs: Configure** and select the **General** tab to define the filename and path of the exported comma-separated variable (CSV) file.
- 2 Type a path and filename or use **Browse** and navigate the file.
- 3 Adjust the length of the event list to meet your needs by increasing or decreasing the value (the default is to record 300 events) in the **Record at most** field.
- 4 Click **OK**.

Configure Log File Email Recipients

Use vCenter Server Heartbeat to email the log to specified personnel at predetermined intervals.

To configure email log notifications

- 1 To configure vCenter Server Heartbeat to email a copy of the log file, click on the **Mail Log File** tab, select **Mail Every**, and configure the day and time to send the log file.
- 2 Specify the recipients. Click **Add** on the top left of the email recipient field and type the email address in the **Add Mail Address** dialog.
- 3 To remove a recipient, select the recipient's email address in the **Mail Log File** pane and click **Remove**.
- 4 Click **OK**.

Review Event Logs





The **Logs: Event Log** pane lists events logged chronologically by default.

The Event log shows the time the event happened, the type, the source, its importance, and its detail. The display order for events can be sorted either descending or ascending by clicking on the column heading. Since the detail in the data grid is truncated, it may be necessary to review the log in more detail.

To review the details

- 1 Double-click the entry in the data grid.
Event Properties displays the full detail and trace of the log that caused the event and the source of the error to aid in troubleshooting.
- 2 Use the **Up** and **Down** arrows in this window to review other logs. This feature is useful where many events have occurred simultaneously and helps to identify the source of the problem.
- 3 Click **Close** to close the **Event Properties**.

Table 9-2. Log Events

Icon	Description
	Errors within the underlying operation of vCenter Server Heartbeat and can be considered critical to the operation of the system.
	Warnings generated for discrepancies within the vCenter Server Heartbeat operational environment that are not deemed critical to the operation of the system.
	System logs are generated following normal vCenter Server Heartbeat operations. You can use these logs to verify the success of processes such as file synchronization.
	Information on operations within the graphical user interface rather than operations on vCenter Server Heartbeat service, such as Test Alert Reporting.

Event Log Filters

The list of logs that vCenter Server Heartbeat displays may be filtered to hide less important logs.

To filter log events by importance




- 1 Click **Filters** to invoke the Event Log Filters dialog.
- 2 Select **Events of at Least**.
- 3 Select the **Show events of at least** check box in the **Importance** group.
- 4 Select the importance level from the drop-down list and click **OK**.

Only logs equal to or above the select severity are displayed.

To filter log events by date and time range

- 1 Select the **Only show events from** check box and adjust the start date, end date, and time.
- 2 Click **OK**.

Table 9-3. Event Log Buttons

Icon	Purpose
	To export the list to a comma-separated variable file, click Export event log at the top left of the Log Details data grid. You can configure the filename and path to export the data in the Configuration tab.
	To immediately email the list, click E-mail .
	To clear the list, click Remove all Entries at the top left of the Log Details data grid.

This chapter includes the following topics:

- [“Troubleshooting Unexpected Behaviors”](#) on page 65
- [“Two Active Servers”](#) on page 65
- [“Two Passive Servers”](#) on page 67
- [“Synchronization Failures”](#) on page 68
- [“Registry Status is Out of Sync”](#) on page 70
- [“Channel Drops”](#) on page 70
- [“Subnet or Routing Issues ”](#) on page 74
- [“MaxDiskUsage Errors”](#) on page 74
- [“MaxDiskUsage Error Messages”](#) on page 75

Troubleshooting Unexpected Behaviors

The following unexpected behaviors illustrate symptoms, causes and resolution for a given scenario.

Two Active Servers

When two identical active servers are live on the same network, vCenter Server Heartbeat refers to the condition as Split-brain syndrome. Two active servers do not occur by design and when detected, must be resolved immediately.

Symptoms

Split-brain syndrome is identified by the following symptoms:

- Both servers in the pair are running and in an active state. The task bar icons display **P / A** (Primary and active) and **S / A** (Secondary and active).
- An IP address conflict occurs on a server pair running vCenter Server Heartbeat on the Principal (Public) IP address.
- A name conflict occurs on a server pair running vCenter Server Heartbeat. In a WAN environment the Primary and Secondary servers connect to the network using different IP addresses. However, if the servers are running with the same name and are visible to each other across the WAN, a name conflict occurs.
- Clients (for example, VI Client, ESX, etc.) cannot connect to the server running vCenter Server Heartbeat.

Causes

The most common causes of two active servers (Split-brain syndrome) are as follows:

- Loss of the VMware Channel connection (most common in a WAN environment)
- The active server is too busy to respond to heartbeats
- Incorrect configuration of the vCenter Server Heartbeat software

You must determine the cause of the Split-brain syndrome and resolve the issue to prevent this condition from recurring.

Resolution

To resolve Split-brain syndrome, identify the server with the most up-to-date data. If you identify the wrong server you risk losing data. You must reinstate the correct server.

To identify the server with the most up-to-date data

- 1 Check the date and time of files on both servers. Make the most up-to-date server the active server.
- 2 From a client PC on a LAN, run `nbtstat -A 192.168.1.1` where the IP address is the Principal (Public) IP address of the server. This can help identify the MAC address of the server currently visible to client machines.

NOTE If both active servers were servicing clients, perhaps at different WAN locations, you can make only one server active. Both servers contain recent data that cannot be merged using vCenter Server Heartbeat. To restart replication, make one server active and one server passive. When replication restarts, the active server overwrites all data on the passive server. You can manually extract the up-to-date data from the passive server prior to restarting replication. Consult the Microsoft knowledge base for information on various tools for this purpose. For further information, contact your VMware support representative.

To resolve two active servers (Split-brain syndrome)

- 1 Identify the server with the most up-to-date data or the server to make active.
- 2 Shut down vCenter Server Heartbeat on both servers (if running).
- 3 On the server to make passive, right-click the Task bar icon, and select the **Server Configuration** wizard.
- 4 Click the **Machine** tab and set the server role to **passive**. Do not change the identity of the server (Primary or Secondary).
- 5 Click **Finish**.
- 6 Restart this server.
- 7 Start vCenter Server Heartbeat, if required, and check that the Task bar icon now reflects the changes by showing **P / -** (Primary and Passive) or **S / -** (Secondary and Passive).
- 8 On the active server, right-click the Task bar icon and select the **Server Configuration** wizard.
- 9 Click the **Machine** tab and verify that the server role is set to **active**. Do not change the identity of the server (Primary or Secondary).
- 10 Click **Finish**.
- 11 Restart this server. As the server restarts, it connects to the passive server and starts replication. The active server overwrites data on the passive server.
- 12 Start vCenter Server Heartbeat, if required, and check that the Task bar icon now reflects the changes by showing **P / A** (Primary and active) or **S / A** (Secondary and active).
- 13 Start vCenter Server Heartbeat Console.
- 14 Check that the servers have connected and replication has started.

Two Passive Servers

Primary and Secondary servers are both passive at the same time. This situation is serious and must be resolved immediately.

Symptom

You are unable to connect to protected applications, and if you configured alerts, you receive notification that replication is not functioning properly.

Causes

The condition of two passive servers results from a sudden failure on the active server. Examples:

- An unexpected termination of the Neverfail Server R2 service
- A transient power failure
- A server reset triggered from the Power or Reset button
- An unclean shutdown. Following an unclean shutdown, an active server assumes the passive role to isolate itself from the network until the failure is investigated.
- The active server fails before the handshake that establishes the VMware Channel connection. The passive server cannot detect that the active server is not responding when the failure occurs and cannot determine the condition of the active server. The active server suffers a transient failure and the passive server cannot respond by failing over into the active role, leaving both servers in the passive role.
- Both Primary and Secondary servers experience a power outage simultaneously, for example, they use the same power source and neither is attached to a UPS. A failover cannot occur and when the servers are restarted, each displays the following error message:

Cannot start replication because previous run did not shutdown properly. Check configuration.

NOTE If you attempt to start vCenter Server Heartbeat without reconfiguring one server in the pair as active, vCenter Server Heartbeat responds with the following warning:

[U16] Serious configuration mismatch between the two servers. Please reconfigure so there is one and only one Primary, and one and only one Active.

Resolution

Two passive servers prevent users from accessing the protected application and should be resolved immediately.

To resolve two passive servers

- 1 Determine the active server.
- 2 Shut down vCenter Server Heartbeat on both servers. Leave any protected applications running on the server to make active.
- 3 On the server to make active, start the **Server Configuration** wizard, and select the active role. Do not change the identity (Primary or Secondary).
- 4 On the server to make passive, start the **Server Configuration** wizard, and confirm the passive server. Do not change the identity (Primary or Secondary).
- 5 Restart the passive server. All protected application services stop.
- 6 Start vCenter Server Heartbeat on both servers.

Synchronization Failures

When you start vCenter Server Heartbeat, a full system check occurs to verify the following:

- All protected registry keys and values from the active server are present on the passive server.
- All protected file and folder structures from the active server are present on the passive server.

After the full system check completes, the File System Status and the Registry Status display as **Synchronized**. However, the File System Status or the Registry Status can also display as **Out of sync** or **Synchronized and busy processing**. Some typical scenarios are described with possible causes and workarounds.

Services Running on the Passive Server

Services running on a passive server is not normal behavior and can prevent synchronization.

Symptom

File System Status is **Out of sync** or **Synchronized and busy processing**.

Cause

A service running on the passive server opens a protected file for exclusive access. If vCenter Server Heartbeat attempts to update this opened file, the Apply component logs the following error message:

```
[N29]The passive VMware vCenter Server Heartbeat server attempted to access the file: {filename}.
This failed because the file was in use by another application. Please ensure that there are no
applications which access protected files running on the passive.
```

NOTE This occurs if the vSphere Client is left running on the passive server.

Services that keep files locked on the passive server are:

- Protected application services
- File level antivirus tool services

NOTE vCenter Server Heartbeat periodically checks for and stops any services running on the passive server.

Resolution

Until the file is closed on the passive server, vCenter Server Heartbeat reports the file status and the File System Status as **Out of sync**.

To resolve the Out of sync status

- 1 Set Protected Application services to **Manual** on both servers and verify that they are not running on the passive server.
- 2 Set Recovery Actions to **Take No Action**. You can set this from the Service Control Manager (SCM) for the Protected Application services. Otherwise, the SCM restarts the Protected Application services.
- 3 Verify that file level antivirus is not part of the protected set as the file level antivirus and the corresponding services are running on both machines.

VMware Channel Incorrectly Configured

If the VMware Channel is not properly configured, it cannot initiate the handshake to establish communications through the VMware Channel connection.

Symptom

Failure to establish the VMware Channel connection prevents a full system check, thereby leaving the File System Status and Registry Status as **Out of sync**.

Causes

The most common VMware Channel configuration errors are as follows:

- VMware vCenter Server Heartbeat Packet Filter is enabled on one or more VMware Channel NICs
- VMware Channel IP addresses are configured in different subnets
- In a WAN implementation, no static routes exist between the VMware Channel NICs

Resolution

The VMware Channel configuration should be reviewed to verify proper configuration.

To resolve a VMware Channel configuration error

- 1 Disable the VMware vCenter Server Heartbeat Packet Filter on VMware Channel NICs.
- 2 Configure the VMware Channel IP addresses properly.
- 3 In a WAN implementation, configure static routes between VMware Channel NICs properly.
- 4 Disable NetBIOS on the VMware Channel NICs.

Incorrect or Mismatched Disk Configuration

When vCenter Server Heartbeat starts, it checks the complete set of file filters for consistency.

Symptom

If any of the entries points to a non-existent drive letter or to a non-NTFS partition, the list of file filters resets to the default value of `C:\Protected**`. This is a safety measure as vCenter Server Heartbeat requires the same drive letter configuration on the Primary and the Secondary servers, and only supports protection of NTFS partitions.

Cause

Different partition structures on Primary and Secondary servers, such that one or more file filters point to drives that cannot be protected on both servers. For example:

- The Primary server has drive G, a valid NTFS partition, but no corresponding drive exists on the Secondary server.
- The Primary server has drive G, a valid NTFS partition. The equivalent drive on the Secondary server is a CD or DVD drive, or a FAT or FAT32 partition that cannot be protected.

In either scenario, if you configure a file filter to protect a directory on drive G, the entire filter set is rejected and the filters are reset to the default value of `<Windows drive>\Protected**`.

Resolution

Follow the steps documented in knowledge base article [1008458](#) (*vCSHB-Ref-500*) *Troubleshooting a set of File Filters that is reset to C:\Protected***.

Passive Server Has Less Available Space than Active Server

Inadequate available disk space on the passive server can cause replication to cease.

Symptom

Replication stops with the following error:

```
[N27]Failed to write information for the file: {filename} to the disk. Either the disk is full or the quota (for the SYSTEM account) was exceeded.
```

Cause

The passive server has less available disk space than the active server, preventing updates from being replicated to the passive server. The quantity of updates from the active server exceeds the passive server's available disk space.

Resolution

Free up some additional disk space on the passive server. Do not delete data from the protected set to prevent data loss in the event of a switchover. You could update the disk subsystem on the passive server. After allocating space, start replication.

Registry Status is Out of Sync

The Registry can be reported as **Out of sync** when one or more Registry keys fail to synchronize.

Resource Issues

Inadequate resources can cause poor performance and prevent the registry from synchronizing.

Symptom

vCenter Server Heartbeat logs the following error message:

```
Call to RegOpenKeyEx failed: on <Reg_Key>: Insufficient system resources exist to complete the requested service.
```

Cause

One or both of the servers are running low on virtual memory.

Resolution

Restart the server to free up virtual memory.

Registry Security Issues

Inability to access the registry prevents replication of the registry.

Symptom

vCenter Server Heartbeat is unable to read, sync, or replicate the registry.

Cause

If a protected registry key has permissions that deny Write access to the System account, this can prevent vCenter Server Heartbeat from synchronizing or replicating it.

Resolution

Change the permissions on the affected registry key to grant the System account Full Control.

Channel Drops

When the VMware Channel loses connection between the servers, the following scenarios can occur.

Performance Issues

Poor performance can be experienced as a result of a channel loss.

Symptom

The message `java.io.IOException: An existing connection was forcibly closed by the remote host` appears in the active server's `NFLog.txt` file, and the VMware Channel connection between the servers is lost.

Causes

This unusual condition points to an application or Windows experiencing a fault on the passive server. A sudden restart of the passive server can occur due to the following causes:

- The server is configured for automatic software update management and some updates force the server to restart
- A software or Operating System issue that occasionally fails and requires a system restart
- The Neverfail Server R2 service experiences problems, does not respond, or terminates unexpectedly

Resolution

To resolve the issue, make the following checks.

- Determine the likely source by examining the Windows event logs.
- If the server does not display evidence of a system restart or unresponsive application, one or both of the VMware Channel NICs could be forcing a channel disconnection. See [“Hardware or Driver Issues on VMware Channel NICs”](#) on page 71 for more information on this topic.

Passive Server Does Not Meet Minimum Hardware Requirements

Inadequate hardware can cause channel drops and result in poor performance.

Symptom

The data rate between the servers is very fast during a Full System Check and the VMware Channel drops.

Cause

The passive server does not meet the recommended hardware requirements for vCenter Server Heartbeat or it meets the requirements, but is much less powerful than the active server. The underpowered server cannot apply the received replication data from the active server at the rate that the data is sent to the passive server.

Resolution

To avoid reinstalling vCenter Server Heartbeat, upgrade the hardware, such as memory or CPU, on the passive server. Establish the identity (Primary or Secondary) of the affected server before you perform the upgrade.

Hardware or Driver Issues on VMware Channel NICs

NIC malfunctions and old or incorrect drivers can cause channel drops resulting in poor performance.

Symptom

The VMware Channel intermittently drops or disconnects and reconnects.

Causes

The following are common causes of NIC problems.

- Old or incorrect VMware Channel NIC drivers
- Hardware failure of the hub or Ethernet switch used for the VMware Channel connection
- Defective Ethernet patch or crossover cables

- Improper configuration of the NICs used for the VMware Channel connection
- ISP problems in a WAN environment

Resolution

When a NIC problem is encountered, the following should be checked.

- Verify that VMware Channel NIC drivers are the correct and latest versions. Known issues are identified with HP/Compaq ProLiant NC67xx/NC77xx Gigabit Ethernet NICs. Check other NIC types. See knowledge base article [1008383 \(vCSHB-Ref-116\) – VMware vCenter Server Heartbeat and Gigabit Ethernet NIC drivers. \(NC77XX\)](#).
- Verify hubs and Ethernet switches are operating properly. Identify and replace any defective components.
- Test for defective Ethernet patch or crossover cables and replace if defective.
- Correctly configure the NICs used for the VMware Channel connection.
- Check the physical link for ISP problems.

Firewall Connection

In a LAN or WAN deployment, the VMware Channel can be connected through one or more Internet firewalls. Because firewalls block unauthorized network traffic, configure firewalls on the route of the VMware Channel to allow channel traffic.

Symptoms

The VMware Channel cannot connect, or continuously connects and disconnects.

Causes

In a WAN deployment, port 57348 or any other port configured for the VMware Channel is closed on one or more firewalls on the route between the VMware Channel NIC on the Primary server and its counterpart on the Secondary server.

Resolution

Open port 57348 and any other port configured for the VMware Channel on all firewalls on the route between the VMware Channel NIC on the Primary server and its counterpart on the Secondary server.

Incorrect VMware Channel Configuration

An incorrectly configured channel connection can prevent proper communication and replication.

Symptoms

The following problems are experienced:

- IP conflicts occur on one of the VMware Channel IP addresses
- The VMware Channel does not connect, or connects and disconnects

Causes

The list below provides the most common misconfigurations.

- Identical IP addresses at each end of the VMware Channel
- IP addresses in different subnets without static routing at each end of the VMware Channel
- VMware Channel NIC configured for DHCP when a DHCP server is not available

During installation, vCenter Server Heartbeat configures the VMware Channel NICs with user-provided information. Incorrect information or incorrectly modifying the VMware Channel NIC configuration after installation causes the VMware Channel to fail communicating.

On rare occasions, if the Primary and Secondary servers have NICs of the same type in a different order, both the name and IP address of a VMware Channel NIC on the Primary server can transfer to the Principal (Public) NIC on the Secondary or the name and IP address of the Principal (Public) NIC can transfer to a VMware Channel NIC. Similarly, the names of the VMware Channel NICs can reverse on the Secondary server. You must reconcile the names of the NICs with their physical identities and assign the correct IP address to each NIC on the Secondary server.

Resolutions

The installation process manually assigns the correct IP addresses to each NIC on the Secondary server. If no VMware Channel connection occurs between the servers, verify the configuration of the IP addresses on the Secondary server's channel NICs. Check the settings for the Principal (Public) NIC. The configuration error can remain unrecognized until you perform a switchover or a failover occurs.

To capture the identities of all of the NICs on the Secondary server prior to installing vCenter Server Heartbeat, open a Windows Command Prompt on that server and execute the following command:

```
ipconfig /all > ipconfig.txt
```

The output of this command saves the name, TCP/IP configuration, and MAC address of each NIC on the Secondary server to a file called ipconfig.txt, which is present on that server after the PnP phase of the vCenter Server Heartbeat install completes. Compare the pre-install and post-install state of each NIC by running `ipconfig /all` from a Windows command prompt and compare the output of this command with the content of ipconfig.txt.

The MAC address of each NIC is connected to the physical identity of each card and never changes. You can identify each NIC by its MAC address and determine its original name and network configuration, even if this was updated by the PnP process.

VMware vCenter Server Heartbeat Packet Filter Is Enabled on the Channel NIC(s)

Proper configuration requires that the packet filter be disabled on the VMware Channel NIC. When the packet filter is enabled on the channel NICs, the following symptoms are encountered.

Symptom

Interference with network traffic across the VMware Channel results in an intermittent channel connection or no channel connection at all.

Cause

During installation, the VMware vCenter Server Heartbeat Packet Filter is installed and enabled on all NICs on both the Primary and Secondary servers. The Packet Filter on the VMware Channel NICs on each server is disabled later in the installation of vCenter Server Heartbeat. If the vCenter Server Heartbeat Packet Filter is left enabled on one or more channel NICs after installation completes, it can interfere with network traffic across the VMware Channel.

Resolution

Click the **Properties** tab for each Channel NIC on both servers and verify that the check box for **vCenter Server Heartbeat Packet Filter** is cleared, so that the Packet Filter is disabled on that NIC.

Subnet or Routing Issues

In a LAN or WAN deployment, the following connection problems can occur.

LAN Deployment

Incorrectly configured subnets or routing can cause channel problems resulting in poor performance or failure to connect.

Symptom

The Channel disconnects or fails to connect in a LAN deployment.

Causes

The Channel disconnects or fails to connect due to the Principal (Public) NIC and/or one or more channels sharing the same subnet.

Resolution

If vCenter Server Heartbeat is deployed in a LAN environment, the Principal (Public) IP address and the VMware Channel IP address on a server must be in separate subnets. When multiple redundant channels are present, each must have its own subnet. Check the network configuration for each NIC on both servers in the pair and correct any issues.

WAN Deployment

Incorrect routing can prevent the active and passive servers from connecting in a WAN environment.

Symptom

The VMware Channel disconnects or fails to connect in a WAN deployment.

Cause

When the VMware Channel disconnects or fails to connect in a WAN deployment, the static route might not be configured or might be configured incorrectly.

When vCenter Server Heartbeat is deployed in a WAN, the Principal (Public) IP address and the VMware Channel IP addresses cannot be in different subnets, because there usually is a single network path between the two servers. Configure a static route between the endpoints to route traffic in the VMware Channel.

Resolution

Refer to knowledge base article [1008451](#) (*vCSHB-Ref-466*) - *Creating a static route for the VMware Channel Connection in a WAN Environment* for a detailed discussion about WAN channel routing issues, and for instructions on configuring a static route for the VMware Channel.

MaxDiskUsage Errors

vCenter Server Heartbeat uses queues to buffer the flow of replication data from the active server to the passive server. This configuration provides resilience in the event of user activity spikes, VMware Channel bandwidth restrictions, or VMware Channel drops across a WAN deployment. Some types of file write activity also require buffering as they can cause a sharp increase in the amount of channel traffic. The queues are called the send queue (on the active server) or the receive queue (on the passive server).

Send Queue

vCenter Server Heartbeat considers the active server's send queue as unsafe because the data in this queue has not yet been replicated across the VMware Channel to the passive server and therefore could be lost in the event of a failover. As a result of failover, some data loss is inevitable, with the exact amount depending on the relationship between available VMware Channel bandwidth and the required data transmission rate. If the required data transmission rate exceeds available VMware Channel bandwidth, the send queue fills. If the available VMware Channel bandwidth exceeds the required data transmission rate, the send queue empties. This situation is most commonly seen in a WAN environment, where VMware Channel bandwidth is restricted. In a LAN that normally has high bandwidth on a dedicated channel, the size of the send queue is zero or near zero most of the time. On a server not protected with vCenter Server Heartbeat, all data is technically unsafe and subject to loss if the server fails.

Receive Queue

The passive server's receive queue is considered safe because the data in this queue already was transmitted across the VMware Channel from the active server, and is not lost in the event of a failover, which applies all updates to the passive server as part of the process.

Both send and receive queues are stored on disk by default in the <VMware vCenter Server Heartbeat Install Directory>\R2\log directory, with a quota configured for the maximum permitted queue size (by default, 1GB on each server). You can configure both the queue location and the quota.

Two methods to set the queue size:

- Using vCenter Server Heartbeat Console
 - a Start vCenter Server Heartbeat
 - b Open the **vCenter Server Heartbeat Console**, and select **Data: Traffic Queues**.
 - c Click **Configure**.
 - d Set the **Allow a maximum** value and click **OK**.
 - e Shut down and restart vCenter Server Heartbeat for the change to take effect. You are not required to stop protected applications.
- Using the Server Configuration wizard
 - a Shut down vCenter Server Heartbeat.
 - b Open the **Server Configuration** wizard and click the **Logs** tab.
 - c Set the **Maximum Disk Usage** value and click **Finish**.
 - d Start vCenter Server Heartbeat.

NOTE vCenter Server Heartbeat is a symmetrical system and can operate with either server in the active role. For this reason, the queue size is always set to the same value for both servers.

MaxDiskUsage Error Messages

The following error messages display when available disk space on the servers is exceeded.

[L9]Exceeded the Maximum Disk Usage (VCChannelExceededMaxDiskUsageException)

This message indicates that you have exceeded the amount of allocated disk space reserved for the queue.

Symptom

vCenter Server Heartbeat exceeds its preconfigured queue size.

Causes

On the active server, the size of the active server queue has exceeded the disk quota allocated for it. On the passive server, the size of the passive server queue has exceeded the disk quota allocated for it.

Resolution

While neither condition is critical, determine the sequence of events that led to the condition.

[L9]Exceeded the Maximum Disk Usage on the ACTIVE Server

This message indicates that you have exceeded the amount of allocated disk space reserved for the active server's send queue.

Symptom

Replication stops and the vCenter Server Heartbeat Event Log displays the error message originating from the active server.

Causes

A temporary interruption in the VMware Channel, or insufficient VMware Channel bandwidth to support the volume of replication traffic starts filling the active server queue. The size of the queue eventually exceeds the configured disk quota.

Resolution

Assuming no other channel connection issues exist (see knowledge base article [1008551](#) (*vCSHB-Ref-992*) - *Troubleshooting VMware vCenter Server Channel Drops*), you can increase the amount of disk space allotted to the queues. The default setting is 1GB, which can be insufficient on servers with a large volume of replication traffic and limited VMware Channel bandwidth. If you have sufficient disk space, set the queue size to zero (unlimited) so vCenter Server Heartbeat can use any free disk space to store the queues.

[L9]Exceeded the Maximum Disk Usage on the PASSIVE Server

This message indicates that you have exceeded the amount of allocated disk space reserved for the passive server's receive queue.

Symptom

Replication stops and the vCenter Server Heartbeat Event Log displays the error message originating from the passive server.

Causes

Two of the most common causes are shown below:

- The bottleneck lies between the VMware Channel NIC and the disk subsystem on the passive server. Replication traffic passes across the VMware Channel faster than it can be written to disk on the passive server. The excess is buffered temporarily in the passive server's receive queue. The size of the queue can eventually exceed the allotted disk quota.
- If the passive server is much less powerful than the active server in terms of processor speed, RAM, or disk performance, it can lag behind the active server during periods of high replication activity. Monitor one or more Windows performance counters to determine the component experiencing sustained high activity. Intensive page file use or persistently large disk queue length can indicate a problem. Upgrade one or more physical components of the server.

Either server can be active or passive. If the Secondary server is more powerful than the Primary server, hardware-related issues can only occur while the Secondary server is in the active role.

Resolution

To resolve this issue:

- If you have multiple physical disks on each server, locate the vCenter Server Heartbeat send and receive queues on a separate physical disk, away from the Windows directory, the Windows page file, and any protected files help to alleviate disk performance issues:
 - a Shut down vCenter Server Heartbeat.
 - b Open the **Server Configuration** wizard and click the **Logs** tab.
 - c Set the path for **Message Queue Logs Location** and click **Finish**.
 - d Start vCenter Server Heartbeat on both servers.

The selected path is applied to all vCenter Server Heartbeat queues on both servers.

- Increase the amount of disk space allotted to the queues. However, if a hardware issue is the root of the problem, correct that problem at the source.
- The size of the passive server's receive queue can increase sharply in response to certain types of file write activity on the active server, such as when vCenter Server Heartbeat is replicating a large number of very small updates of a few bytes each. The volume of update traffic can be far greater than the physical size of the files on the disk, and the receive queue can become disproportionately large. You can see this pattern of disk activity during the population of Full-Text Catalogs in Microsoft SQL Server. Increase the amount of disk space available for the queues. Move the queues to their own physical disk, upgrade the memory or the disk subsystem.
- vCenter Server Heartbeat requires a certain amount of system resource for its own basic operations and requires some additional resources for processing replication traffic. This is in addition to the resources used by Windows and other applications running on the server, including critical applications protected by Heartbeat. Allocate sufficient resources for all the applications and services running on such a server to provide maximum performance, stability, and resilience for changing client, server, and network activity.

[L20]Out of Disk Space (VCChannelOutOfDiskSpaceException)

This message indicates that one of the servers in the pair has run out of disk space without reaching its preset quota.

Symptom

Replication stops and the vCenter Server Heartbeat Event Log displays the error message originating from either server in the pair.

Cause

One of the queues has exceeded the amount of physical disk space available for it without reaching its quota limit. For example, if the maximum queue size is set to 5GB, but only 3GB of physical disk space remains, this error message is reported if one of the queues exceeds 3GB in size.

Resolution

Free up more disk space or move the queues to a disk with sufficient free space to accommodate queue sizes up to the limit configured for Maximum Disk Usage.

Application Slowdown

Operations performed by the application can take longer to complete, and in turn, can affect the time required to log in to a remote client, or to open or save a file. This is true for both servers running vCenter Server Heartbeat and for servers running any other application. vCenter Server Heartbeat can monitor system performance counters and display warnings when predefined thresholds are exceeded, but it does not actively manage system resources for other applications. Like any other application, it also requires a finite amount of resources for its own operations in addition to the resources used by the operating system and the protected application.

The machines hosting vCenter Server Heartbeat must meet recommended hardware requirements and must be powerful enough to support the load, the protected applications, and any other critical applications running on the same server pair.

Poor Application Performance

When applications are competing for resources, one or more applications can perform poorly.

Symptom

Neither server in the pair can accommodate the load placed upon it during normal operation.

Cause

The Primary server's resource usage in one or more areas reached close to the maximum before vCenter Server Heartbeat was installed.

Resolution

Heartbeat Diagnostics can report these conditions and issues warnings if CPU usage or memory usage exceed a certain percentage of the available resource. Information provided by Heartbeat Diagnostics can minimize the risk of application slowdown by identifying needed hardware upgrades on the Primary server.

Both Servers Can Accommodate the Initial Load but the Load Has Increased

Any software installed on a server or workstation consumes a finite amount of system resources when it runs and it must share the resources it uses with any other applications running at the same time. Increased demand caused by additional user activity can have an impact on the server performance.

Symptom

Increased user activity slows application response time.

Causes

The server pair operates normally when vCenter Server Heartbeat is first installed, but performance decreases due to increased user activity. For example, users on the SQL Server system increase or the typical usage pattern becomes more intense. This can be a gradual and sustained increase over time, or transient if a specific event triggers a temporary surge in user activity.

Resolution

If the situation is sporadic, it can correct itself when the load decreases. If the increase is sustained and permanent, upgrade the server hardware.

One Server Can Provide Adequate Resource Support, but the Other Cannot

If the total resource requirements of the applications exceed the available physical resources, the operating system attempts to provide resources, but leaves some applications under resourced. When this situation occurs, an application cannot obtain enough memory to operate normally, or a process must wait before accessing the hard disk

Symptom

Applications operate normally when the Primary server is active but operate slowly when the Secondary server is active (or the reverse).

Cause

A large discrepancy occurs in the processing power between the Primary and Secondary servers. One server can handle the operational load while the other cannot. The load on a server is greater while in the active role when the protected application starts. Applications on the server pair run successfully when the Primary server is active, but experience performance issues when the Secondary is active (or the reverse). Problems can arise even when the more powerful server is active.

Resolution

Both servers must have approximately equivalent processing power, RAM and disk performance. Upgrade the hardware on one server in the pair so that the two servers have roughly the same performance.

Scheduled Resource Intensive Tasks

Scheduling multiple resource intensive tasks at the same time can adversely impact server performance and affect application performance.

Symptom

Resource-intensive scheduled tasks impact performance at certain times.

Cause

Two or more resource-intensive processes run simultaneously or one process performs actions that increase the load on vCenter Server Heartbeat by triggering additional and sometimes unnecessary replication traffic. Examples: processes such as backups, database maintenance tasks, disk defragmentation, or scheduled virus scans.

Resolution

Schedule operations so that they do not overlap and schedule them outside regular working hours, when fewer users are accessing the protected application and consequently less load on the server.

Glossary

A **Active**

The functional state or role of a server visible through the network by clients running protected applications and servicing client requests.

Alert

A notification sent to a user or entered into the system log indicating an exceeded threshold.

Active Directory (AD)

Presents applications with a single, simplified set of interfaces so users can locate and use directory resources from a variety of networks while bypassing differences among proprietary services. vCenter Server Heartbeat switchovers and failovers require no changes to AD, resulting in switchover and failover times measured in seconds.

Active – Passive

The coupling of two servers: one server visible to clients on a network and providing application service, the other server not visible and not providing application service.

Active Server Queue

The staging area of the active server used to store intercepted data changes before being transported across the VMware Channel to the passive server.

Advanced Configuration and Power Interface (ACPI)

A specification that dictates how the operating system can interact with hardware using power saving schemes. Primary and Secondary servers must have the same ACPI compliance.

Asynchronous

A process whereby replicated data is applied (written) to the passive server independently of the active server.

B **Basic Input/Output System (BIOS)**

The program a personal computer's microprocessor uses to start the computer system after you turn it on. It also manages data flow between the computer's operating system and attached devices such as the hard disk, video adapter, keyboard, mouse, and printer.

C **Cached Credentials**

Locally stored security access credentials used to log in to a computer system when a Domain Controller is not available.

Channel Drop

An event in which the dedicated communications link between the Primary and Secondary server fails, often resulting in the passive server becoming active and consequently creating a split-brain syndrome.

Channel NIC (Network Interface Card)

A dedicated subnet used by the VMware Channel.

Cloned Servers

Two servers in a pair with the same configuration settings, names, applications, Security Identifiers (SIDs) and IP addresses, following the installation of vCenter Server Heartbeat.

Cloning Process

The vCenter Server Heartbeat process whereby all installed applications, configuration settings, the machine name, security identifier (SID), and IP address are copied to a second server.

Crossover Cable

A network cable that crosses transmit and receive lines.

D Data Replication

The transmission of protected data changes (files and registry) from the active to the passive server through the VMware Channel.

Device Drivers

A program that controls a hardware device, linking it to the operating system.

Disaster Recovery (DR)

A term indicating how you maintain and recover data in light of a disaster such as a hurricane or fire. vCenter Server Heartbeat achieves DR protection by placing the Secondary server at an offsite facility and replicating the data through a WAN link.

DNS (Domain Name System) Server

Responsible for providing a centralized resource for clients to resolve NetBIOS names to IP addresses.

Domain

A logical group of client server based machines where the administration rights across the network are maintained in a centralized resource called a domain controller.

Domain Controller (DC)

The server responsible for maintaining privileges to domain resources, sometimes called AD controller in Windows 2000 and above domains.

F Failover

The process by which the passive server assumes the active role when it no longer detects that the active server is alive as a result of a critical unexpected outage or server crash.

Full System Check (FSC)

The internal process programmatically started at the initial connection of a server pair or manually triggered through the vCenter Server Heartbeat Console. The FSC verifies the files and registry keys, and synchronizes the differences.

Fully Qualified Domain Name (FQDN)

Also known as an absolute domain name, a FQDN specifies its exact location in the tree hierarchy of the Domain Name System (DNS). It specifies all domain levels, including the top-level domain, relative to the root domain. Example: `somehost.example.com.`, where the trailing dot indicates the root domain.

G Graceful (Clean) Shutdown

vCenter Server Heartbeat shuts down with no data loss after completing replication using the vCenter Server Heartbeat Console.

- H** **Hardware Agnostic**
A key vCenter Server Heartbeat feature enabling the use of servers from different manufacturers, models, and processing power in a single vCenter Server Heartbeat server pair.
- Heartbeat**
The packet of information issued by the passive server across the VMware Channel, which the active server responds to, indicating its presence.
- Heartbeat Diagnostics**
The umbrella name for the VMware process and tools used to check the production server health and applicability to the implementation of the vCenter Server Heartbeat solution.
- High Availability (HA)**
Keeping users seamlessly connected to their applications, regardless of the nature of a failure. LAN environments are ideally suited for HA.
- Hotfix**
A single, cumulative package that includes one or more files used to address a problem in a product.
- I** **Identical Nodes**
The use of two servers identical in name, IP address, and security identifier (SID).
- Identity**
The reference of a server's position in the server pair based upon hardware, either the Primary server or the Secondary server.
- L** **Low Bandwidth Module (LBM)**
A vCenter Server Heartbeat Module that compresses and optimizes data replicated between a Primary and Secondary server, thereby delivering maximum data throughput and improving application response time on congested WAN links.
- M** **Machine Name**
The Windows or NETBIOS name of a computer.
- Management IP Address**
An additionally assigned unfiltered IP address used for server management purposes only.
- Many-to-One**
One physical Secondary server (hosting more than one virtual server) can provide protection to multiple physical Primary servers.
- N** **Network Monitoring**
Monitoring the active server's capability to communicate with the rest of the network by polling defined nodes around the network at regular intervals.
- Non-Identical Nodes**
Two servers in a cluster using differing names, Passive IP addresses, and security identifiers (SIDs).
- P** **Passive**
The functional state or role of a server that is not delivering service to clients and is hidden from the rest of the network.
- Passive Server Queue**
The staging area on the passive server used to store changes received from the active server before they are applied to the passive server's disk or registry.

Pathping

A route-tracing tool that sends packets to each router on the way to a final destination and displays the results of each hop.

Plug and Play (PnP)

A standard for peripheral expansion on a PC. When starting the computer, Plug and Play (PnP) configures the necessary IRQ, DMA and I/O address settings for the attached peripheral devices.

Plug-in

An optional module that can be installed into a vCenter Server Heartbeat server to provide additional protection for a specific application.

Pre-Installation Checks

A list of system and environmental checks performed before the installation of vCenter Server Heartbeat.

Principal IP address

An IP address used by clients to contact the server through drive mappings, UNC paths, DNS resolved paths, to access the server's services and resources.

Principal NIC

The network card that hosts the Principal IP address.

Protected Application

An application protected by vCenter Server Heartbeat.

Q**Quality of Service (QoS)**

An effort to provide different prioritization levels for different types of traffic over a network. For example, vCenter Server Heartbeat data replication can have a greater priority than ICMP traffic, as the consequences of interrupting data replication are more obvious than slowing down ICMP traffic.

R**Remote Desktop Protocol (RDP)**

This multi-channel protocol connects to a computer running Microsoft Terminal Services.

Replication

The generic term given to the process of intercepting changes to data files and registry keys, transporting the changed data across the VMware Channel, and applying them to the passive server so both servers are maintained in a synchronized state.

Role

The functional state of the server in the pair that can be either active or passive.

Rule

A set of actions vCenter Server Heartbeat to perform when defined conditions are met.

S**Security Identifier (SID)**

A unique alphanumeric character string that identifies each operating system and each user in a network of Windows NT, Windows Server 2000, Windows Server 2003, and Windows Server 2008 systems.

Server Monitoring

Monitoring the active server by the passive server, using a heartbeat message, to ensure that the active server is functional.

Server Pair

The generic term used to describe the coupling of the Primary and Secondary server in vCenter Server Heartbeat.

Shared Nothing

A key vCenter Server Heartbeat feature whereby hardware is not shared between the Primary and Secondary servers, thus preventing a single point of failure.

SMTP

A TCP/IP protocol used in sending and receiving e-mail between or among servers.

Split-brain Avoidance

A unique feature of vCenter Server Heartbeat that uses various checks to overcome a scenario where both Primary and Secondary servers attempt to become active at the same time, leading to an active-active rather than an active-passive model.

Split-brain Syndrome

A situation where both the Primary and Secondary servers in a vCenter Server Heartbeat server pair are operating in the active mode and attempting to service clients, causing different data updates to be applied independently to each server.

SSL

(Secure Sockets Layer) establishes a secure session by electronically authenticating each end of an encrypted transmission.

Subnet

A division of a network into an interconnected but independent segment or domain, to improve performance and security.

Storage Area Network (SAN)

A high-speed special-purpose network or (sub-network) that interconnects different kinds of data storage devices with associated data servers on behalf of a larger network of users.

Switchover

The graceful transfer of control and application service to the passive server.

Synchronize

The internal process of transporting 64KB blocks of changed files or registry key data, through the VMware Channel from the active server to the passive server. The data on the passive server is a mirror image of the protected data on the active server, a required condition for data replication on a vCenter Server Heartbeat server pair.

System State

Data that comprises the registry, COM+ Class Registration database, files under Windows File Protection, and system boot file. Other data can be included in the system state data.

T**Task**

An action performed by vCenter Server Heartbeat when defined conditions are met.

Time-To-Live (TTL)

The length of time that a locally cached DNS resolution is valid. The DNS server must be re-queried after the TTL expires.

Traceroute

A utility that records the route through the Internet between the computer and a specified destination computer.

U**Ungraceful (Unclean) Shutdown**

A shutdown of vCenter Server Heartbeat resulting from a critical failure or by shutting down Windows without first performing a proper shutdown of vCenter Server Heartbeat, resulting in possible data loss.

Unprotected Application

An application that is not monitored or its data replicated by vCenter Server Heartbeat.

V VMware Channel

The IP communications link used by vCenter Server Heartbeat for heartbeat and replication traffic.

VMware vCenter Server Heartbeat

The core replication and system monitoring component.

VMware vCenter Server Heartbeat Packet Filter

The network component installed on both servers that controls network visibility.

VMware vCenter Server Heartbeat Switchover and Failover Process

A vCenter Server Heartbeat unique process whereby the passive server gracefully (Switchover) or unexpectedly (Failover) assumes the role of the active server providing application services to connected clients.

Virtual Private Network (VPN)

A private data network that uses the public telecommunication infrastructure, maintaining privacy through the use of a tunneling protocol and security procedures.

VMware Web Site

The VMware web site dedicated to support partners and customers providing technical information, software updates, and license key generation.

W Windows Management Instrumentation (WMI)

A management technology using scripts to monitor and control managed resources throughout the network. Resources include hard drives, file systems, operating system settings, processes, services, shares, registry settings, networking components, event logs, users, and groups.