

Administrator's Guide

For
Neverfail Heartbeat v6.7



You can find the most up-to-date technical documentation on the Neverfail Extranet at:

<http://extranet.neverfailgroup.com>.

The Neverfail Extranet also provides the latest product updates. If you have comments about this documentation, submit your feedback to:

docfeedback@neverfailgroup.com

The Neverfail Group has taken all reasonable care to ensure the information in this document is accurate at the date of publication. In relation to any information on third party products or services, the Neverfail Group has relied on the best available information published by such parties. The Neverfail Group is continually developing its products and services, therefore the functionality and technical specifications of the Group's products can change at any time. For the latest information on the Neverfail Group's products and services, please contact us by email (info@neverfailgroup.com) or visit our Web site www.neverfailgroup.com).

Heartbeat is a product trade mark of the Neverfail Group Ltd. Neverfail products are protected, in whole or in part by U.S. and foreign patents, which include US. Patent No. 7,409,577 and 7,788,524 and European Patent No. 1,397,744.

All third party product names referred to in this document are acknowledged as the trade marks for their respective owner entities.

© 2002-2013 Neverfail Group Ltd. All rights reserved.

Contents

Preface: About This Book.....	vii
Part I: Getting Started.....	9
Chapter 1: Neverfail Heartbeat Concepts.....	11
Overview.....	11
Architecture.....	11
Protection.....	12
Neverfail Heartbeat Networking Configuration.....	13
Neverfail Heartbeat Communications.....	15
Neverfail Heartbeat Switchover and Failover Processes.....	16
Recovery from a Failover.....	19
Chapter 2: Configuring Neverfail Heartbeat.....	21
Configure Server Wizard.....	21
Configure Machine Identity.....	22
Configure Server Role.....	22
Change the Client Connection Port.....	22
Configure Channel IP Routing.....	23
Configure the Default Channel Port.....	23
Configure Low Bandwidth Optimization.....	24
Configure Public IP Addressing.....	24
Add/Remove an Neverfail Heartbeat License Key.....	25
Configure the Message Queue Logs.....	26
Configure Maximum Disk Usage.....	27
Chapter 3: Status and Control.....	29
Managing Neverfail Heartbeat Clusters.....	29
Starting the Neverfail Heartbeat Management Client	30
Navigating the Neverfail Heartbeat Management Client	31
Changing the Neverfail Heartbeat Management Client Font.....	33
Add a Neverfail Heartbeat Group.....	34
Remove an Neverfail Heartbeat Group.....	35
Add a New Connection.....	35
Edit a Connection.....	37
Move a Connection.....	38
Remove a Connection.....	39
Business Application Groups.....	39
Edit User Name and Passwords.....	50
Review the Status of Neverfail Heartbeat Clusters and Groups.....	51
Exit Neverfail Heartbeat Management Client	52
Shutdown Windows with Neverfail Heartbeat Installed.....	52
Controlled Shutdown.....	52
Neverfail vSphere Client.....	53
Neverfail SRMXtender.....	64
Part II: Protection.....	69

Chapter 4: Server Protection.....	71
Overview.....	71
Monitoring the Status of Servers.....	71
Configure Neverfail Heartbeat Settings.....	72
Common Administrative Tasks in Neverfail Heartbeat	75
Forcing a Switchover.....	77
Failover versus Switchover.....	77
Split-brain Avoidance.....	82
Chapter 5: Network Protection.....	85
Overview.....	85
Configure Public Network Monitoring.....	86
Chapter 6: Application Protection.....	89
Applications Environment.....	89
Applications: Summary.....	89
Applications: Services.....	95
Applications: Tasks.....	99
Application: Rules.....	103
Applications: Plug-ins.....	103
Chapter 7: Performance Protection.....	107
Rules Overview.....	107
Check a Rule Condition.....	108
Edit a Rule.....	109
Chapter 8: Data Protection.....	111
Data: Replication.....	111
Data: File Filters.....	118
Data: Rollback.....	121
Data: Traffic/Queues.....	129
Data: WAN Compression.....	131
Part III: Reference.....	135
Appendix A: Other Administrative Tasks.....	137
Configure Alerts.....	137
Configure Alert Reporting.....	138
Add and Remove Email Alert Recipients.....	139
Using WScript to Issue Alert Notifications.....	140
Test Alert Reporting.....	141
Configure Event Log Files.....	141
Review Event Logs.....	142
Appendix B: Troubleshooting.....	145
Two Active Servers.....	145
Two Passive Servers (In a Pair).....	146

Three Passive Servers (In a Trio).....	147
Invalid Neverfail Heartbeat License.....	148
Synchronization Failures.....	149
Services Running on the Passive Server.....	149
Neverfail Channel Incorrectly Configured.....	150
Incorrect or Mismatched Disk Configuration.....	151
The First Passive Server has Less Available Space than the Active Server.....	151
Unprotected File System Features.....	152
Registry Status is Out-of-Sync.....	152
Channel Drops.....	153
Performance Issues.....	153
Passive Server Does Not Meet Minimum Hardware Requirements.....	153
Hardware or Driver Issues on Channel NICs.....	154
Firewall Connection.....	154
Incorrect Neverfail Channel Configuration.....	155
Neverfail Heartbeat Packet Filter is Enabled on the Channel NIC(s).....	156
Subnet/Routing Issues — In a LAN.....	156
Subnet/Routing Issues — In a WAN.....	156
MaxDiskUsage Errors.....	157
[L9]Exceeded the Maximum Disk Usage on the ACTIVE Server.....	158
[L9]Exceeded the Maximum Disk Usage on a PASSIVE Server.....	158
[L20]Out of disk space (NFChannelOutOfDiskSpaceException).....	159
Application Slowdown.....	160
Poor Application Performance.....	160
Servers Could Accommodate the Initial Load but the Load has Increased.....	161
One Server is Able to Cope, but the Other Cannot.....	161
Scheduled Resource Intensive Tasks.....	161
Glossary.....	163

About This Book

The Administrator Guide provides information about configuring and performing the day-to-day management of Neverfail Heartbeat when deployed in a Pair or Trio over a Local Area Network (LAN) or Wide Area Network (WAN). Additionally, this guide provides information about configuring network protection, application protection, data protection, split-brain avoidance, and more. To help you protect your applications, this guide provides an overview of the protection offered by Neverfail Heartbeat and the actions that Neverfail Heartbeat can take in the event of a network, hardware, or application failure.

Intended Audience

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

Using the Administrator's Guide

This guide is designed to provide information related to the daily management of your Neverfail Heartbeat Cluster after successful installation. To help you protect your applications, this guide provides an overview of the protection offered by Neverfail Heartbeat and the actions that Neverfail Heartbeat can take in the event of a network, hardware, or application failure. The information contained in this guide is current as of the date of printing.

Overview of Content

This guide is designed to give guidance on the configuration and administration of Neverfail 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 — Neverfail Heartbeat Concepts presents an overview of Neverfail Heartbeat architecture and the five levels of protection provided by Neverfail Heartbeat.

Chapter 2 — Configuring Neverfail Heartbeat discusses how to configure Neverfail Heartbeat using the *Configure Server Wizard*.

Chapter 3 — Status and Control describes how to connect to Neverfail Heartbeat using the Neverfail Heartbeat Management Client and the Virtual Availability Director to review the status of and manage a Cluster.

Chapter 4 — Server Protection discusses how the Neverfail Heartbeat solution protects users from server system failure or server hardware crash.

Chapter 5 — Network Protection describes how Neverfail Heartbeat protects against network failure by ensuring that the network identity of the production server, IP address, etc. are provided to users.

Chapter 6 — Application Protection discusses how Neverfail Heartbeat maintains the protected application environment ensuring that applications and services stay alive on the network.

Chapter 7 — Performance Protection describes how Neverfail Heartbeat monitors system and application attributes to prevent an unexpected system or application failure.

Chapter 8 — Data Protection discusses how Neverfail Heartbeat intercepts all data written by users and protected applications and maintains a copy of this data for use in case of failure.

Appendix A — Other Administrative Tasks discusses additional tasks for the administrator to configure system logging and alerting functions.

Appendix B — Troubleshooting discusses common issues that may appear and techniques to troubleshoot the issue and includes two active servers, two passive servers in a pair and three passive servers in a trio, application slowdown, channel drops, and MaxDiskUsage errors.

Document Feedback

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

Abbreviations Used in Figures

The figures in this book use the abbreviations listed in the table below.

Table 1: Abbreviations

Abbreviation	Description
Channel	Neverfail Channel
NIC	Network Interface Card
P2P	Physical to Physical
P2V	Physical to Virtual
V2V	Virtual to Virtual
SAN	Storage Area Network type datastore

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 <http://www.neverfailgroup.com/services/technical-support.html>.

Online and Telephone Support

Use online support to view your product and contract information, and to submit technical support requests. Go to <http://www.neverfailgroup.com/services/technical-support.html>.

Support Offerings

To find out how Neverfail Support offerings can help meet your business needs, go to <http://www.neverfailgroup.com/services/technical-support.html>.

Neverfail Professional Services

Neverfail Professional 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 on site, in the classroom, and live online. For the day-to-day operations of Neverfail Heartbeat, Neverfail Professional Services provides offerings to help you optimize and manage your Neverfail Heartbeat servers. To access information about education classes, certification programs, and consulting services, go to <http://www.neverfailgroup.com/services/professional-services.html>.

Getting Started

Chapter 1

Neverfail Heartbeat Concepts

Overview

Neverfail Heartbeat is a Windows based system specifically designed to provide [High Availability \(HA\)](#) and [Disaster Recovery \(DR\)](#) to server configurations in one solution that does not require any specialized hardware. To appreciate the full capabilities of Neverfail Heartbeat you must understand the basic concepts under which Neverfail Heartbeat operates and the terminology used.

Note: In this document, the term “Cluster” refers to a Neverfail Heartbeat Cluster. Refer to the Glossary for more information about Neverfail Heartbeat Clusters.

Architecture

Neverfail Heartbeat provides a flexible solution that can be adapted to meet most business requirements for deployment. When discussing Neverfail Heartbeat, it is beneficial to use common terminology to describe its configuration and what is occurring.

Neverfail describes the organization of servers based upon Clusters, Cluster status, and relationships between Clusters. Neverfail refers to a Cluster of two servers as a Neverfail Heartbeat Pair and three servers as a Neverfail Heartbeat Trio. Installing Neverfail Heartbeat on the servers and assigning an Identity (Primary, Secondary, or Tertiary) to the servers results in a Neverfail Heartbeat Pair or a Neverfail Heartbeat Trio.

Each server is assigned both an [Identity](#) ([Primary](#) , [Secondary](#) , or [Tertiary](#)) and a [Role](#) ([Active](#) or [Passive](#)). Identity is used to describe the physical instance of the server while the Role is used to describe what the server is doing. When the Identity is assigned to a server it normally will not change over the life of the server whereas the Role of the server is subject to change as a result of the operations the server is performing. When Neverfail Heartbeat is deployed on a Pair of servers, Neverfail Heartbeat can provide all five levels of protection (Server, Network, Application, Performance, and Data) and can be deployed for High Availability in a Local Area Network (LAN) or Disaster Recovery over a Wide Area Network (WAN).

In its simplest form, Neverfail Heartbeat operates as a Neverfail Heartbeat Pair with one server performing an active role (normally the Primary server) while the other server performs a passive role (normally the Secondary server). The server in the active role provides application services to users and serves as the source for replication while the server in the passive role serves as the standby server and target for replicated data. This configuration supports replication of data between the active and passive server over the Neverfail Channel.

When more than one passive server exists in the Cluster (such as a Trio), Neverfail refers to them as first passive and second passive. The first passive server is the server that is receiving replication data from the active server and sending replication data to the second passive server. When deployed in

a Trio, any of the servers can operate as active or passive but the Tertiary server can never operate as a first passive receiving data from the active and replicating it to the second passive.

When deployed for High Availability, a LAN connection is used. Due to the speed of a LAN connection (normally 100 Mb or more) bandwidth optimization is not necessary.

When deployed in a WAN for Disaster Recovery, Neverfail Heartbeat can assist replication by utilizing WAN Compression with the built-in WAN Acceleration feature.

When deployed in the Trio configuration, Neverfail Heartbeat provides continuous availability of protected applications to end-users in the event of hardware failure or an operating system crash on the local network, and provides Disaster Recovery with a third server located at a remote DR site. The Trio configuration allows the Primary and Secondary servers to operate in a LAN providing High Availability while the Secondary and Tertiary servers operate over a WAN to provide Disaster Recovery. This architecture uses a daisy chain where the active server replicates to the first passive (normally the Secondary) server and the first passive server replicates to the second passive (normally the Tertiary) server.

Protection

Neverfail Heartbeat provides five levels of protection to ensure that end-user clients remain connected in the event of a failure.

- **Server Protection** — Neverfail Heartbeat continues to provide availability to end-user clients in the event of a hardware failure or operating system crash. When deployed, Neverfail Heartbeat provides the ability to monitor the active server by sending “I’m alive” messages from the first passive server to the active server which reciprocates with an acknowledgment over a network connection referred to as the Neverfail Channel. Should the first passive server detect that the process or heartbeat has failed, it can then initiate a failover.

A failover occurs when the first passive server detects that the active server is no longer responding. This can be because the active server’s hardware has crashed or because its network connections are lost. Rather than the active server being gracefully closed, it has been deemed to have failed and requires no further operations. In a failover, the first passive server is brought up immediately to take on the role of the active server. The mechanics of failover are discussed later in this guide.

- **Network Protection** — Neverfail Heartbeat proactively monitors the ability of the active server to communicate with the rest of the network by polling up to three defined nodes around the network, including by default, the default gateway, primary DNS server, and the Global Catalog server at regular intervals. If all three nodes fail to respond, for example, if a network card or local switch fails, Neverfail Heartbeat can gracefully switch the roles of the active and first passive servers (referred to as a switchover) allowing the previously first passive server to assume an identical network identity to that of the previously active server. After the switchover, the newly active server then continues to service the clients.

- **Application Protection** — Neverfail Heartbeat running on the active server locally monitors the applications and services it has been configured to protect through the use of plug-ins. If a protected application should fail, Neverfail Heartbeat will first try to restart the application on the active server. If a restart of the application fails, then Neverfail Heartbeat can initiate a switchover.

A switchover gracefully closes down any protected applications that are running on the active server and restarts them on the first passive server along with the application or service that caused the failure. The mechanics of switchover are discussed in more detail later in this guide.

- **Performance Protection** — Neverfail Heartbeat proactively monitors system performance attributes to ensure that your protected applications are actually operational and providing service to your end users, and that the performance of those applications is adequate for the needs of those users.

Neverfail Heartbeat Plug-ins provide these monitoring and preemptive repair capabilities. Neverfail Heartbeat Plug-ins monitor application services to ensure that protected applications are operational, and not in a 'hung' or 'stopped' state. In addition to monitoring application services, Neverfail Heartbeat can also monitor specific application attributes to ensure that they remain within normal operating ranges. Similar to application monitoring, various rules can be set to trigger specific corrective actions whenever these attributes fall outside of their respective ranges.

- **Data Protection** — Neverfail Heartbeat ensures the data files that applications or users require in the application environment are made available should a failure occur. Once installed, Neverfail Heartbeat can be configured to protect files, folders, and even the registry settings of the active server by mirroring these protected items, in real-time, to one or more passive servers. This means that if a failover occurs, all files that were protected on the failed server will be available to users on the server that assumes the active role after the failover.

Updates to protected files are placed in a queue on the active server (the send queue), ready to be sent to the first passive server with each request numbered to maintain its order in the queue. Once the send queue reaches a specific configured size, or the configured time duration has expired, the update is sent to the first passive server, which places all the requests in an array of log files termed the receive queue. The first passive server then confirms the changes have been logged by sending the active server an acknowledgment.

The first passive server's receive queue is then read in numerical order and a duplicate set of file operations are applied to the disk of the first passive server.

If a Trio configuration is deployed, the first passive server then repeats the process replicating protected data to the second passive server via the WAN connection.

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

Neverfail Heartbeat Networking Configuration

The server IP address used by a client to connect to the active server, the Principal (Public) IP address, must be a static IP address (not DHCP-enabled).

In the example below, the Principal (Public) IP address is configured as 192.168.1.5.

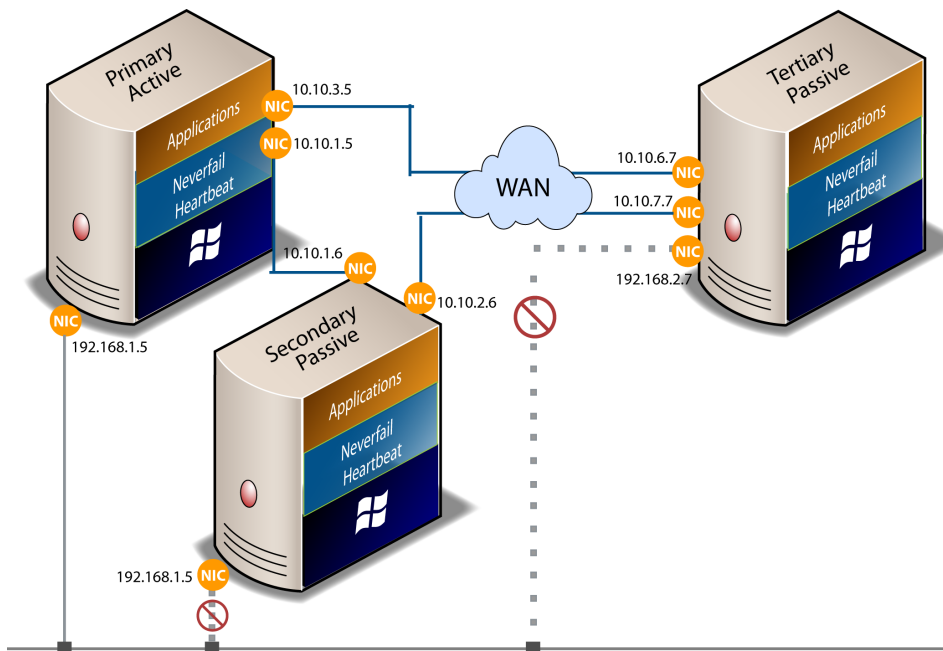


Figure 1: Neverfail Heartbeat Network Configuration

Note: The IP addresses of all NICs on the server can be obtained using a Windows command prompt and typing `ipconfig /all`.

In a High Availability configuration, the Principal (Public) NIC on the first passive server uses the same IP address as the active server but is prevented from communicating with the live network through an IP packet filtering system installed with Neverfail Heartbeat. This packet filter prevents traffic using the Principal (Public) IP 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.

When configured for Disaster Recovery (DR) to a remote site with a different subnet, Neverfail Heartbeat must be configured to use a different Principal (Public) IP address for the Primary and Secondary servers. When a switchover is performed, the DNS server will be updated to redirect users to the new active server at the DR site. These updates are not required when the same subnet is used in the Disaster Recovery Site. Neverfail Heartbeat uses DNS Update task to update Microsoft Windows 2003, 2008, 2008 R2, and 2012 DNS servers with the new Principal (Public) IP address. DNS Update runs the `NFDNSUpdate.exe` to perform the following actions:

- First, *NFDNSUpdate* must unregister the current address with all DNS servers that have an entry for the server (this may not be all DNS servers in the enterprise). Unregistering the address involves removing the 'A host record' from the Forward lookup zone and removing the 'PTR record' from any relevant reverse lookup zones.
- Next, *NFDNSUpdate* must register the new address with all DNS servers that need an entry (again this may not be all DNS servers in the enterprise). Registering the address involves adding the 'A host record' to the Forward lookup zone and adding the 'PTR record' to the pertinent reverse lookup zone.
- Finally, where secondary DNS servers are present, *NFDNSUpdate* must instruct them to force a replication with the already updated Primary servers.

The NICs on the Primary, Secondary, and Tertiary servers intended for use by the Neverfail Channel must be configured so that they use IP addresses outside of the Principal (Public) Network subnet range. These addresses are termed the Neverfail Channel addresses.

Important: NetBIOS must be disabled for the Neverfail Channel(s) on the active and passive servers because the Primary, Secondary, and Tertiary servers use the same NetBIOS name. When Neverfail Heartbeat installation is complete (runtime), NetBIOS will automatically be disabled across the channel(s) preventing NetBIOS name conflicts.

The NICs that allow the connectivity across the Neverfail Channel can be standard 100BaseT or Gigabit Ethernet cards providing a throughput of 100Mbps per second or more across standard Cat-5 cabling.

Note: A dedicated channel requires no hubs or routers, but any direct connection requires crossover cabling.

When configured for a WAN deployment, the Neverfail Channel is configured using static routes over switches and routers to maintain continuous communications independent from traffic on the Principal (Public) Network.

Neverfail Heartbeat Communications

The Neverfail Channel is a crucial component of the setup and is configured to provide dedicated communications between the servers. When deploying in a pair configuration, each server in the pair requires at least one network card (see Single NIC configuration in the Installation Guide) although two network cards are recommended (one NIC for the Principal (Public) Network connection and at least one NIC for the Neverfail Channel connection). An additional pair of NICs may be used for the Neverfail Channel to provide a degree of redundancy. In this case, the Neverfail Channel is said to be *Dualed* if more than one dedicated NIC is provided for the Neverfail Channel on each server.

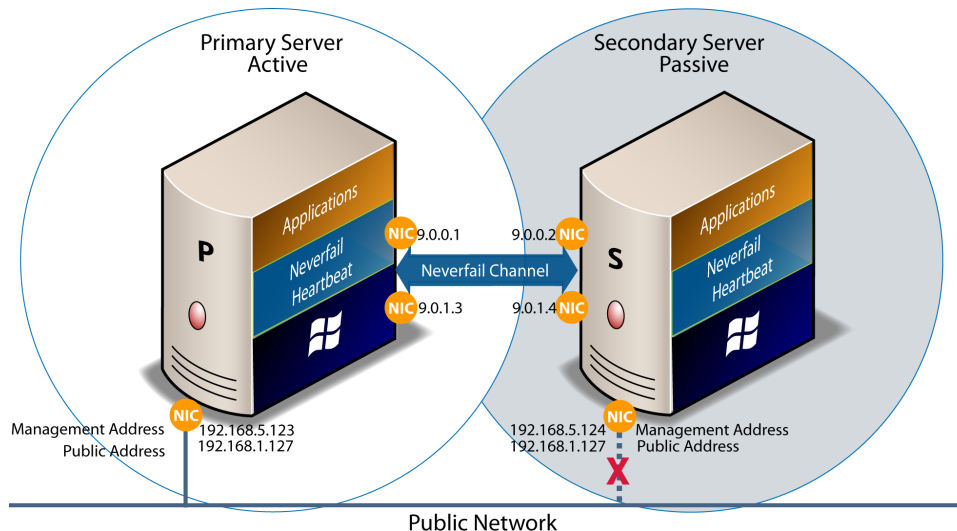


Figure 2: Neverfail Heartbeat Pair Communications

Note: To provide added resilience, the communications for the second channel must be completely independent from the first channel, for example, they must not share any switches, routers, or WAN connection.

When deploying in a Trio configuration, each server in the Trio requires at least one network card (see Single NIC configuration in the Installation Guide) although three NICs are recommended: one

for the Principal (Public) Network connection and one separate NIC for each Neverfail Channel. In a Tertiary configuration, each server communicates with the two adjacent servers as illustrated below.

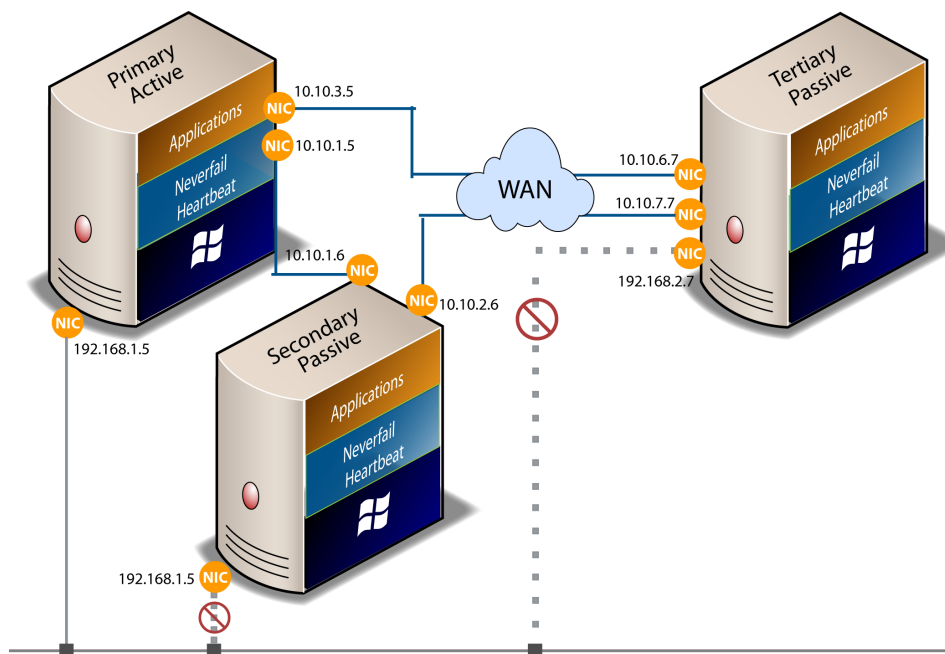


Figure 3: Neverfail Heartbeat Trio Communications

Neverfail Heartbeat Switchover and Failover Processes

Neverfail Heartbeat uses four different procedures to change the role of active and passive servers depending on the status of the active server.

***Note:** This section illustrates the simpler cases of switchover and failover in a Neverfail Heartbeat Pair. These processes are similar in a Neverfail Heartbeat Trio, except that the passive servers in a Trio configuration are referred to as “*First Passive*” and “*Second Passive*.” Under most failover and switchover conditions, the First Passive server assumes the active role.*

The Managed Switchover Process

A managed switchover can be initiated manually from the Neverfail Heartbeat Management Client **Server Summary** page by selecting the server to make active and clicking the **Make Active** button. When a managed switchover is initiated, the running of protected applications is transferred from the active machine to a passive machine in the Cluster - the server roles are reversed.

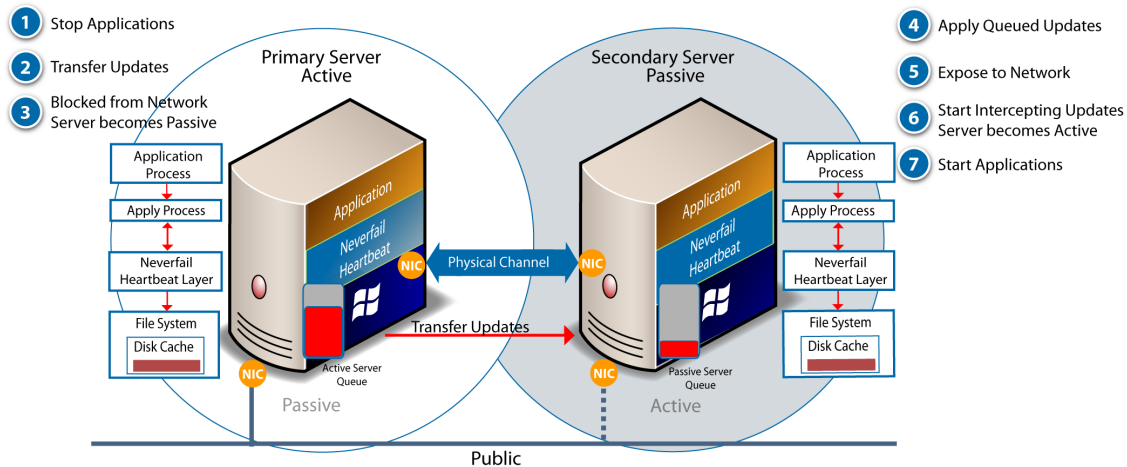


Figure 4: Neverfail Heartbeat Switchover Process

The automatic procedure executed during a managed switchover operation includes the following steps:

1. Stop the protected applications on the active server. Once the protected applications are stopped, no more disk updates are generated.
2. Send all updates that remain queued on the active server to the passive server. After this step, all updates are available on the passive server.
3. Change the status of the active server to '*switching to passive*'. The server is no longer visible from the network.
4. Apply all queued updates on the passive server.
5. Change the status of the passive server to '*active*'. After this step, the new active server starts intercepting disk I/Os and queues them for the new passive server. The new active server becomes visible on the network with the same identity as the old active server.
6. Change the status of the old active server from '*switching to passive*' to '*passive*'. The new passive server begins accepting updates from the active server.
7. Start the same protected applications on the new active server. The protected applications become accessible to users.

The managed switchover is complete.

The Automatic Switchover Process

An automatic-switchover (auto-switchover) is triggered automatically if a protected application, which the system is monitoring, fails.

An auto-switchover is different from a managed switchover in that although the server roles are changed, Neverfail Heartbeat is stopped on the previously active server to allow the administrator to verify the integrity of the data on the newly passive server and to investigate the cause of the auto-switchover.

Auto-switchovers are similar to failover (discussed next) but initiated upon the failure of a monitored application. Once the cause for the auto-switchover is determined and corrected, the administrator can use the **Configure Server Wizard** to change the server roles to their original state.

The Automatic Failover Process

When a passive server detects that the active server is no longer running properly, it assumes the role of the active server.

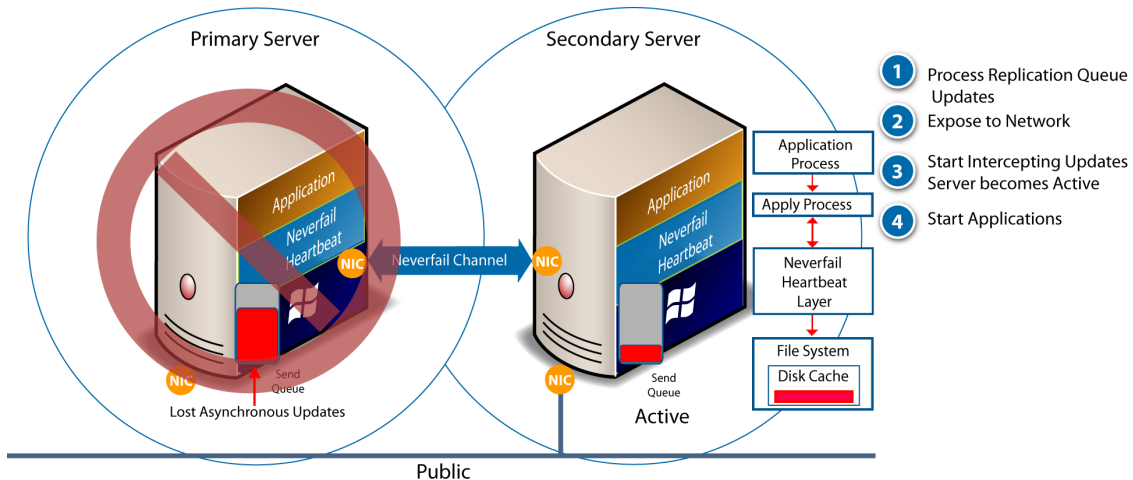


Figure 5: Neverfail Heartbeat Failover Process

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

1. It applies any intercepted updates that are currently saved in the passive server receive queue as defined by the log of update records that are saved on the passive but not yet applied to the replicated files.

The length of the passive server receive queue affects the time the failover process takes to complete. If the passive server 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. When there are no more update records that can be applied, it discards any update records that it is unable to apply (an update record can only be applied if all earlier update records are applied, and the completion status for the update is in the passive server receive queue).

2. It switches its mode of operation from passive to active.

It 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 only be enabled on one of the systems at any time. When the public identity is enabled, any clients previously connected to the server before the automatic failover are able to reconnect.

3. It starts intercepting updates to the protected data. Updates to the protected data are saved in the send queue on the local server.
4. It starts all the 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 stage, the originally active server is “off the air,” and the originally passive server assumes the role of the active server and runs the protected applications. Because the originally active server stopped abruptly, the protected applications may lose some data, but the updates that completed before the failover are retained. The application clients can reconnect to the application and continue running as before.

The Managed Failover Process

A managed failover is similar to an 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 occurs until the system administrator chooses to trigger this operation manually.

Recovery from a Failover

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

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

When Neverfail 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 at boot time, and halts Neverfail Heartbeat so that the issues that caused the failure can be resolved.

The following steps are used to restore the previously failed server to the active role.

1. Correct incident that caused the failover.
2. Verify the integrity of the disk data on the failed server.
3. Restart the failed server.
4. Once all such issues are resolved, run the **Configure Server Wizard** on the failed server to set the correct server role.
5. Start Neverfail Heartbeat on failed server.

At this stage, the instances of Neverfail Heartbeat running on the servers connect and start to resynchronize the data on the Primary server.

6. Allow Neverfail Heartbeat to fully synchronize.

When synchronization is complete, you can continue running with this configuration (for example, the Secondary is the active server and the Primary is the passive server), or initiate a managed switchover to reverse the server roles in the Neverfail Heartbeat Pair (for example, giving the Primary and Secondary the same roles that they had before the failover).

7. Perform a managed switchover.

Chapter 2

Configuring Neverfail Heartbeat

Configure Server Wizard

The Neverfail Heartbeat - **Server Configuration Wizard (Configure Server Wizard)** helps you set up and maintain communications between Neverfail Heartbeat servers. Configuration information includes the IP address for the Neverfail Channel(s) and Principal (Public) addresses on all servers in the Pair or Trio. The Identity of a server (Primary, Secondary, or Tertiary) describes the physical hardware of the machine and should not be confused with what the server is doing (the Role).

Prerequisites

Prior to making changes using the Neverfail Heartbeat **Configure Server Wizard**, you must stop Neverfail Heartbeat.

Procedure

- Once Neverfail Heartbeat is stopped, navigate to **Start > All Programs > Neverfail Heartbeat > Configure Server Wizard** to launch the *Configure Server Wizard*.

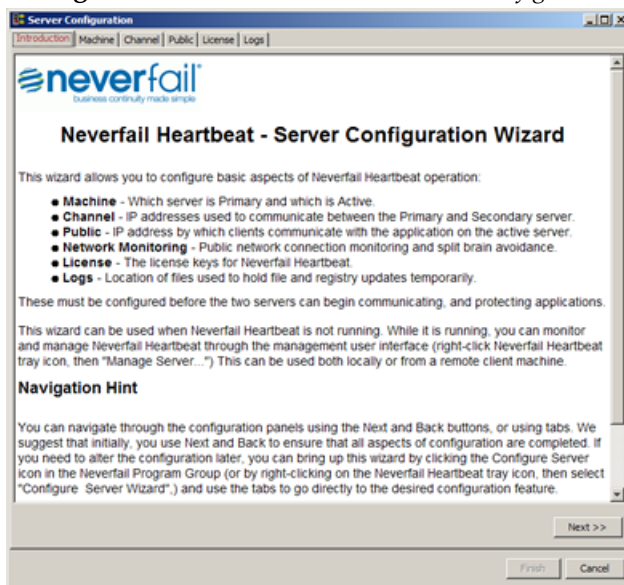


Figure 6: Configure Server Wizard - Introduction Tab

Configure Machine Identity

The Identity of a server (Primary, Secondary, or Tertiary) describes the physical hardware of the machine and should not be confused with what the server is doing (the Role).

Prerequisites

Prior to making changes using the **Configure Server Wizard**, you must stop Neverfail Heartbeat.

Procedure

- To change the machine *Identity*, select the **Machine** tab of the **Configure Server Wizard** and select the *Physical Hardware Identity* of the local machine and click **Next** or **Finish**.

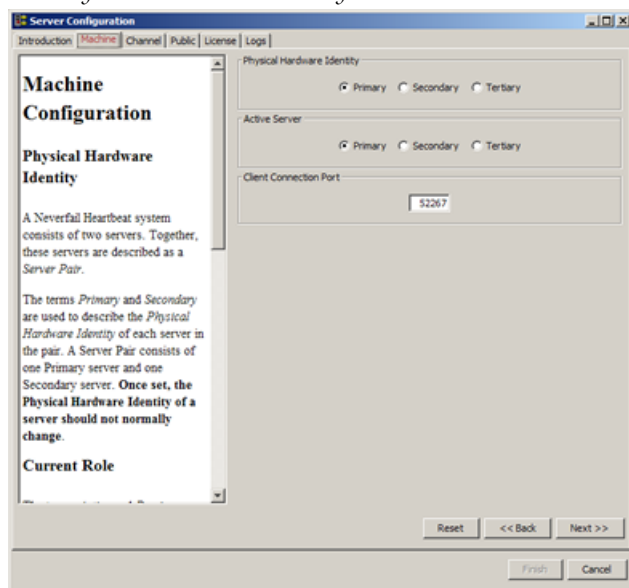


Figure 7: Configure Server Wizard - Machine Tab

Configure Server Role

The server's role describes what the server is currently doing.

Prerequisites

Before changing the *Role* of the local server to active, verify that no other server (including remote servers) in the Cluster is active.

Procedure

- To change the *Role* of the server, select the **Machine** tab of the **Configure Server Wizard** and specify which server in the Cluster is active. Click **Next** or **Finish**.

Change the Client Connection Port

The *Client Connection Port* specifies the port through which clients (such as the Neverfail Heartbeat Management Client) connect to Neverfail Heartbeat.

Procedure

- To change the *Client Connection Port*, select the **Machine** tab of the **Configure Server Wizard** and type a new value in the text box. Click **Next** or **Finish** to accept changes.

Note: Do not change this port unless the default port (52267) is required by another application.

Configure Channel IP Routing

Channel IP routing defines the IP addresses used to communicate between the local server (such as the Primary) and the adjacent servers (such as the Secondary and Tertiary). Each link uses two addresses, one for the local server and one for the remote server.

Procedure

- To add a channel after installing and configuring the NICs, select the **Channel** tab of the **Configure Server Wizard**. Add the new IP addresses for the local server and the remote server to the *Neverfail Channel IP Routing* table by clicking the **Add Row** icon. The drop-down list shows the IP addresses available on the local server. Manual entry of the IP addresses for remote servers is required

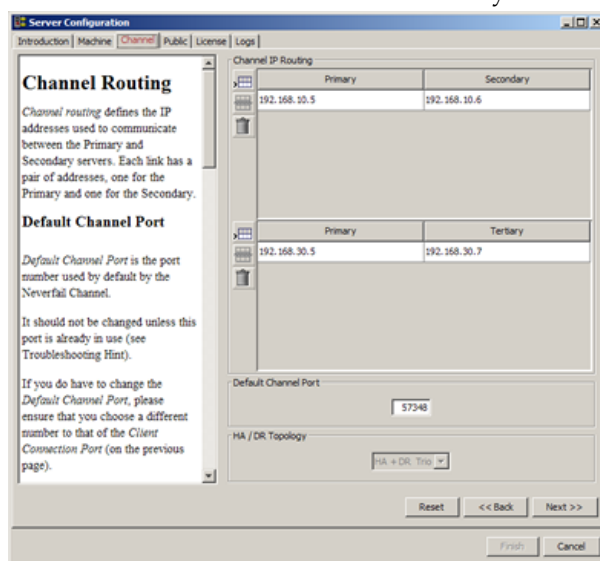


Figure 8: Configure Server Wizard - Channel IP Routing

- To change the channel IP addresses, select and edit the entry in the table. Click **Next** or **Finish** to accept changes.

Configure the Default Channel Port

The Neverfail Channel uses the *Default Channel Port* to communicate between the Primary and Secondary, Secondary and Tertiary, and Tertiary and Primary servers. Do not change this port unless required by another application.

Procedure

- To change the *Default Channel Port*, select the **Channel** tab of the **Configure Server Wizard** and edit the default entry (57348). Click **Next** or **Finish** to accept changes.

Configure Low Bandwidth Optimization

Low Bandwidth Optimization is configured automatically during installation based upon the configuration options selected during Setup. Low Bandwidth Optimization can be configured for: High Availability (HA) when deployed as a pair in a LAN, DR when deployed in a WAN, or HA and DR when deployed as a trio.

In a High Availability (HA) server pair, the queues and buffers are optimized for a high-speed local area network (LAN) connection, compression is disabled, and automatic failover between servers is enabled. In a Disaster Recovery (DR) pair, the queues and buffers are optimized for a low-bandwidth wide area network (WAN) connection, compression may be used, and automatic failover between servers is disabled. In a trio configuration, the only standard configuration is HA between Primary and Secondary, and DR between Primary and Tertiary and between Secondary and Tertiary. In a server pair you can choose HA or DR topology, but in a server trio you normally only have the option of HA+DR topology. However, if you have manually configured a non-standard topology, for example, by changing the Auto-Failover setting, then "Non-Standard" will appear in the menu and you can choose to leave the non-standard topology option as it is, or reset it to one of the standard topologies.

Note: *The same HA/DR configuration must be set on all servers in the pair or trio.*

- To change Low Bandwidth Optimization after installation, select the **Channel** tab of the **Configure Server Wizard** and use the *HA/DR Topology* drop-down to select the appropriate topology. Click **Next** or **Finish** to accept changes.

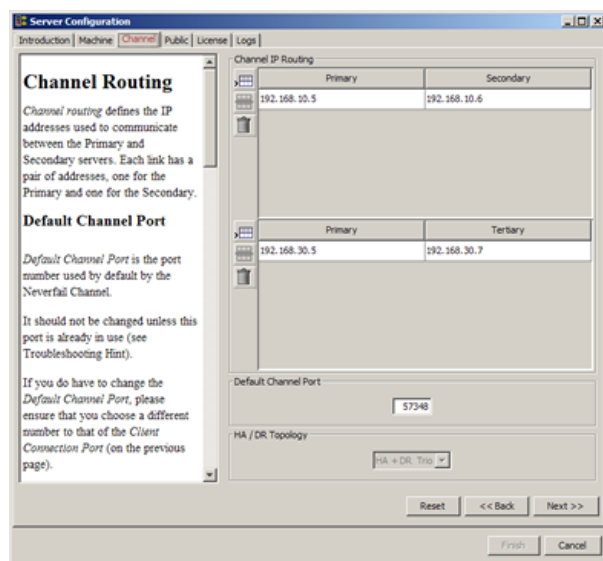


Figure 9: Configure Server Wizard - Channel tab

Configure Public IP Addressing

A typically configured Neverfail Heartbeat Cluster uses only one Principal (Public) IP address when deployed as a pair or on a LAN, but can be configured with more than one Principal (Public) IP address. These are the addresses by which clients of the protected application connect to the application. Typical installations configure the same Principal (Public) IP address on the Primary and Secondary servers. The Tertiary server may have the same Principal (Public) IP address if on a LAN, but may require a different address when deployed over a WAN. All traffic to and from these Principal (Public) IP addresses is passed through to the active server but blocked on the passive server(s). When the server

roles are switched, the IP filtering mode also switches, so client systems always connect to the Principal (Public) IP addresses on whichever server is currently active. When the Neverfail Heartbeat service is shut down, the filtering remains in place to prevent IP address conflicts between servers.

Procedure

1. To configure Principal (Public) IP addressing, select the **Public** tab of the **Configure Server Wizard** and list all of the addresses intended for use as Principal (Public) IP addresses.

Note: An address must not appear more than once, and no Principal (Public) IP address may appear in the list of IP addresses on the Channel tab.

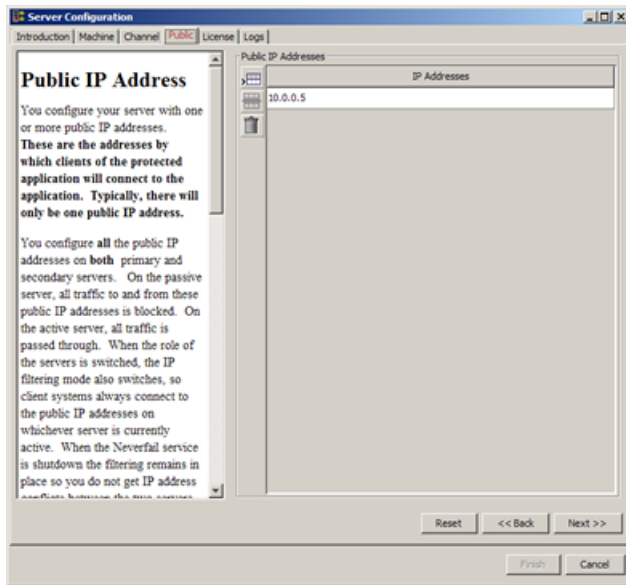


Figure 10: Configure Server Wizard - Public Tab

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

Add/Remove an Neverfail Heartbeat License Key

When adding additional licensed features to Neverfail Heartbeat after installation, use the *Configure Server Wizard* to add additional License Keys.

Procedure

1. To manage Neverfail Heartbeat License Keys, select the **License** tab of the **Configure Server Wizard**.
2. To add an entry to the *License Keys* table, manually type or paste (using **Ctrl+V**) your license key into the table. Alternatively, click **Import** on the tool bar to import a license file (.txt). License keys are available from Neverfail or your distributor.

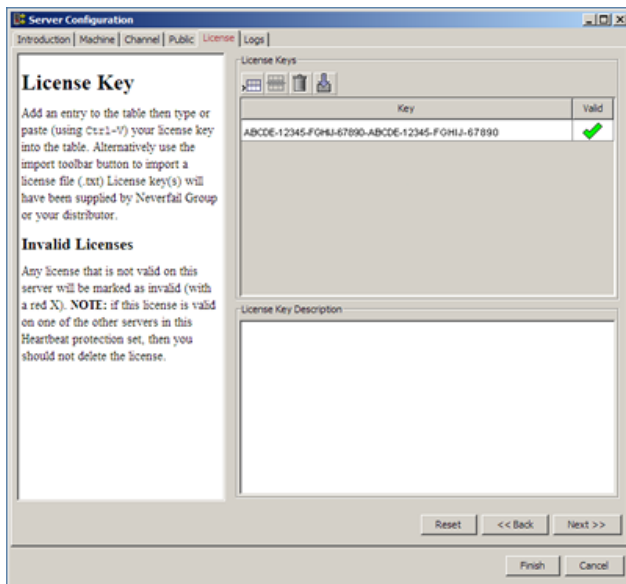


Figure 11: Configure Server Wizard - License Tab

- After entering your license keys click **Next** or **Finish**.

Configure the Message Queue Logs

The configured message queue logs location determines where the local server temporarily stores replication data received (the receive queue) and the replication data waiting to send (the send queue). This configuration affects only the local server; logs can be in different locations on the Primary, Secondary, and Tertiary servers.

Procedure

- To configure the location of the message queue logs, select the **Logs** tab of the **Configure Server Wizard**. Click **Browse** to open an Explorer-like window. Navigate to and select the folder for storing the message queue logs, and click **Finish** to accept the location.

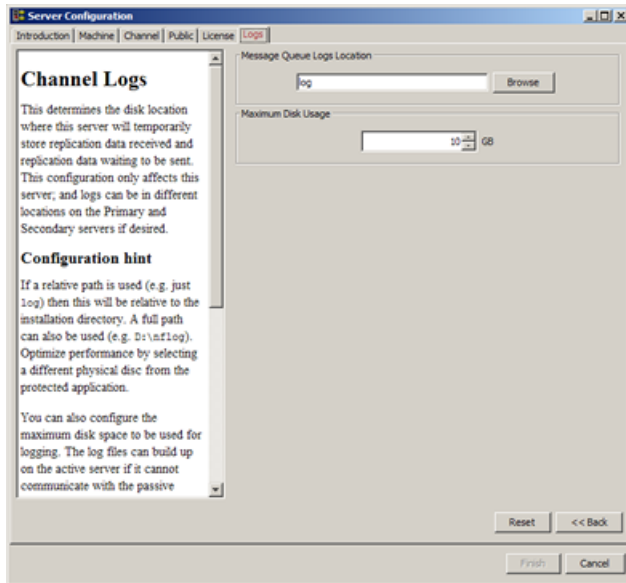


Figure 12: Configure Server Wizard - Logs Tab

Configure Maximum Disk Usage

You can configure the maximum disk space allocated for logging. Log files accumulate when the active server cannot communicate with the passive server, when a passive server is performing certain operations, or when a server is under heavy load. Configuring this value is important because when the value set for maximum disk usage is reached, replication stops, and your system is no longer protected. If your system uses a dedicated disk for log files, consider disabling the maximum disk usage setting.

Procedure

- If your system uses a dedicated disk for log files, consider disabling the maximum disk usage setting. To do this, set *Maximum Disk Usage* to zero (0).

Note: When *Maximum Disk Usage* is disabled, there is a risk that Neverfail Heartbeat may run out of physical disk space, and when this happens, a shutdown and restart may be required before replication can resume.

- Neverfail recommends a *Maximum Disk Usage* setting that leaves a little overflow space to enable Neverfail Heartbeat to stop replicating gracefully. To configure *Maximum Disk Usage*, select the **Logs** tab of the **Configure Server Wizard** and enter the maximum dedicated disk space allocated for message queue log files and click **Finish** to accept the changes.

Chapter 3

Status and Control

Managing Neverfail Heartbeat Clusters

Neverfail Heartbeat operates in Clusters of two or three servers with each Cluster administered as a single entity using the Neverfail Heartbeat Management Client. The Neverfail Heartbeat Management Client, which can be run from any server in the Cluster or remotely from another machine in the same subnet, simplifies routine administration tasks for one or more Clusters.

Note: *The controlling workstation must have Neverfail Heartbeat or Neverfail Heartbeat Client Tools installed.*

Install Client Tools

Neverfail Heartbeat allows installation of Neverfail Heartbeat Client Tools for remote management of Neverfail Heartbeat clusters.

Prerequisites

When installing Neverfail Heartbeat Client Tools on Windows XP, the following Service Pack levels are required.

- Windows XP 32 bit SP3
- Windows XP 64 bit SP2

Note: *Neverfail Heartbeat Client Tools requires that Microsoft™ .Net Framework 4 be installed prior to running Setup.exe. If .Net Framework 4 is not installed when you attempt to initiate Setup, Neverfail Heartbeat Client Tools will prevent installation until .Net Framework 4 is installed.*

Procedure

1. To install Neverfail Heartbeat Client Tools, download, extract the contents, and initiate Neverfail Heartbeat setup.exe.
 - a) Download Neverfail Heartbeat to a suitable location. .zip
 - b) Extract the contents of the Neverfail Heartbeat .zip file.
 - c) Double-click the setup.exe file to initiate the installation process.
2. On the **Setup Type** page, select *Install Client Tools Only* and click **Next**.
3. Read the license agreement carefully and select *I accept terms of the License Agreement*. Click **Next**.
4. Configure the installation paths. The default installation location is C:\Program Files\Neverfail\ but can be changed by manually typing a path to another install location.

Alternatively, click **Browse** to select one of these locations. Select *Create icons on Desktop* and click **Next**.

Note: The path of the Neverfail Heartbeat Client Tools installation folder cannot contain Unicode characters. The path of the Neverfail Heartbeat Client Tools installation folder can only contain lower and upper case letters A to Z, digits from 0 to 9, and the following special characters: space \ _ - () . :

Additionally, Neverfail Heartbeat Client Tools does not support file or folder names ending with a period "." or space " ".

5. Review the summary of options and configuration information for the installation. Click **Next**. Pre-install checks run to ensure that the installation can continue. The *Report* pane displays the results of the pre-install checks. If some pre-install checks are unsuccessful, go back through the wizard, make the necessary changes, and run the pre-install checks again.
6. If the pre-install checks are successful, click **Next**. The next page displays the progress of the installation. During this process, Neverfail Heartbeat Setup installs the necessary files and folders onto your system and applies the configuration you specified.
7. Click **Next** after Neverfail Heartbeat Client Tools components are complete. The **Client Tools Installation Complete** page is displayed.
8. Click **Finish**.

Starting the Neverfail Heartbeat Management Client

You can start the Neverfail Heartbeat Management Client from any server in the Neverfail Heartbeat Cluster.

Procedure

1. Right-click the Neverfail Heartbeat interactive status icon in the Windows tool tray (located on the right side of the Windows tool bar).



The Neverfail Heartbeat quick access menu opens.




Figure 13: Neverfail Heartbeat Quick Access Menu

2. Select **Manage Server**. The Neverfail Heartbeat Management Client opens in a new window and shows the Neverfail Heartbeat Management Client Servers (overview) page.

Alternatively you can start the Neverfail Heartbeat Management Client from the Neverfail Heartbeat program group on the Windows Start menu. This is the only method supported if the Neverfail Heartbeat Management Client has been installed on a workstation that is not part of the Cluster.

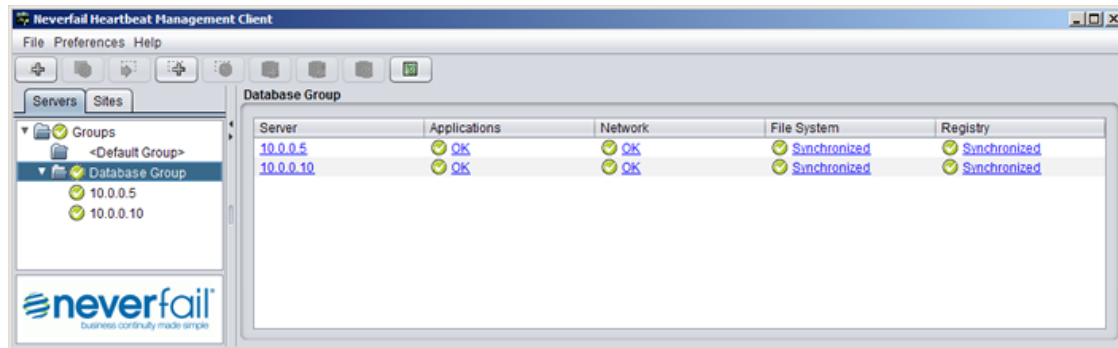


Figure 14: Neverfail Heartbeat Servers (Overview) page

Navigating the Neverfail Heartbeat Management Client

After the Neverfail Heartbeat Management Client is running, use the navigation panel on the left of the Neverfail Heartbeat Management Client page to view and select the Groups, and Cluster Connections to manage with the Neverfail Heartbeat Management Client.

Note: A Group is an arbitrary collection of Neverfail Heartbeat Clusters used for organization.

A Connection, or Cluster Connection, allows Neverfail Heartbeat Management Client to communicate with a Neverfail Heartbeat Cluster, either on the same machine or remotely.

The Group or Cluster you select make in the navigation panel “point” the Neverfail Heartbeat Management Client to that Group or Cluster, and the Neverfail Heartbeat Management Client provides information related to only the selected Group or Cluster. To avoid confusion, pay particular attention to the selection in the navigation panel when you are managing more than one Group or Cluster.

Note: Groups and Clusters are not automatically detected by the Neverfail Heartbeat Management Client. Each Group and Cluster you want to manage must be added to the Neverfail Heartbeat Management Client before you can use it to view status or change settings for that Group or Cluster Connection.

Select a Cluster in the navigation panel of the Neverfail Heartbeat Management Client to show a set of tabs and sub-tabs that offer detailed status and control of the associated Neverfail Heartbeat servers in the Cluster. The table below provides an outline of these pages.

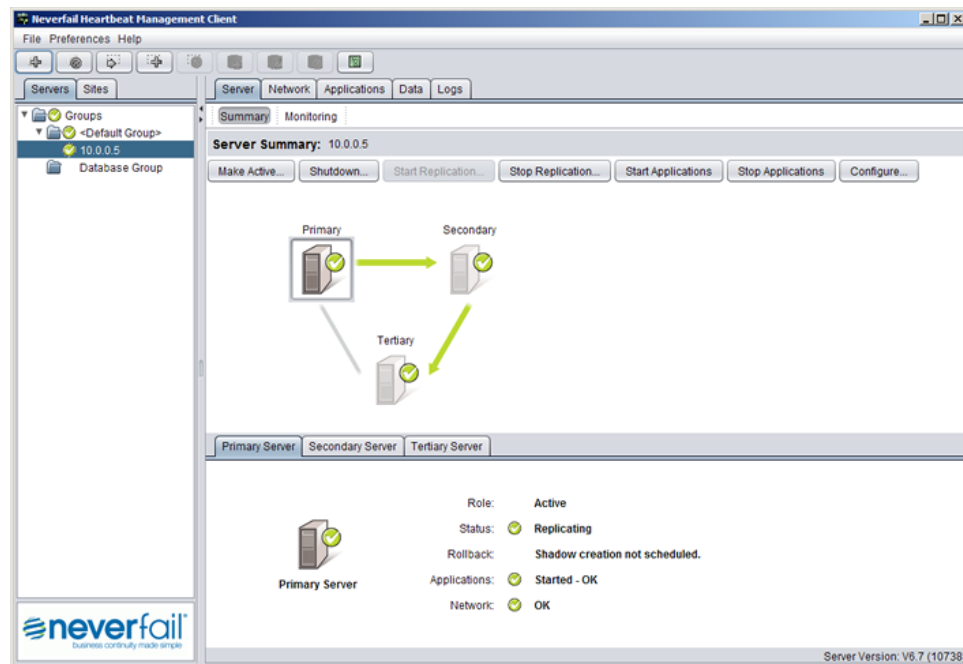


Figure 15:

Table 2: Neverfail Heartbeat Management Client Interface Features

Tab	Sub-tab	Interactive Features
Server	Summary	<p>Main frame features selectable servers (click on a server to select) and graphical status</p> <p>Tabs (Primary, Secondary, and Tertiary) in bottom frame provide status of servers in the Cluster (role, replication status, rollback status, file synchronization, registry synchronization, and recover point delay)</p> <p>Buttons (actions apply to selected server) include Make Active, Shutdown, Start Replication, Stop Replication, Start Applications, Stop Applications, and Configure.</p>
	Monitoring	<p>Graphical display of connectivity status between servers in the Cluster.</p> <p>Buttons include Configure Pings, Configure Failover, and Configure Response Times</p>
Network	----	<p>Graphical display of network connectivity status</p> <p>Summary pane identifying the Active server and status of network connectivity</p> <p>Ping Targets from Currently Active Server pane listing targets and associated IP addresses</p> <p>Buttons include Configure Pings and Configure Auto-switchover</p>
Applications	Summary	<p>Detail of application health and Log</p> <p>Application Health pane - Clear button</p> <p>Applications pane - Edit, Remove, Start Applications, Stop Applications, and Configure buttons</p> <p>Applications Log pane - Filters, Properties buttons</p>

Tab	Sub-tab	Interactive Features
	Services	List of services and current status Add, Edit, Remove, Move Up and Move Down buttons
	Tasks	List of tasks Add, Edit, Remove, Move Up, Move Down, Run Now , and User Accounts buttons
	Rules	List of rules Edit, Check Now buttons
	Plug-ins	List of Plug-ins Install, Uninstall, Edit buttons
Data	Replication	Sub-tabs - <i>Active to First Passive, Passive to Second Passive</i> Configuration button - Displays the <i>Orphaned Files, Fast Check</i> , and <i>Controlled Shutdown</i> tabs Registry Synchronization Status pane - Full Registry Check button File Hierarchy pane - Full System Check button, Synchronize button + <i>Including Subdirectories</i> check box, Verify and Synchronize button + <i>Including Subdirectories</i> check box.
	File Filters	List of file filters and state (<i>Effective</i> or <i>Not Effective</i>) Buttons - Add Inclusion Filter, Add Exclusion Filter, Edit, Remove
	Rollback	Cancel button for Primary and Secondary Shadow Creation List of existing shadows - Create, Delete, Rollback, Configure buttons
	Traffic Queues	Graphical status of traffic - Configure button Tabs in bottom frame provide detailed information about <i>Primary Server</i> and <i>Secondary Server</i> send and receive queues status (<i>P>S Connection</i>)
	WAN Compression	Status of WAN and WAN Compression/Acceleration, Configure button
Logs	----	Detailed Event Log Buttons: Export event log as comma-separated text, Mail event log to recipients immediately, Remove all entries from the event log, Filters, Configure, Properties, Mail Settings, Configure Alerts, Test Alert Reporting

Changing the Neverfail Heartbeat Management Client Font

The font face and font size can be changed using the Neverfail Heartbeat Management Client interface.

Procedure

1. Select **Font Selection** from the *Preferences* menu. The *Font Selection dialog* opens.

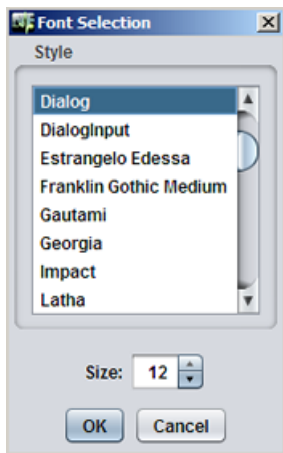


Figure 16: Font Selection Dialog

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** to accept the changes and close the dialog.
A confirmation message appears.
5. Click **Yes** to confirm the changes and restart the Neverfail Heartbeat Management Client to apply the new font settings. Click **No** to restart later; the changes will be applied the next time the Neverfail Heartbeat Management Client is started. Click **Cancel** to exit without making any changes

Add a Neverfail Heartbeat Group

The Add Group feature in the Neverfail Heartbeat Management Client allows you to add new Neverfail Heartbeat Groups to manage.

Procedure

1. Open the Neverfail Heartbeat Management Client 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.




Figure 17: Add Group - Tool Bar Button

2. Type a name for the new group into the text box and click **OK** to create the group. The newly created group appears in the navigation panel on the left of the Neverfail Heartbeat Management Client window.



Figure 18: Add Group Dialog




Figure 19: Group Added to Navigation

Remove an Neverfail Heartbeat Group

The Remove Group feature in the Neverfail Heartbeat Management Client allows you to remove existing Neverfail Heartbeat Groups from management.

Procedure

1. Select the Group to be removed in the navigation panel of the Neverfail Heartbeat Management Client. Click **Remove Group** in the tool bar, or select *Remove Group* from the *File* menu.




Figure 20: Remove Group - Tool bar Button

A Confirmation message appears.

2. Click **Yes** to remove the selected Group or **No** to cancel removal and close the message window. The selected Group is removed.

Add a New Connection

When you created a Neverfail Heartbeat Group using the instructions in *Add a Neverfail Heartbeat Group*, you created an empty container. Next you must add the connections to the Cluster or Clusters that comprise your new Neverfail Heartbeat Group. The *Add Connection* feature in the Neverfail Heartbeat Management Client allows you to add a new Cluster Connection to an existing Neverfail Heartbeat Group.

Procedure

1. In the navigation panel, select the Neverfail 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** from the menu.




Figure 21: Add Connection - Tool bar Button

The **Add Connection** dialog appears.

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).




Figure 22: Add Connection Dialog

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




Figure 23: Add Connection Login

3. Click the **Alternate IPs** button to add additional IPs for the server.

Note: The *Add Connection* dialog allows you to add additional Management IP addresses to the connection for the server. In the event that the main IP address becomes unavailable, the Neverfail Heartbeat Management Client will attempt to use the alternate IP addresses to connect to Neverfail Heartbeat.

Important: To use alternate IP addresses, you must configure Management IP addresses for the server.




Figure 24: Add Alternate IPs Dialog




Figure 25: Add Alternate IP Dialog

4. Click **OK** on each dialog to create the connection.
The newly created connection appears in the navigation panel on the left of the Neverfail Heartbeat Management Client window, and the **Server Summary** page updates to represent any existing network relationships of the added server.
5. Enter the remaining connections necessary to define the new Neverfail Heartbeat Group.

Edit a Connection

The Edit Connection feature in the Neverfail Heartbeat Management Client allows you to change the *Port Number* for existing connections.

Procedure

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.

Note: When a configured connection is not found, an error message similar to the one shown in below may be displayed. Click **Edit Connection** to reconfigure the connection.

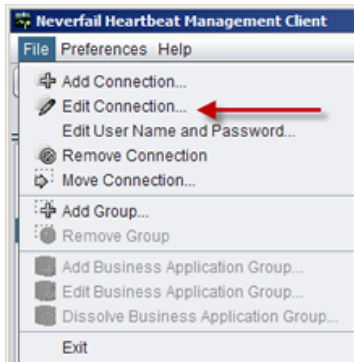


Figure 26: Edit Connection - File Menu

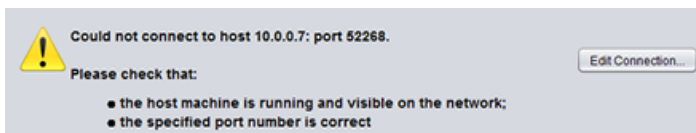


Figure 27: Edit Connection - Error Message



Figure 28: Edit Connection Dialog

2. Type the new value for the *Port Number* into the text box, or use the up and 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 the Neverfail Heartbeat Management Client allows you to reassign an existing Cluster to a different Group.

Procedure

1. Select the Cluster 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.




Figure 29: Move Connection - Tool bar Button

The Move Connection dialog appears.




Figure 30: Move Connection Dialog

2. Select the destination Group to receive the Connection from the drop-down list and click **OK** to move the Connection, or click **Cancel** to dismiss the dialog without making any changes.

Remove a Connection

The Remove Connection feature in the Neverfail Heartbeat Management Client allows you to remove an existing Connection.

Procedure

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.




Figure 31: Remove Connection - Tool bar Button

A confirmation message appears.

2. Click **Yes** to remove the selected Connection or **No** to cancel removal and close the message window.

Business Application Groups

Neverfail Heartbeat offers the ability to group application servers together creating a Business Application Group. Business Application Groups are a grouping of servers that share a common purpose such as Microsoft Exchange servers, BlackBerry Enterprise servers, or Microsoft SQL servers

for monitoring and management purposes. With Neverfail for Business Application Plug-in installed, Neverfail provides the ability to manage groups of servers as a single entity and perform switchovers of a complete group from one site to another.

Installing the Business Application Plug-in

The Business Application Plug-in (`BusinessApplicationNFPlugin.dll`) is installed after installing Neverfail Heartbeat.

Prerequisites

Prior to installing and configuring the Business Application Plug-in, complete the following:

- If you are not using the same host name for all servers in a Cluster, you must configure Alternate IP addresses on all servers in the Secondary sites.
- Configure persistent static routes for the Neverfail Channel between the servers within a Business Application Group site as explained below:
 - Configure persistent static routes between all of the Primary servers within the Business Application Group at the Primary HA site.
 - Configure persistent static routes between all of the Secondary servers within the Business Application Group at the Secondary HA site.
 - Configure persistent static routes between all of the servers within the Business Application Group at the DR site.

Note: *Add persistent routes with a lower metric to allow them to be attempted first.*

1. Download the `Z-SW-BusinessApplicationPlugin.201.5.[n].zip` file from the Neverfail Extranet *Products / Downloads* page www.neverfailgroup.com to a temporary location on the active server in the cluster.

Note: *The `BusinessApplicationNFPlugin.dll` must be downloaded and installed on each cluster server to be included in the Business Application Group.*

2. Extract the archive `.zip` file.
3. Launch the Neverfail Heartbeat Management Client and navigate to the **Applications: Plug-ins** page.

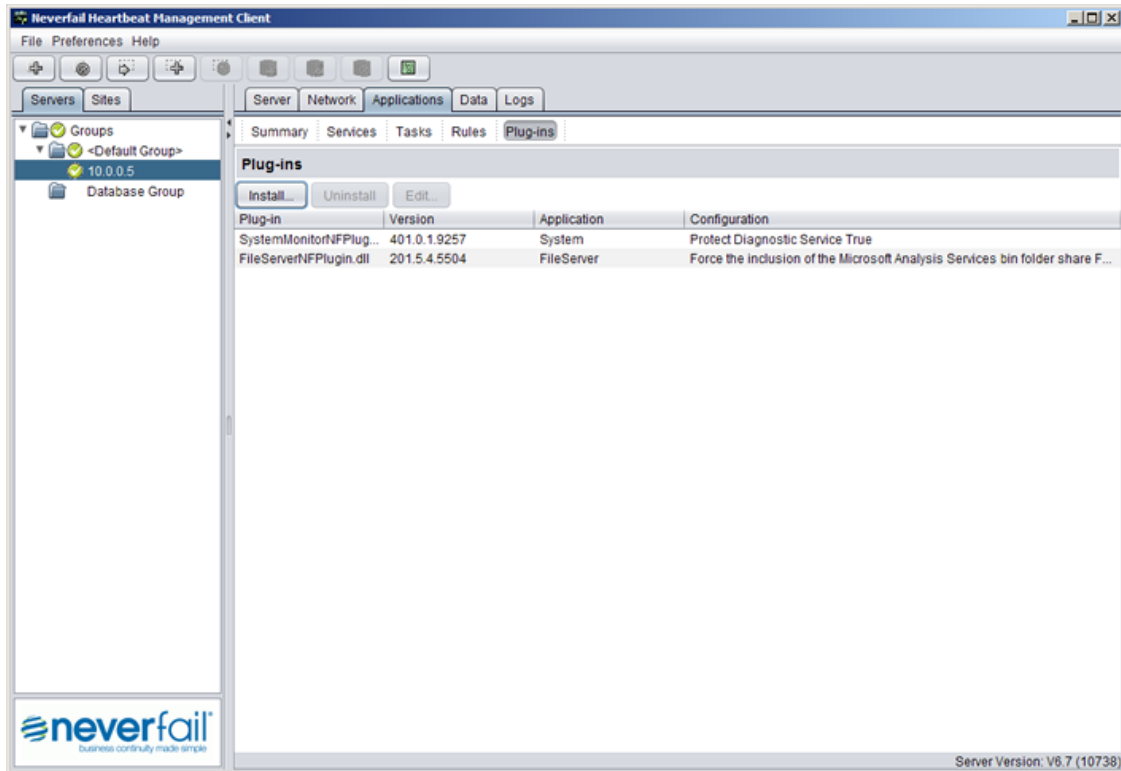


Figure 32: Applications: Plug-ins Page

4. Click **Install**.
5. Browse to the location of the BusinessApplicationNFPlugin.zip file and select the file.




Figure 33: Install Plug-in Dialog

6. Click **OK**.
7. Repeat the process on each Cluster to be included in the Business Application Group.

Important: Once the Business Application Plug-in has been installed, Neverfail recommends that you do **NOT** edit the Business Application Plug-in directly but rather use the **Edit Business Application Group Wizard** to make changes to the plug-in parameters.

Creating a Business Application Group

When the Neverfail Heartbeat Business Application Plug-in is installed it is initially in an unconfigured state. The Unconfigured icon appears in the left pane of the Neverfail Heartbeat Management Client under Servers. All servers listed in the Unconfigured category are available as Business Application

Group candidates and may be added to a Business Application Group. Add the appropriate servers to a Business Application Group to monitor or manage servers with a common function or purpose as a group.

Prerequisites

The Neverfail Heartbeat Management Client requires that you have access to a minimum of two Neverfail Heartbeat clusters displayed in the *Servers* pane as Unconfigured to create a new Business Application Group.

1. Launch the Neverfail Heartbeat Management Client.
2. Navigate to **File > Add Business Application Group**.
The **Business Application Group Wizard** is displayed.

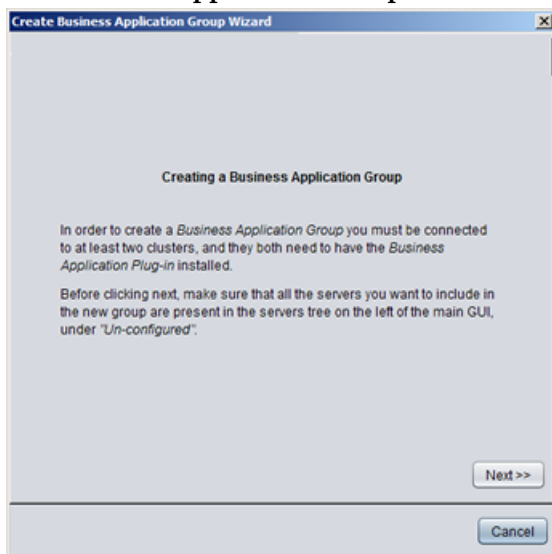


Figure 34: Business Application Group Wizard

3. Review the information in the **Create Business Application Group Wizard** and click **Next**. The **Enter Basic Group Information** page is displayed.



Figure 35: Enter Basic Group Information Page

4. Enter a name for the Business Application Group into the text field.

Note: The name of the Business Application Group cannot exceed 15 characters.

5. Add the name of the Primary Site.
6. Add the name of the Secondary (DR) site and click **Next**.
The **Add Servers to Business Application Group** page is displayed. A list of available servers is displayed in the left pane of the dialog.



Figure 36: Add Servers to Business Application Group Page

7. Select the servers to join the Business Application Group and click the > button to add the servers to the Business Application Group. Click **Next**.
The **Select First Server to Switch** page is displayed.
8. Select the server you want to be the first to switch within the Business Application Group. Click **Next**.

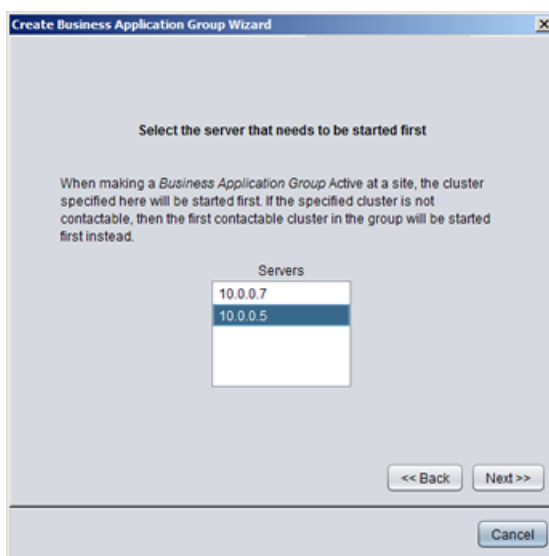


Figure 37: Select First Server to Switch Page

Note: *Neverfail Heartbeat will attempt to switch the server indicated in step 8 above but in the event that the server is unavailable, Neverfail Heartbeat will continue to switch other servers in the Business Application Group.*

The **Create Business Application Wizard Complete** page is displayed.

9. The **Create Business Application Wizard Complete** page informs you that you have successfully created a Business Application Group and can now take advantage of Neverfail Heartbeat's Site Switchover capabilities discussed in [Site Switchover](#). Click **Finish**.




Figure 38: Create Business Application Wizard Complete Page

Editing a Business Application Group

Neverfail Heartbeat Management Client allows you to edit the configuration of an existing Business Application Group.

1. Navigate to **File > Edit Business Application Group** or click on the **Edit Business Application Group** button.
The **Edit Business Application Group Wizard** is displayed.

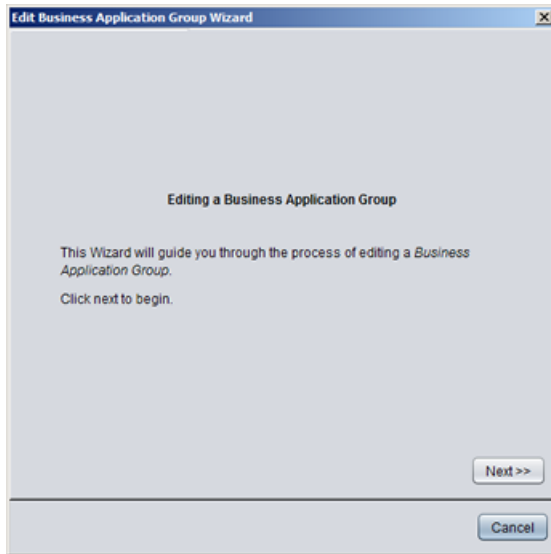


Figure 39: Edit Business Application Group Wizard

2. Click **Next**.
The **Enter Basic Group Information** page is displayed.

Figure 40: Enter Basic Group Information

3. Edit the name of the Business Application Group, Primary Site, and/or the Secondary Site and click **Next**.
The **Select First Server to Switch** page is displayed.

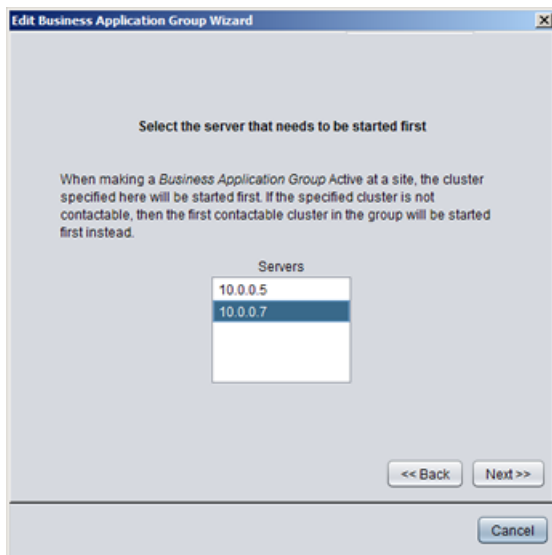


Figure 41: Select First Server to Switch Page

4. Select the server you want to be the first to switch within the Business Application Group and click **Next**.
The *Edit Business Application Wizard* page is displayed.
5. Click **Finish**.

Dissolve a Business Application Group

The Dissolve Business Application Group feature of the Neverfail Heartbeat Management Client allows you to remove a Business Application Group without removing the servers from the Neverfail Heartbeat Management Client.

1. Using the Neverfail Heartbeat Management Client, select the Business Application Group to be dissolved.

***Note:** If you do not intend to recreate the Business Application Group, you must remove the Business Application Plug-in from each server in the Group.*

2. Navigate to **File > Dissolve Business Application Group**.




Figure 42: Dissolve Business Application Group - Tool bar Button

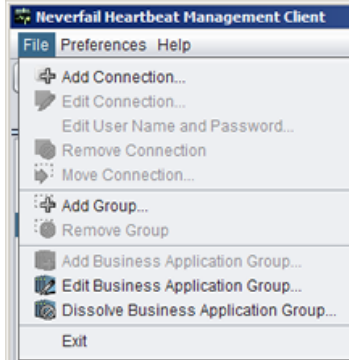


Figure 43: Dissolve Business Application Group - File Menu

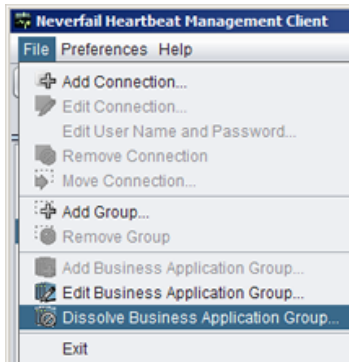


Figure 44: Dissolve Business Application Group - Navigation Panel

A dialog is displayed asking if you are sure you want to dissolve the Business Application Group.



Figure 45: Dissolve Business Application Group - Confirmation Dialog

3. Click **Yes** to dissolve the Business Application Group.

Business Application Switchover

Neverfail Heartbeat provides the ability to perform a managed switchover of a Business Application Group thereby allowing the administrator to transfer the load of the active servers in the Business Application Group to a secondary site with a single operation.

In the event that one of the servers in the Business Application Group should fail, the administrator can perform a managed switchover to the secondary site thereby maintaining continuous availability for users. Additionally, for maintenance and management purposes, the administrator can perform a managed switchover to the secondary site for all servers in the Business Application Group with the click of a single button.

The **Business Application Group Summary** page provides an overview of all servers within the Business Application Group. Selecting an individual server within the group displays information that is specific to the selected server.

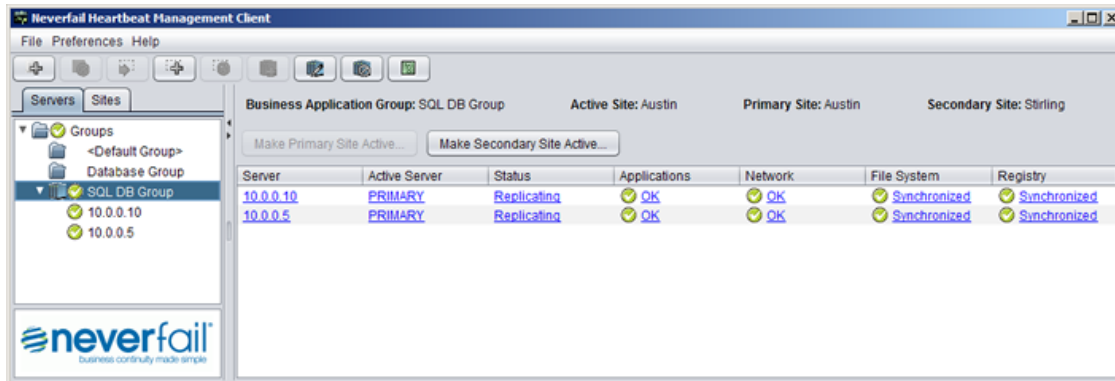


Figure 46: Business Application Group Summary Page

Performing a Business Application Switchover

1. Launch the Neverfail Heartbeat Management Client.
2. Select the *Servers* tab in the left pane.
3. In the *Servers* pane, select the Business Application Group to switch. The **Business Application Group Summary** page is presented.
4. To perform a managed switchover, click:

Option	Description
Make Secondary Site Active	Switches the active operational load from the current (Primary) site to an alternate Secondary site
Make Primary Site Active	Switches the active operational load from the current (Secondary) site to the Primary site

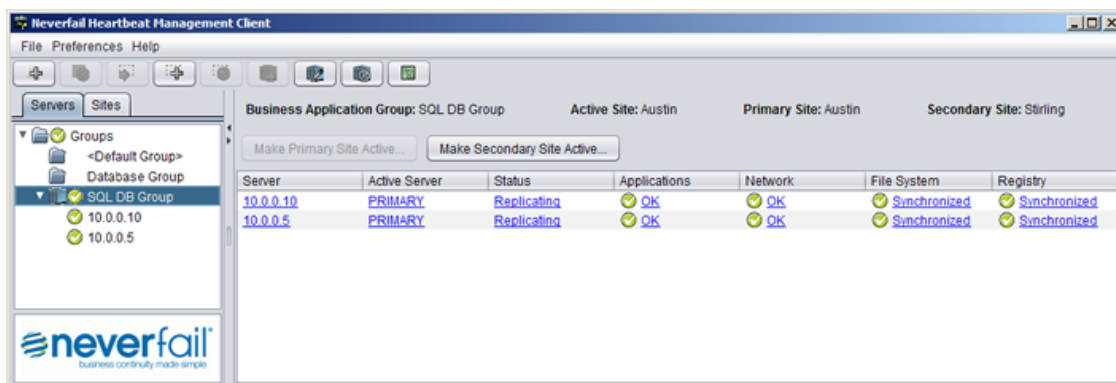


Figure 47: Business Application Group Summary Page

The active servers at the current site become passive and the passive servers at the opposing site become active.

Note: If there are multiple Tertiary Clusters in a Business Application Group and a switchover is initiated on the High Availability connection of the Cluster, a switchover will also be triggered on other Tertiary Clusters High Availability connections in the same Business Application Group.

Site Switchover

When Neverfail Heartbeat is deployed for Disaster Recovery in a pair or High Availability and Disaster Recovery in a Trio, Neverfail Heartbeat can be configured to perform a managed switchover at the site level.

When the Business Application Plug-in is installed and Business Application Groups are configured, Neverfail Heartbeat can provide a single button action to switch the active load of all Business Application Groups in a single site to a Standby Site and back again as required.

This feature can be used when a Business Application Group member server has failed, an application running on one of the servers has failed and cannot be restored, or a total site outage has occurred.

If the server that fails is the server configured to switch first, the Neverfail Heartbeat Management Client will be unable to connect to the host name and after a retry, will attempt to connect via the Alternate IP address. If the Alternate IP address has not been configured, then the connection will drop out of the group and commands to switchover cannot be sent.

In the event of a WAN outage, the administrator needs to ensure that if the standby site is made active, then the administrator must shut down the previously active site to prevent both sites from being simultaneously active. To prevent both sites from being active at the same time, the administrator should shut down the active site prior to making the Standby Site active. A site switchover assumes that the Primary Site has experienced a total failure and that the servers in the Primary Site are not longer running. If this is not the case, the administrator is responsible for shutting down the previously active site.

Performing a Site Switchover

Procedure

1. Launch the Neverfail Heartbeat Management Client.
2. Select the *Sites* tab in the left pane.

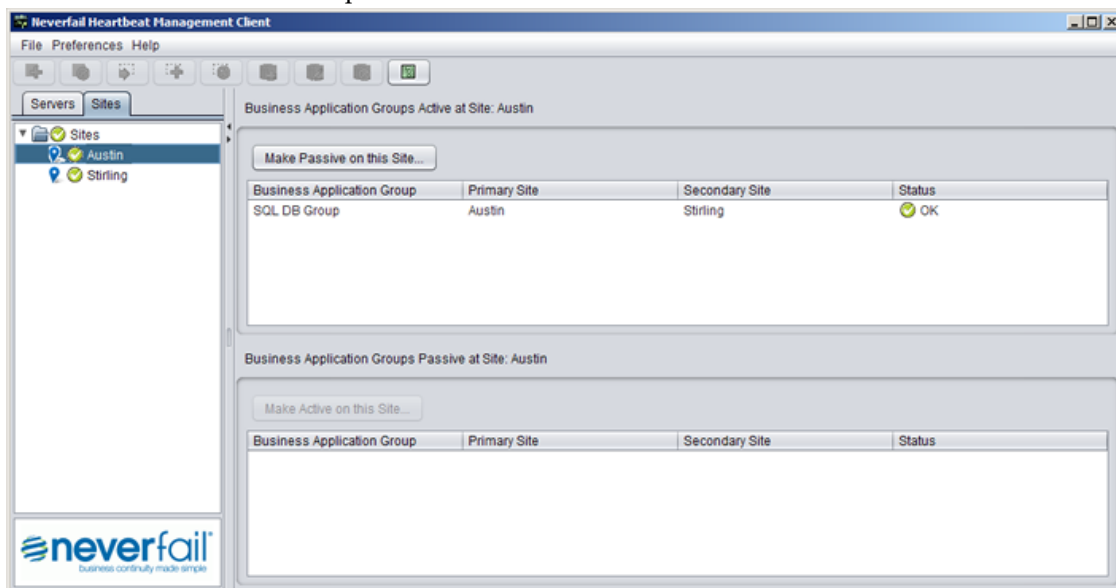


Figure 48: Neverfail Heartbeat Sites Overview Page

3. Select the Site to change.
4. Click:

Option	Description
Make Passive on this Site	The currently active site
Make Active on this Site	The currently passive site

Note: If you select the currently active site, only the **Make Passive on this Site** button is available. If you select the currently passive site, only the **Make Active on this Site** button is available.

Perform a Site Switchover when the First Server to Switch is Unavailable

In the event that the First to Switch server in the Business Applications Group can not be contacted to perform a switchover, you can perform a switchover by performing the steps below:

Procedure

1. Launch the Neverfail Management Client.
2. Login to Neverfail Heartbeat on the Disaster Recovery server of the First to Switch Cluster.
3. Navigate to the **Server: Summary** page.
4. Select the Disaster Recovery server icon.
5. Click the **Make Active** button.
The Disaster Recovery server of the First to Switch Cluster becomes active.

Uninstall the Business Application Plug-in

The Neverfail Heartbeat Management Client allows you to uninstall the Business Application Group Plug-in on-the-fly without stopping Neverfail Heartbeat.

Prerequisites

If the Business Application Plug-in must be uninstalled for any reason, you must first dissolve the Business Application Group and then uninstall the Business Application Plug-in. After uninstalling the Business Application Plug-in, you can then reinstall the plug-in and create a new Business Application Group.

1. After dissolving the Business Application Group, select the server to have the Business Application Plug-in uninstalled.
2. Navigate to the **Applications: Plug-ins** page of the Neverfail Heartbeat Management Client.
3. Select the server on which to uninstall the Business Application Plug-in.
4. Select the BusinessApplicationNFPlugin.dll
5. Click **Uninstall**.
The Business Application Plug-in is uninstalled.

Note: When upgrading the Business Application Plug-in on a server in a Business Application Group, you must upgrade the Business Application Plug-in on all other servers in the Business Application Group.

Edit User Name and Passwords

The Edit User Name and Password feature in the Neverfail Heartbeat Management Client allows you to change the *User Name* and *Password* settings used to connect to a given Cluster.

Procedure

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.

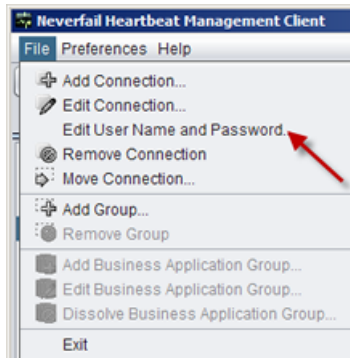


Figure 49: Edit User Name and Password - Navigation panel

The **Edit Username and Password** dialog is displayed.




Figure 50: Edit User Name and Password Dialog

2. Type new values for *User Name* and *Password* into the corresponding text boxes and click **OK** to accept the changes or **Cancel** to close the dialog without making any changes.

Review the Status of Neverfail Heartbeat Clusters and Groups

Procedure

- Click on the top level of the Neverfail Heartbeat Management Client, *Groups*, to view a list of all managed Clusters and a quick status of the protected applications, network, file system, and registry settings for each Cluster. In the example below, two Clusters are identified and both are operating as expected.

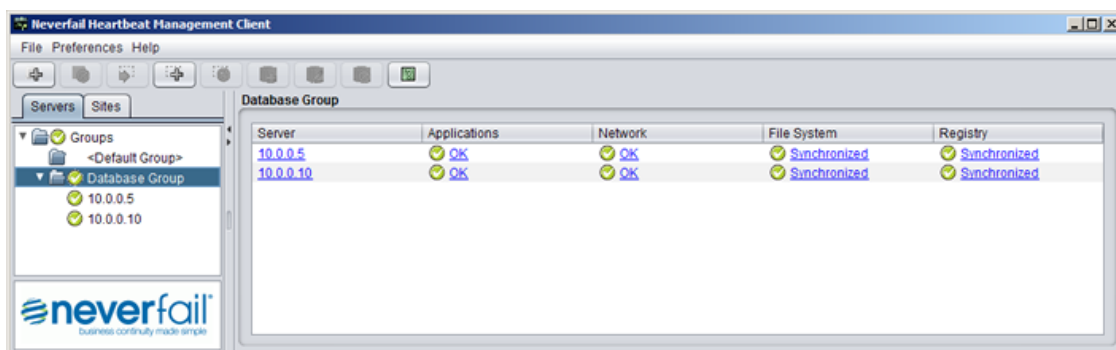


Figure 51: Neverfail Heartbeat Servers Overview page

The status hyperlinks in the overview page link to pages that provide more specific, related information and management controls.

- Click:

Option	Description
Server	To view the Server: Summary page
Applications	To view the applications status on the Applications: Summary page
Network	To view the network status on the Network Monitoring page
File System	To view the File System status on the Data: Replication page
Registry	To view the Registry status on the Data: Replication page

Exit Neverfail Heartbeat Management Client

Procedure

1. Click **Exit** on the **File** menu.
The **Confirm Exit** message appears.
2. Click **Yes** to close the Neverfail Heartbeat Management Client window or **No** to dismiss the message without exiting the Neverfail Heartbeat Management Client.

Shutdown Windows with Neverfail Heartbeat Installed

Procedure

- Always stop Neverfail Heartbeat before attempting to shut down Microsoft Windows. If an attempt is made to shut down Windows without stopping Neverfail Heartbeat, Neverfail Heartbeat will not stop in a graceful manner.

Controlled Shutdown

A Controlled Shutdown is a process where the Neverfail Heartbeat service is able to delay a system shutdown for a sufficient period to perform all of the necessary steps required to stop the applications and replication in a synchronized state. The Controlled Shutdown is intended for situations where an unattended planned shutdown of the server is necessary. When configured in the Neverfail Heartbeat Management Client **Data: Replication** page, this feature allows Neverfail Heartbeat to gracefully shutdown in the absence of the administrator.

Procedure

1. Navigate to the **Data: Replication** page of the Neverfail Heartbeat Management Client.
2. Click the **Configure** button.
3. Select the *Controlled Shutdown* tab of the **Replication Configuration** dialog.
4. Select the servers on which to enable Controlled Shutdown.
5. Select the days and hours parameters under which the server(s) will perform Controlled Shutdown.
6. Configure the length of time for the server(s) to wait for the Controlled Shutdown.

Note: The ability to configure the length of time for the server(s) to wait for the Controlled Shutdown is configurable on Windows Server 2008 and 2012 but is not configurable on Windows Server 2003.

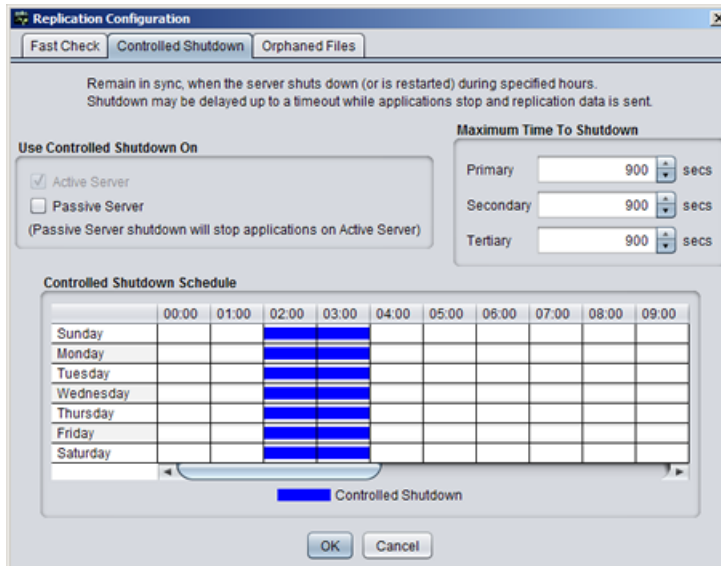


Figure 52: Controlled Shutdown

7. click **OK**.

Note: When the Fast Check process is enabled in addition to the Controlled Shutdown process, Neverfail Heartbeat can be scheduled to perform unattended restarts of the system while maintaining synchronization of data. For more information about Fast Check, see [Configure Fast Check](#).

Neverfail vSphere Client

Neverfail vSphere Client is an optional feature that employs the Neverfail vSphere Client Extension Plug-in and integrates Neverfail vSphere Client with VMware's vSphere Client. Neverfail vSphere Client allows administrators to manage Neverfail Heartbeat Clusters and Groups from within the VMware vSphere Client. Neverfail vSphere Client allows administrators to Discover Neverfail Clusters and Groups, manually Add and Remove Neverfail Heartbeat servers, perform a switchover using the Make Active command, Start Replication, Stop Replication, and Shutdown Neverfail Heartbeat.

Additionally, the Neverfail vSphere Client allows administrators to view the current status of Neverfail Heartbeat servers previously discovered or added to Neverfail vSphere Client. Administrators can view Replication status, Application health, File and Registry Synchronization status, and the role of the server.

Note: If the Neverfail Management on vCenter Server features was not previously included in your Neverfail Heartbeat license, you must regenerate your license and install it using the Configure Server Wizard to add the new feature. The Neverfail Management on vCenter Server feature will only be available for installation once the feature has been added to your license and the license installed.

Installing Neverfail vSphere Client

Installation of Neverfail vSphere Client is performed on the server running vCenter Server using Neverfail Heartbeat Setup similar to installation of Neverfail Heartbeat.

Prerequisites

- Neverfail Heartbeat v6.6 or later installed on all physical Neverfail Heartbeat servers
- An appropriate Neverfail license for the Pair or Trio for Management on vCenter Server
- vCenter Server installed on a single server to manage virtual servers in a the site

Note: If vCenter Server Heartbeat is installed on the vCenter Server, then Neverfail Heartbeat Management on vCenter Server must be installed on both the active and passive servers. If Neverfail Heartbeat Management of vCenter Server is installed on the Primary vCenter Server before installing vCenter Server Heartbeat, the cloning process used during the installation of vCenter Server Heartbeat will automatically install Neverfail Management of vCenter Server on the Secondary server.

1. Download the Neverfail Heartbeat .zip file to a suitable location on the server running vCenter Server.

Note: Before attempting to unzip the Neverfail Heartbeat .zip file, verify that the file is not blocked. Right-click the Neverfail Heartbeat .zip file and select Properties. If the file is shown as blocked, click the **Unblock** button.

2. Extract the contents of the Neverfail Heartbeat .zip file.
3. Double-click the setup.exe file to initiate the installation process.
The **Setup Type** page is displayed.

Note: If you click **Exit** any time after Setup.exe starts, you are prompted to save the current settings. When you run Setup.exe later, you are offered the option to use the previously saved settings.

If .Net 4.0 is not currently installed on the server, Neverfail Heartbeat Setup prompts you to download and install .Net 4.0 before allowing you to install Neverfail Heartbeat Management on vCenter Server. You must exit the installation process, download and install .Net 4.0, and restart the installation process.

The left panel of each page in the Neverfail Heartbeat Setup Wizard provides information about the setup process. Select the configuration options in the panel on the right.

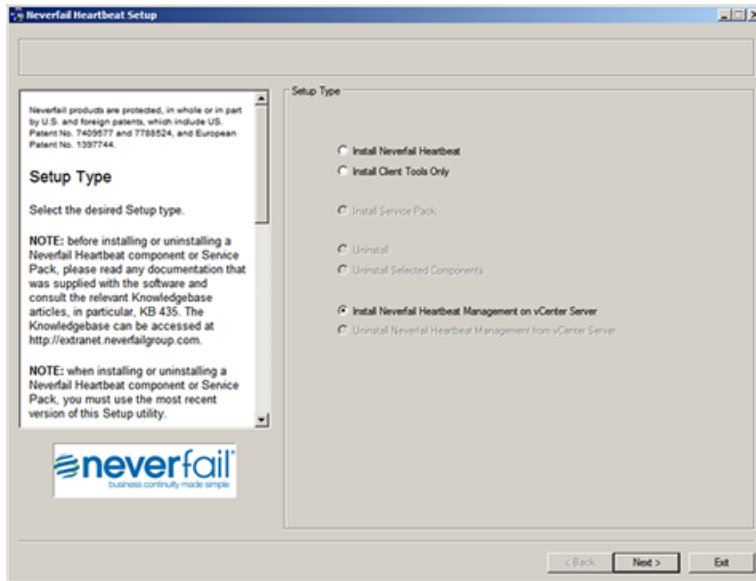


Figure 53: Setup Type page

4. Because the Neverfail vSphere Client Extension Plug-in is being installed, in the **Setup Type** page, select *Install Neverfail Heartbeat Management on vCenter Server*. Click **Next**. The **End User License Agreement** page is displayed.

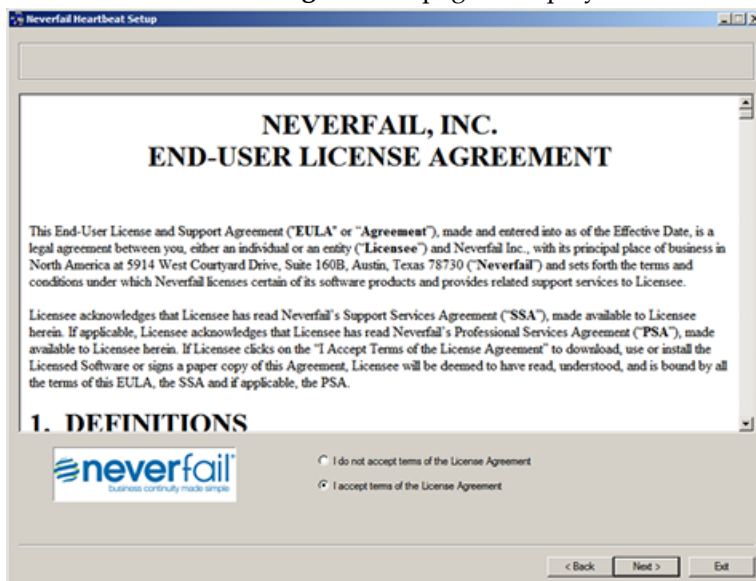


Figure 54: End User License Agreement page

5. Review the contents of the End User License Agreement and select *I accept terms of the License Agreement* to continue installation of Neverfail vSphere Client Plug-in. If you select *I do not accept terms of the License Agreement*, the **Next** button remains inactive and the dialog allows only the **Back** and **Exit** button to be active. Click **Next** to continue. The **Installation Paths** page is displayed.

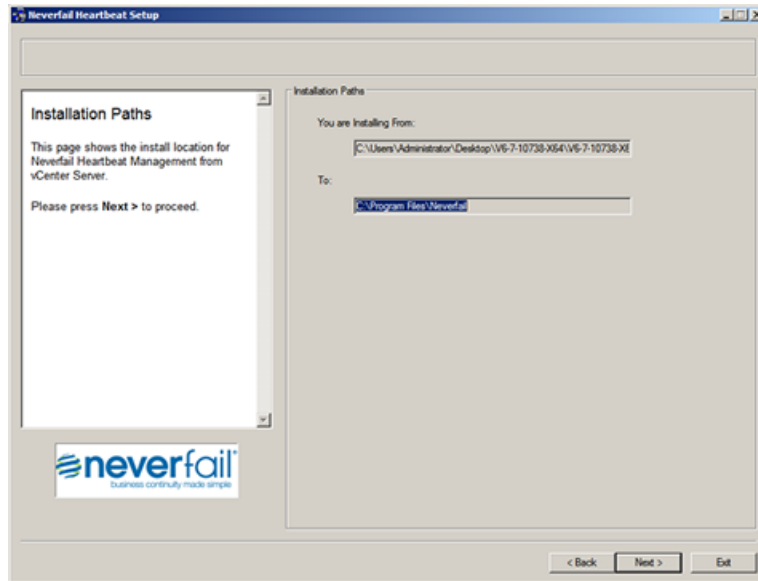


Figure 55: Installation Paths page

6. The **Installation Paths** page displays the location where Neverfail Heartbeat Management on vCenter Server is to be installed. Click **Next**
The **vCenter Management Tool Registration** page is displayed.

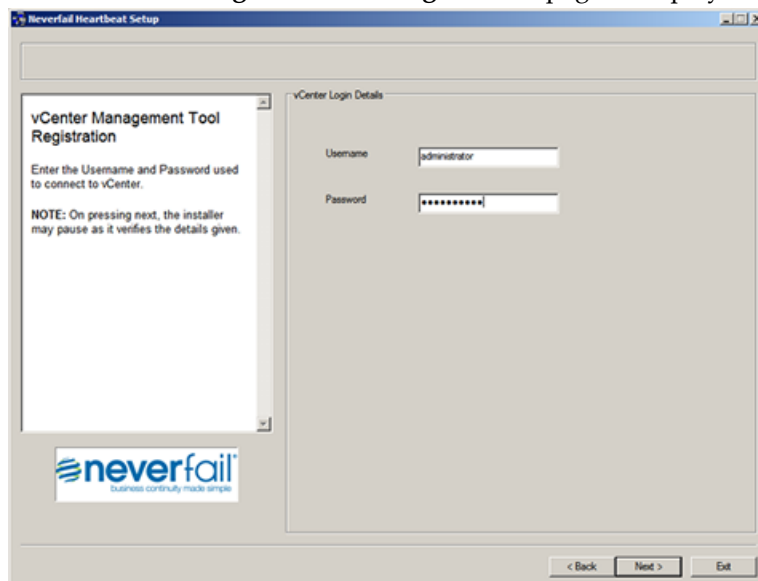


Figure 56: vCenter Management Tool Registration page

7. Enter the *Username* and *Password* used to connect to vCenter Server and click **Next**.
Neverfail Setup verifies the details provided. After the details have been verified, the **Installation Summary** page is displayed.

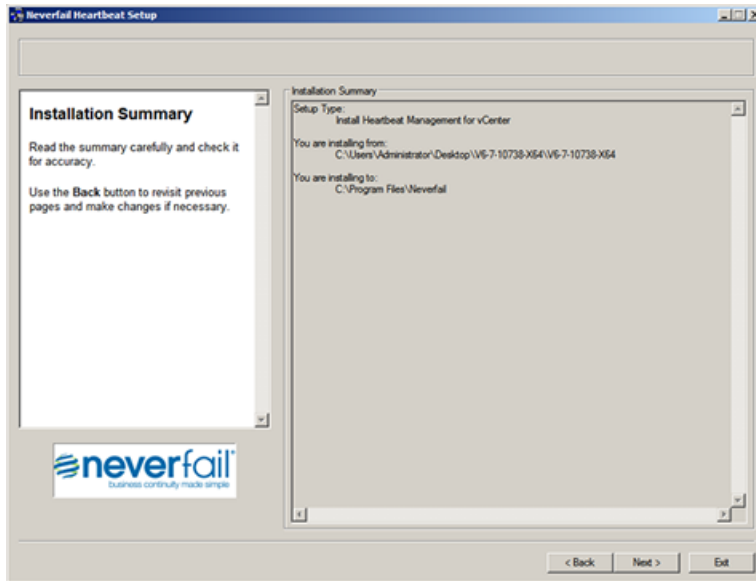


Figure 57: Installation Summary page

8. Review the results of the **Installation Summary** page. Click **Next**.

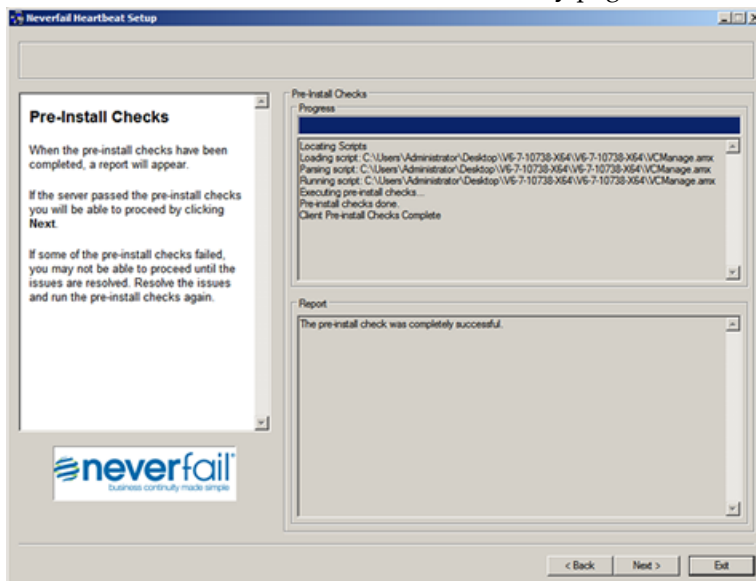


Figure 58: Pre-Install Checks page

9. Neverfail Heartbeat Setup automatically performs pre-install checks to ensure that the installation can continue. If a problem is encountered, click **Back** to return to the appropriate screen to make the changes required to resolve the issue. After making changes, run the pre-install check again to verify that the problem has been resolved and no new issues were created. Once the pre-install check successfully completes, click **Next** to proceed with the installation. The **Install** page is displayed.

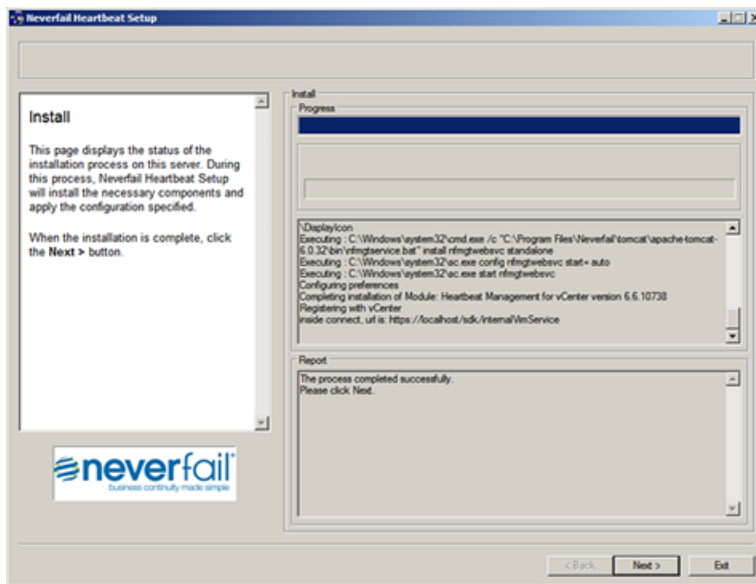


Figure 59: Install page

10. The progress of the installation is displayed in the upper pane of the **Install** page. When the installation has completed, the results of the installation are displayed in the lower pane of the Install page. Once installation has successfully completed, click **Next** to continue. The **Installation Complete** page is displayed.

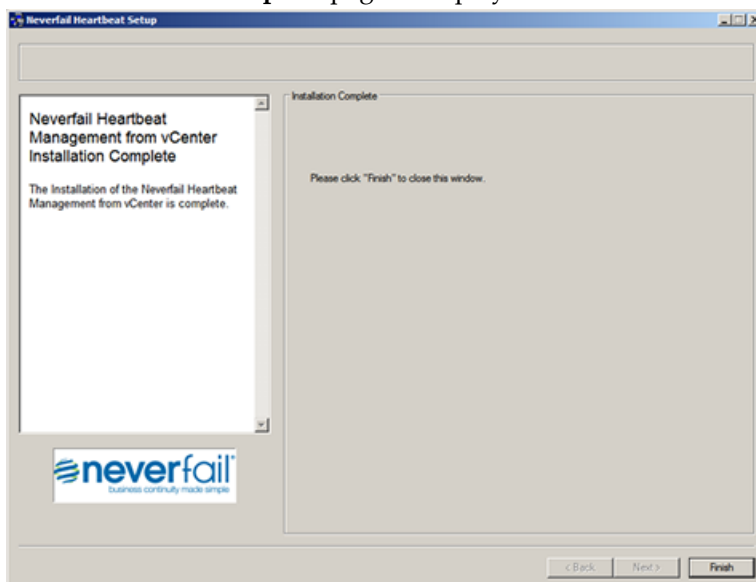


Figure 60: Installation Complete page

Launch Neverfail vSphere Client

Prerequisites

Use of Neverfail vSphere Client requires that Adobe Flash Player 10.0 is installed.

Note: If Adobe Flash Player 10.0 is not installed prior to installation of Neverfail vSphere Client Extension Plug-in, selecting the double-clicking the Neverfail Heartbeat tab in vSphere Client for the first time will provide an opportunity to download Adobe Flash Player 10.0 from the internet and install it.

1. Login to VMware vSphere Client.



Figure 61: vSphere Client Login

2. A security certificate is presented. Select the check box to install the security certificate.

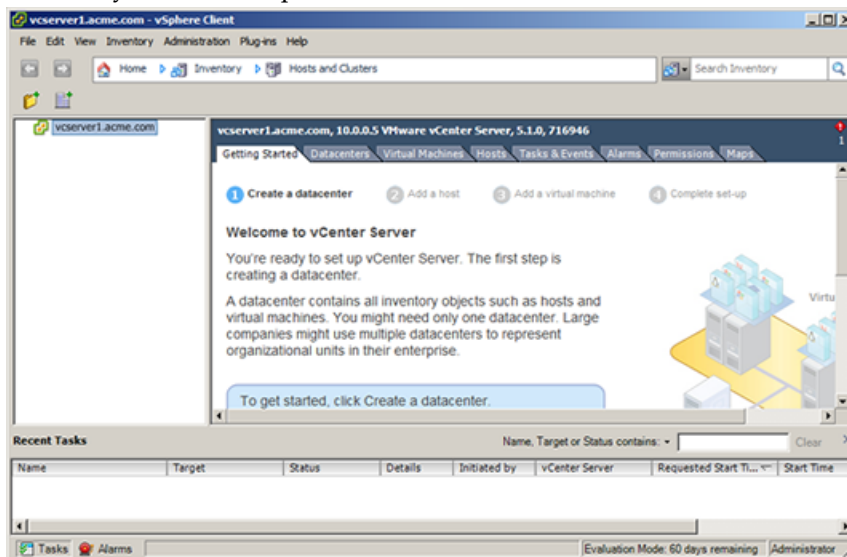


Figure 62: vSphere Client

3. Navigate to **Home > Solutions and Applications**.

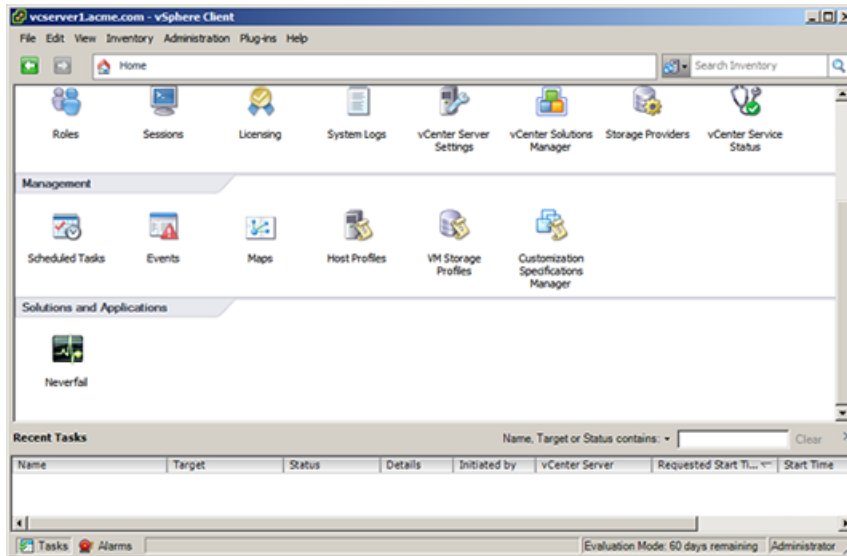


Figure 63: vSphere Client Solutions and Applications

4. Double-click the Neverfail Heartbeat icon in the *Solutions and Applications* pane. Neverfail vSphere Client displays the Neverfail Heartbeat page.

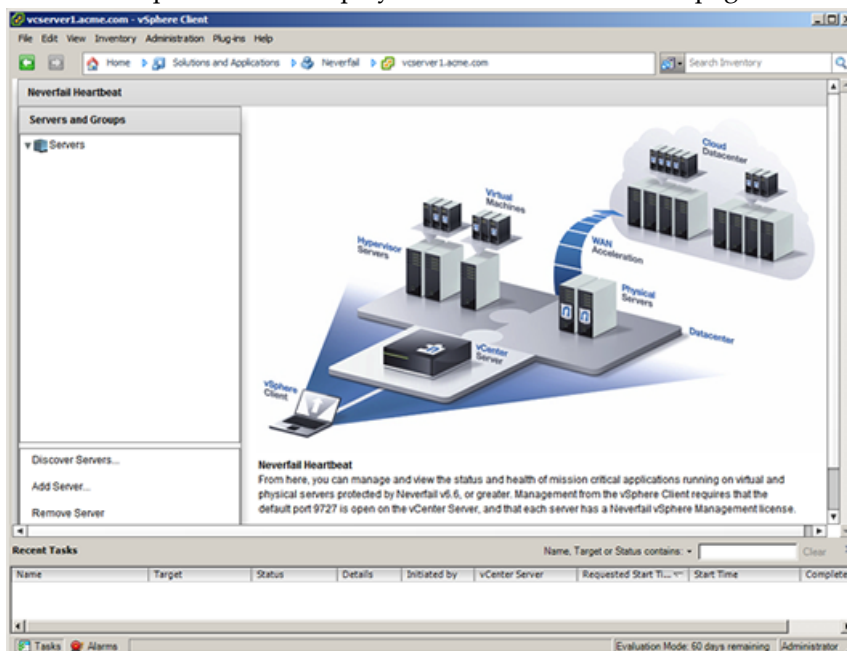


Figure 64: Neverfail Heartbeat page

Discover Neverfail Heartbeat Servers

Neverfail vSphere Client provides the ability to run discovery to identify all Neverfail Heartbeat Clusters and Groups.

Prerequisites

Each server to be managed by Neverfail vSphere Client requires a Neverfail vSphere Management license.

1. From the **Neverfail Heartbeat** page, click **Discover Servers** in the left pane of the page. The **Discover Server** dialog is displayed.

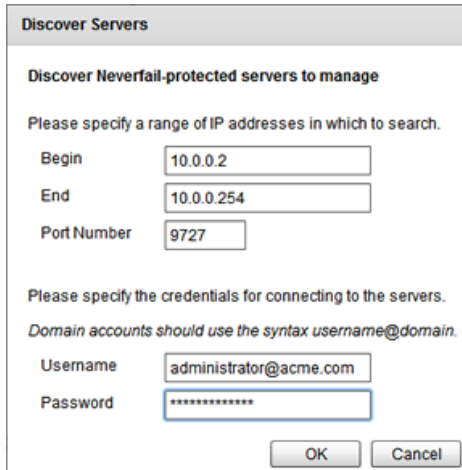
The image shows a 'Discover Servers' dialog box. It has a title bar 'Discover Servers'. Below the title bar is a section 'Discover Neverfail-protected servers to manage'. Under this section, there is a prompt 'Please specify a range of IP addresses in which to search.' followed by three input fields: 'Begin' with the value '10.0.0.2', 'End' with the value '10.0.0.254', and 'Port Number' with the value '9727'. Below these fields is another prompt 'Please specify the credentials for connecting to the servers.' followed by a note 'Domain accounts should use the syntax username@domain.' and two input fields: 'Username' with the value 'administrator@acme.com' and 'Password' with a masked password '*****'. At the bottom right of the dialog are two buttons: 'OK' and 'Cancel'.

Figure 65: Discover Servers dialog

2. Identify the IP address range to search by adding a beginning and ending IP address in the *Begin* and *End* fields.
Neverfail recommends leaving the *Port Number* field with the default port unless the default port is in use by another application and a custom port has been configured.
3. Add a username and password used to connect to Neverfail Heartbeat in the *Username* and *Password* fields.

Note: *If the username is a domain account, use the following format: username@domain.xxx*

4. Click **OK** to run Neverfail Heartbeat server discovery.
Neverfail vSphere Client displays all Neverfail Clusters and Groups discovered.

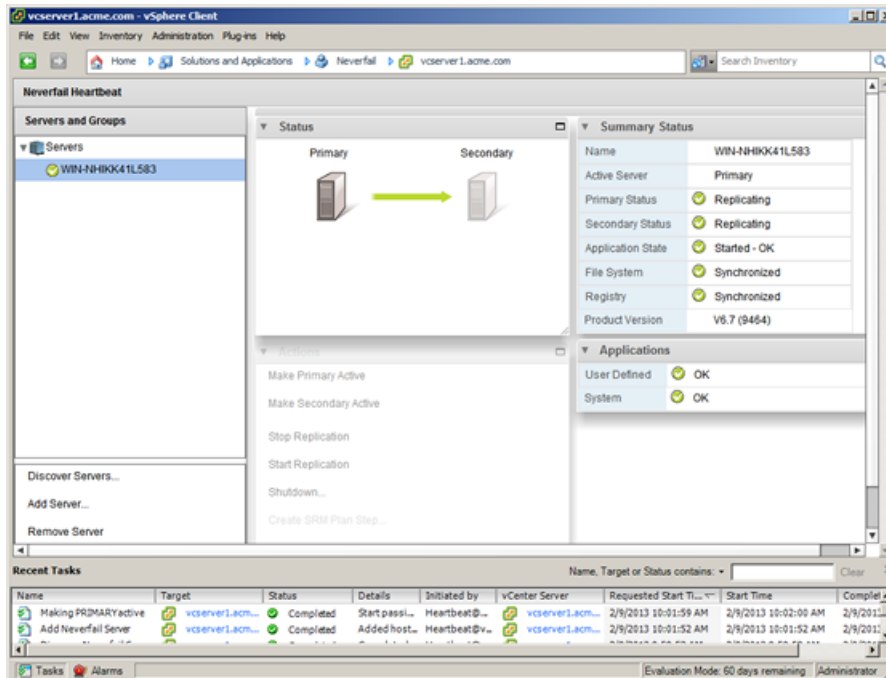


Figure 66: Neverfail Cluster in Neverfail vSphere Client

Managing Neverfail Heartbeat with Neverfail vSphere Client

Neverfail vSphere Client allows administrators to manage Neverfail Heartbeat Clusters and Groups similar to the Neverfail Heartbeat Management Client. Neverfail vSphere Client provides the ability to perform a switchover, Start Replication, Stop Replication, and Shutdown Neverfail Heartbeat.

Additionally, Neverfail vSphere Client identifies the current active server and provides the Replication status of the servers in the Cluster, the Application State, and the File System and Registry State.

***Note:** Neverfail Tasks and Events are not logged in vCenter 4.0 and earlier so will not appear in the vSphere Client.*

Perform a Switchover

- To make the Primary server of the Neverfail Cluster active, click the **Make Primary Active** button. The **Make Primary Active** dialog asks you to verify that you want to make the Primary server active. Click **OK** to make the Primary Server Active.
- To make the Secondary server of the Neverfail Cluster active, click the **Make Secondary Active** button. The **Make Secondary Active** dialog asks you to verify that you want to make the Secondary server active. Click **OK** to make the Secondary Server Active.
- If a Tertiary server is installed, you can make the Tertiary server active by click the **Make Tertiary Active** button. The **Make Tertiary Active** dialog asks you to verify that you want to make the Tertiary server active. Click **OK** to make the Tertiary server active.

Start Replication

When replication is stopped, click the **Start Replication** to initiate replication between the servers. Neverfail Heartbeat responds by starting replication between the servers.

Stop Replication

To stop replication, click the **Stop Replication** button. The **Stop Replication** dialog asks you to verify that you want to stop replication. Click **OK** to stop replication.

Shutdown Neverfail Heartbeat

To shutdown Neverfail Heartbeat, click the **Shutdown** button. The **Shutdown Options** dialog is displayed. Select one or more servers in the Neverfail Cluster to shutdown. Click **OK** to stop Neverfail Heartbeat on the selected servers in the Cluster.

Add Neverfail Clusters

1. Neverfail vSphere Client allows you to add individual Clusters which may be part of a Group. Click **Add Server** in the left pane of Neverfail vSphere Client to add a server. The **Add Server** dialog is displayed.

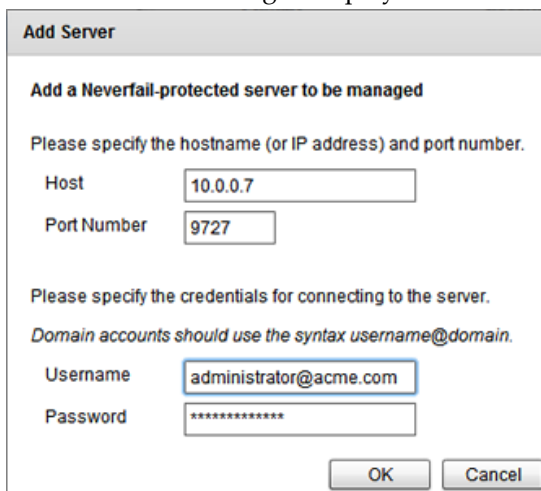


Figure 67: Add Server dialog

2. Enter the hostname or IP address of server to be added in the *Host* field. Neverfail recommends leaving the *Port Number* field with the default port unless the default port is in use by another application and a custom port has been configured.
3. Add a username and password used to connect to Neverfail Heartbeat in the *Username* and *Password* fields.

Note: *If the username is a domain account, use the following format: username@domain.xxx.*

4. Click **OK** to add the Neverfail Cluster or Group. Neverfail vSphere Client adds the Neverfail Heartbeat Cluster or Group to the left pane of the Neverfail Heartbeat page.

Remove Neverfail Clusters

Neverfail vSphere Client allows the removal of specific Neverfail Clusters from the Neverfail vSphere Client using the steps below.

1. Select the server Cluster to be removed from Neverfail vSphere Client.
2. Click **Remove Server**.
The intended Neverfail Cluster is removed from Neverfail vSphere client.

Note: To define and configure Groups, you must use the Neverfail Heartbeat Management Client.

Neverfail SRMXtender

Neverfail's SRMXtender works to extend capabilities of VMware's Site Recovery Manager (SRM). While SRM provides the ability to failover virtual servers to a secondary site, SRMXtender integrates Neverfail Heartbeat physical servers into the failover process as a natural step in the SRM Site Recovery Plan executed by SRM. SRMXtender operates as an additional feature of Neverfail vSphere Client and allows the administrator to create an SRM Step that can be added to the SRM Site Recovery Plan thereby allowing servers protected by Neverfail Heartbeat to participate in failover of servers protected by Site Recovery Manager.

Each Neverfail protected server to be included in the SRM recovery plan needs to have the SRMXtender feature enabled. For instructions on how to install Neverfail Heartbeat Management of vCenter Server, see [Installing Neverfail vSphere Client](#).

Note: If vCenter Server Heartbeat is installed on the vCenter Server, then Neverfail Heartbeat Management on vCenter Server must be installed on both the active and passive servers. If Neverfail Heartbeat Management of vCenter Server is installed on the Primary vCenter Server before installing vCenter Server Heartbeat, the cloning process used during the installation of vCenter Server Heartbeat will automatically install Neverfail Management of vCenter Server on the Secondary server.

Creating an SRM Step

Neverfail SRMXtender allows administrators to create an SRM Recovery Plan Step to integrate failover of physical Neverfail Heartbeat servers within a SRM Recovery Plan.

Prerequisites

- The Neverfail Management Component installed on vCenter Server (as described in [Neverfail SRMXtender](#)) in the Recovery and Protected Sites
- Microsoft PowerShell 2.0 installed on all SRM servers that will run SRMXtender command files, for example the SRM Servers in the Recovery and Protected sites
- The PowerShell Execution Policy must be set to *RemoteSigned* on all SRM Servers, use the following PowerShell command:

```
PS C:\> Set-ExecutionPolicy RemoteSigned
```

1. Launch VMware vSphere Client and connect to the Recovery vCenter Server.
2. Navigate to **Home > Solutions and Applications > Neverfail**.

Note: If Neverfail Solution is not available, then the Neverfail Management Component has not been installed on vCenter (refer to [Neverfail SRMXtender](#)).

3. Select a Neverfail Server in the left pane to be added to the SRM Site Recovery Plan.

Important: If the server is a member of a Group, then select the server from the Group which is to switchover first. All members of a group will switchover when a single member server receives the switchover command.

4. Click the **Create SRM Plan Step** button.
The *Create SRM Plan Step* dialog is displayed.

Note: This step requires Powershell V2 on the SRM server and the execution policy set to RemoteSigned.

Create SRM Plan Step

Create a command to make the server on the recovery or protected site active as part of an SRM plan.

Requires Powershell V2 and permission for powershell scripts to run locally without signing.

For servers which are members of Business Application Groups, all members of a group will failover or switchover together. It is recommended to add only the 'First to switch' server of a group to the SRM plan.

☒ Make Active On Recovery Site
☐ Make Active On Protected Site

You can choose to wait for the server to become active by entering a maximum time as a number of seconds greater than zero.

Maximum time to wait:

You can list alternate IP addresses by which the SRM server can reach servers which are passive. Multiple addresses are separated by commas.

Alternate IP addresses:

Use the Save As... button to save it as a batch file. The command should then be added at the required point in the SRM plan.

Figure 68: Create SRM Plan Step

5. Select the server to be controlled by the SRM Plan. This depends on which server is located at the site for which you are creating a plan. To make the server active on either site, you will require two scripts - one for each option.

Note: If the SRM Plan Step is being created on the site where the Primary server is located, select *Make Primary Server Active*. If the SRM Plan Step is being created on the site where the Secondary (or Tertiary) server is located, select *Make Secondary (or Tertiary) server active*.

6. If you want the SRM plan to wait for the Neverfail server to switchover and become active before the plan continues with the next step, enter the number of seconds to wait in the *Maximum time to wait* field.

Note: If the *Maximum time to wait* is set to zero, execution of the SRM Plan will continue without waiting for the Neverfail server to become active.

7. Management IP addresses are configured on each server in the Neverfail Cluster so that SRM can switch the servers even when the Protected Site cannot be contacted, for example in times of disaster. Enter the Management IP address that will be used by SRM to contact the Neverfail server in the *Alternate IP addresses* field, separate multiple IP addresses with a comma.

8. If you want to log the script output to a file on the SRM server, enter a path in the *Log file for command:* field (recommended for SRM 5.0), otherwise, leave the field blank.
9. Generate two scripts using the SRMXtender Plug-in.
 - a) Generate one script with *Make Primary Server Active* selected.
 - b) Generate one script with *Make Secondary (or Tertiary) Server Active* selected.
10. The scripts should be saved as .bat files with each being saved to a file share on the SRM server in the same site as the server being made active. Click the **Save As** button to save the script as a .bat file.

Note: For SRM 5.0, the scripts must have identical names and locations on each SRM server.

11. Navigate to **Home > Solutions and Applications > Site Recovery Manager** and select the intended **Recovery Plan**.
12. Select the *Recovery Steps* tab.

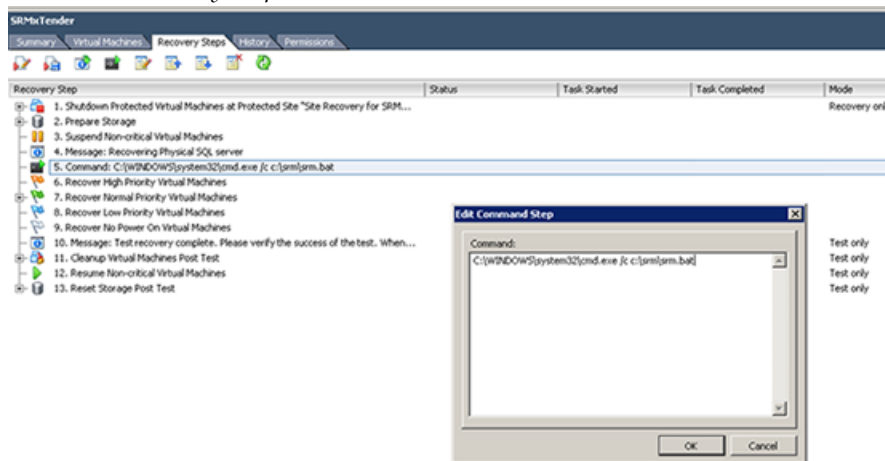


Figure 69: SRM Edit Command Step

13. Add a *Command Step* at the desired point in the Recovery Plan, for example before the *Recover High Priority Machines* Step if the applications running on these servers depend upon the physical server.
14. In the **Add Command Step** dialog enter:

C:\WINDOWS\system32\cmd.exe /c <path_to_saved_file>\<file_name>.bat

Note: <path_ to_saved_file> is the path where you have copied the \<file_name>.bat file at step 10.

15. Click **OK**.

Note: Repeat the step creation process for each Neverfail Cluster that is to participate in the Site Recovery Plan.

Uninstall SRMXtender

Under normal operations it is not necessary to uninstall SRMXtender. In the event that you must uninstall SRMXtender, perform the following steps:

-
1. Navigate to **Start > Neverfail > Uninstall or Modify**.
 2. Select the option to *Uninstall Neverfail Management for vCenter Server* and click **Next**.
Neverfail Management for vCenter Server is uninstalled.
 3. Launch a web browser and navigate to `http://<vcenter_server_name or IP>/mob`
 4. Click on *Content*.

Note: The *SupportLogs* directory is also left behind. This is intentional and should not be deleted in the event you need to submit a support report.

5. Click on *ExtensionManager*.
6. In the *Properties* pane, identify the values `extensionlist["com.neverfail.srmextender"]` and `extensionlist["com.neverfail.management"]`
7. In the *Methods* pane, click the **UnregisterExtension** option and a new window will appear.
8. In the *Value* field, type `com.neverfail.srmextender` and click **Invoke Method** to remove the plug-in.
9. In the *Value* field, type `com.neverfail.management` and click **Invoke Method** to remove the plug-in.
10. Close the pop-up window.
11. Refresh the *Managed Object Type: ManagedObjectReference:ExtensionManager* window and the plug-in should be removed from the list.

Upgrading SRMXtender

Upgrading SRMXtender is a simple multi-step process and requires you to uninstall and reinstall SRMXtender and then rerun server discovery.

1. Uninstall SRMXtender using the procedure found in [Uninstall SRMXtender](#).
2. To install the new version of the SRMXtender Plug-in, follow the procedure [Installing Neverfail vSphere Client](#).
3. Launch the Neverfail vSphere Client using the procedure described in [Launch Neverfail vSphere Client](#)
4. Rerun discovery of Neverfail Heartbeat Servers using the procedure described in [Discover Neverfail Heartbeat Servers](#)

Protection

Chapter 4

Server Protection

Overview

Protection against operating system or hardware failure affecting the active server is facilitated by multiple instances of Neverfail Heartbeat that monitor one another by sending “I am alive” messages and reciprocating with acknowledgments over the Neverfail Channel. If a passive server detects that this process (heartbeat) has failed, an automatic-failover is initiated.

Monitoring the Status of Servers

The **Server Monitoring** page provides information about the status of communications between the servers within the Cluster. 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, a red-dashed channel icon indicates that communications are not operational between the indicated servers, and an orange icon with an exclamation mark on it indicates that the channel has just disconnected and Neverfail Heartbeat will wait for the configured amount of time before determining that the channel is disconnected. In addition to the heartbeat sent between the servers, Neverfail Heartbeat also sends a ping to ensure that the servers remain visible to one another.

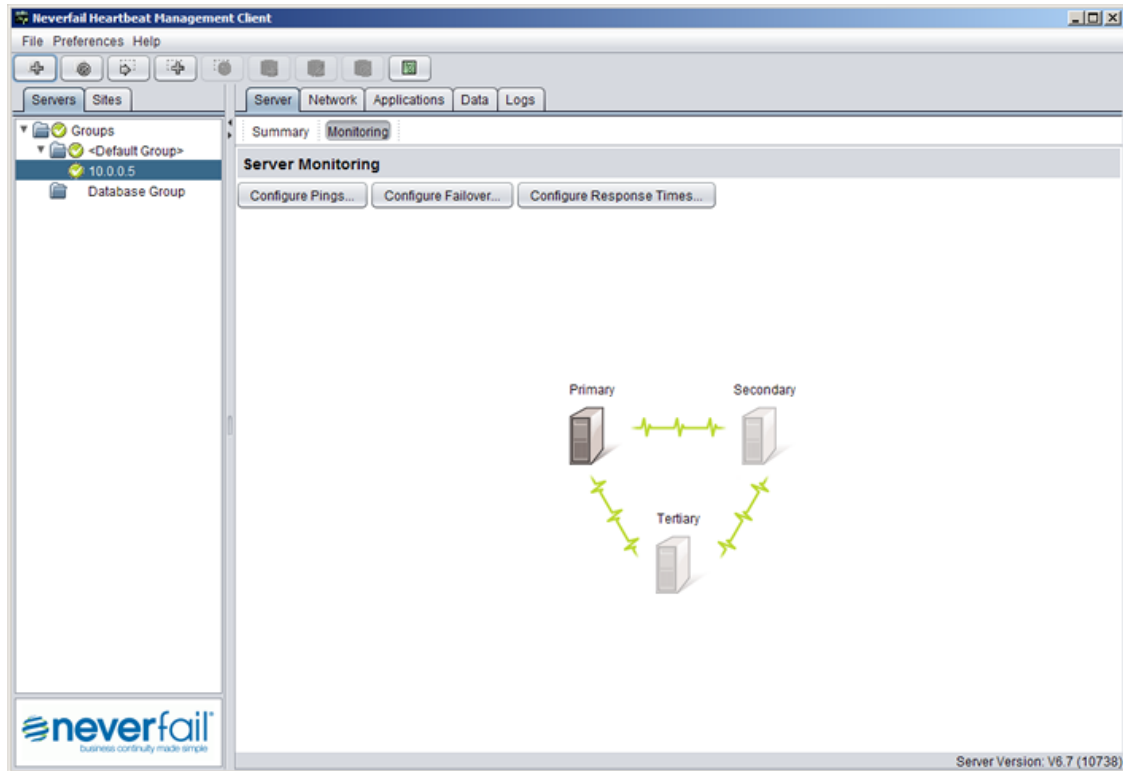


Figure 70: Server Monitoring page

Configure Neverfail Heartbeat Settings

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

Configure Pings

The **Server Monitoring Ping Configuration** dialog allows you to configure the *Ping Interval* and the *Ping Echo Timeout* used to conduct ping operations between servers. Additionally, ping routing can be configured to add additional ping targets by selecting the *Ping Routing* tab of the dialog. The IP addresses of all NICs used for the Neverfail Channel were identified during installation and do not need to be added. 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 Neverfail Heartbeat to send pings across the Neverfail Channel and the Principal (Public) Network in addition to the heartbeat ("I am alive" messages) to confirm that the server is still operational and providing service.

Procedure

- Click **Configure Pings** to open the **Server Monitoring Ping Configuration** dialog.

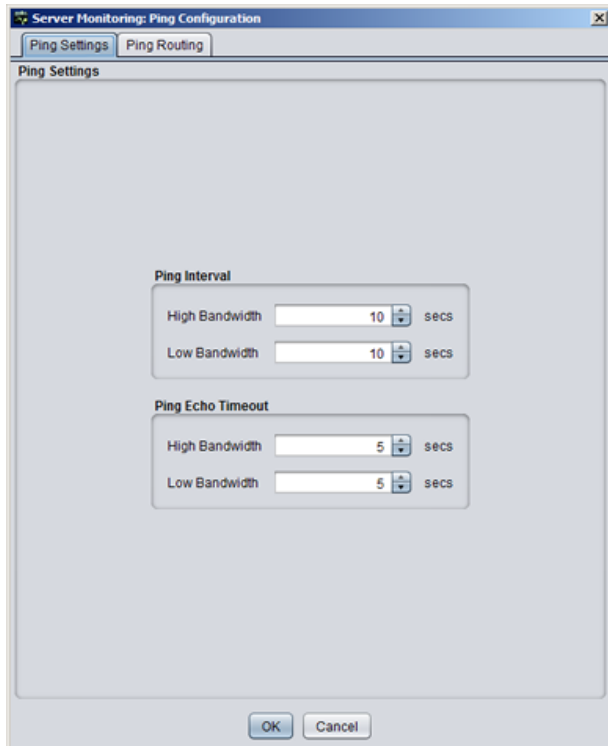


Figure 71: Server Monitoring: Ping Configuration: Ping Settings Tab

- Select the *Ping Routing* tab and enter the auxiliary IP addresses of the appropriate servers.

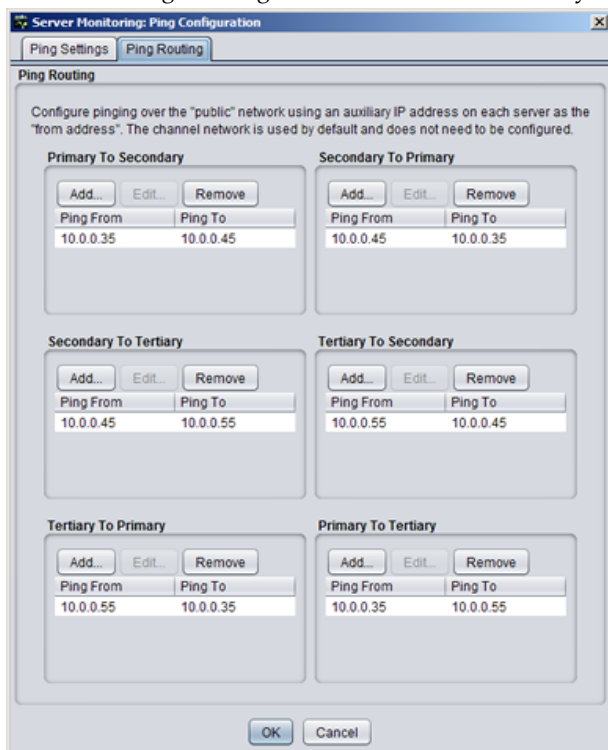


Figure 72: Server Monitoring: Ping Configuration: Ping Routing Tab

Configure Failover

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

Procedure

1. To configure the *Failover timeout*, click **Configure Failover** to open the **Server Monitoring: Failover Configuration** dialog.

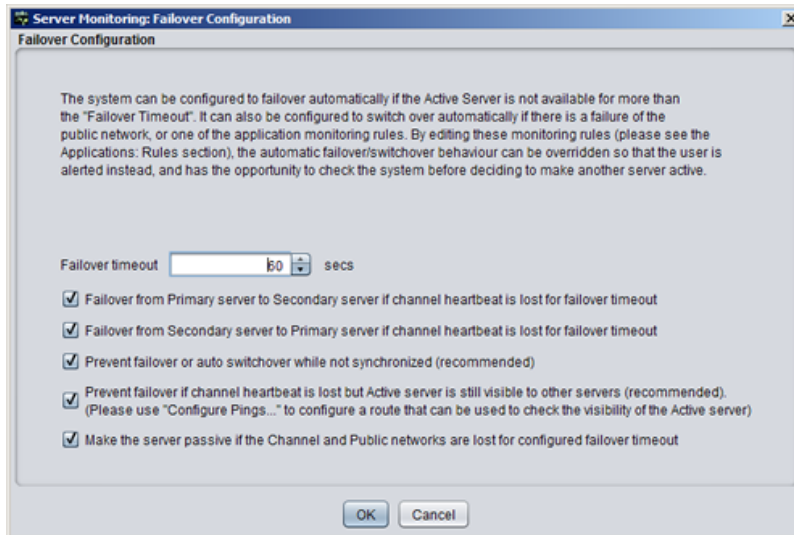


Figure 73: Server Monitoring: Failover Configuration

2. Type a new numeric value (seconds) in the *Failover timeout* text box or use the arrow buttons to set a new value.
3. Select or clear the check boxes to select the actions to take if the specified *Failover timeout* is exceeded.

Note: For more information about configuring options for failover, see [Split-brain Avoidance](#).

4. Click **OK** to accept the changes or **Cancel** to dismiss the dialog without making any changes.

Note: The default configuration for a WAN installation is with the automatic switchover (spontaneous failover) **DISABLED**. To enable Auto-switchover in a WAN pair, select **Network > Configure Auto-Switchover**, select the check box and set the missed ping failover count.

Note: In a TRIO, the installation default is set with automatic failover **ENABLED** between Primary and Secondary servers (which can be changed) and automatic failover to and from the Tertiary is **DISABLED** (and cannot be changed).

Configure Response Times

Neverfail Heartbeat also allows you to configure channel connection timeouts.

Procedure

1. Click **Configure Response Times** to open the **Server Monitoring: Response Times** dialog.
The following options are available:

- Time to wait following channel connection before starting replication
- Time to wait following channel disconnection before stopping replication
- Time for active Primary or Secondary to wait for channel connection to local peer before starting replication directly to Tertiary (when deployed).




Figure 74: Server Monitoring: Response Times

2. Type new numeric values (second) into the text boxes or use the arrow buttons to select new values.
3. Click **OK** to accept the changes or **Cancel** to dismiss the dialog without making any changes.

Common Administrative Tasks in Neverfail Heartbeat

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

- **Make Active** — Opens the **Make Active Options** dialog.




Figure 75: Make Active Options

Select a server and click **OK** to make the selected server active, or click **Cancel** to close the dialog without making any changes.

- **Shutdown** — Opens the **Shutdown Options** dialog.

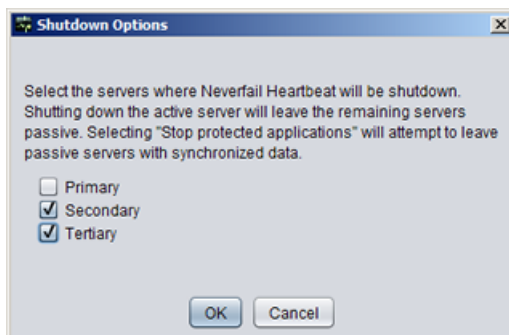


Figure 76: Shutdown Options

Select the server(s) to shut down. If you select the *Active* server (Primary, Secondary, or Tertiary), additional options to stop or not stop protected applications appear in the dialog. Click **OK** to perform the selected shutdown options, or click **Cancel** to close the dialog without shutting down.

- **Start Replication** — Opens the **Start Replication Options** dialog.

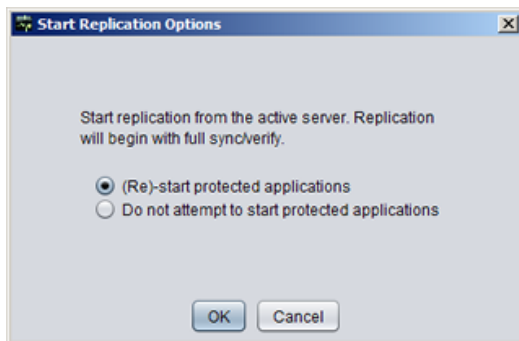


Figure 77: Start Replications Options

Select to start or not start the protected applications and click **OK** to start replication from the *Active* server. By default, all protection modes start when Neverfail 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.

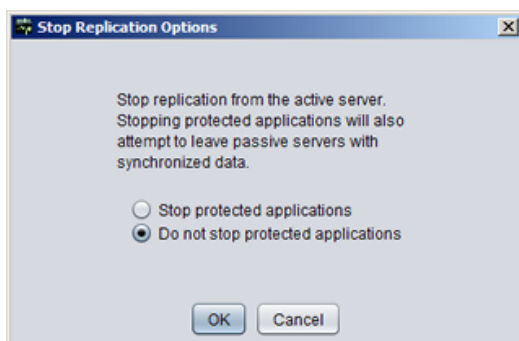


Figure 78: Stop Replication Options

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 Neverfail 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.




Figure 79: Configure (Shutdown)

Select the radio button corresponding to whether you want to stop or leave running the protected applications when Neverfail 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** to accept the changes or **Cancel** to dismiss the dialog without making any changes.

Forcing a Switchover

After Neverfail Heartbeat is configured to protect all required applications and data, it allows the Secondary or Tertiary server (if deployed) to 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.

Prior to performing work on the Primary server, a managed switchover can be triggered by selecting the server to make active and then clicking **Make Active** in the **Server: Summary** page. This changes the server roles such that the active server becomes passive and the selected server becomes active. This action also changes the replication chain depending on which 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.*

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 Neverfail Heartbeat Management Client.

Failover versus Switchover

Do not confuse a failover with a switchover.

A switchover is a controlled switch (initiated from the Neverfail Heartbeat Management Client or automatically by Neverfail Heartbeat when pre-configured) between the active and passive servers. A failover may happen when any of the following fail on the active server: power, hardware, or Channel communications. The passive server waits a pre-configured period of 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.

Configuring Failover and Active Server Isolation

Neverfail Heartbeat continuously monitors the servers in the Cluster and the network to ensure availability and uses native logic and a combination of elapsed time, administrator configured rules, current server network status, and configured ping routing to determine if failover or isolation of the active server is warranted should the servers experience missed heartbeats.

Procedure

To configure failover:

Note: For information on configuring ping routing, see [Configure Pings](#) and [Configure Public Network Monitoring](#).

1. Navigate to **Server: Monitoring > Configure Failover** to open the **Server Monitoring: Failover Configuration** dialog.
2. The *Failover timeout* can be customized by changing the default value (60 seconds) to a custom value. Type a new numeric value (seconds) in the *Failover timeout* text box or use the arrow buttons to configure how long Neverfail Heartbeat waits for a missed heartbeat before it takes a pre-configured action to failover or isolate the active server from the network.
3. Select or clear check boxes for the items listed below to select the actions to take if the specified *Failover timeout* is exceeded.
When the configured *Failover timeout* value has elapsed, Neverfail Heartbeat will evaluate, in order, the following pre-configured rules before taking action:

Note: If a rule is not selected, Neverfail Heartbeat will skip the rule and move to the next rule in the list. After all selected rules have been evaluated Neverfail Heartbeat will take action.

- Failover from Primary server to Secondary server if channel heartbeat is lost for failover timeout
- Failover from Secondary server to Primary server if channel heartbeat is lost for failover timeout
- Prevent failover or auto switchover while not synchronized
- Prevent Failover if channel heartbeat is lost but Active server is still visible to other servers
- Make the server passive if the Channel and Public networks are lost for the configured failover timeout

Note: You must configure Management IP addresses on the Principal (Public) network cards of each server to allow the passive server to send a ping via the Principal (Public) network. 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. For information about how to configure Management IP addresses, see [Configuring Management IP Addressing](#).

4. Click **OK**.

Important: If either *Server: Monitoring Ping Routing* or *Network Monitoring Ping Routing* is misconfigured, unpredictable behavior can occur.

Typical Failover and Active Server Isolation Scenarios

The following scenarios assume that Neverfail Heartbeat is deployed in a LAN with all rules selected in the **Server: Monitoring > Configure Failover > Failover Configuration** dialog.

Failover

The following scenario assumes the active server has failed and is no longer available.

Upon detection of missed heartbeats, Neverfail Heartbeat on the passive server performs the following steps:

-
1. As soon as the passive server detects that the Neverfail Channel is experiencing missed heartbeats, it will determine it itself is a valid failover target to the currently active server.
 2. As soon as the passive server detects that the Neverfail Channel is experiencing missed heartbeats. It will attempt to ping the active server's Management IP address via the Public network using the passive server's NIC configured with the Management IP address. If the ping is successful, the passive server will veto the failover. If the ping is unsuccessful, it will continue to the next step.

Note: *Since the passive server assumes that active server has failed, the passive server will not attempt to verify synchronization with the active server.*

3. At this point, the passive server checks the configured value of the *Failover timeout* and starts a "Heartbeat lost" countdown. The passive server continues with the next step.
4. At this point, failover to the passive server is postponed until the value of the *Failover timeout* has elapsed.
5. The passive server changes its role to active, removes the packet filter, and starts all services.
6. As the new active server, it will begin accepting traffic from clients.

Active Server Isolation

The figure below illustrates a scenario where the active server has lost connection with the passive server via the Neverfail Channel.



Figure 80: Network Isolation Workflow Diagram

Upon detection of missed heartbeats Neverfail Heartbeat performs the following steps:

1. As soon as the active server detects that the Neverfail Channel is experiencing missed heartbeats, it will determine *if a valid failover target (the passive server) is present*.

Simultaneously, once the passive server detects missed heartbeats, it will determine *if it is a valid failover target*.

2. Next, the active server will determine if it is synchronized with the failover target (the passive server). If synchronized, it will continue to the next step. If it is not synchronized, it will veto a failover.

Simultaneously, the passive server checks to see if it is synchronized with the active server. If synchronized, it will continue to the next step. If it is not synchronized, it will veto a failover.

3. At this point, both the active and passive servers check the configured value of the *Failover timeout* and start a "Heartbeat lost" countdown. Both servers should start the countdown at approximately the same time.
4. Failover or isolation of the active server is postponed until the configured *Failover timeout* value (in seconds) has elapsed and it is during this period that both servers accomplish steps 1 & 2.

-
5. Once the configured *Failover timeout* period has elapsed, the active server assumes the Neverfail Channel is lost and will attempt to ping the failover target (passive server) via the Public network. If the ping is successful, active server isolation is vetoed. If the attempt to ping the failover target is unsuccessful, the active server will proceed to the next step.

Simultaneously, the passive server assumes the Neverfail Channel is lost and attempts to ping the active server via the Public network. If the ping is successful, failover is vetoed. If the ping attempt is unsuccessful, the passive server proceeds to the next step.

Note: *If the servers have reached this point, then neither server can see the other server.*

6. The active server checks only its own network connectivity to the Public network. If the active server has lost connectivity to the Public network, it will isolate itself by making itself passive (potential active).
7. Both the active and passive servers will check their connectivity to the Public network. If the active server has lost connectivity to the Public network, it will isolate itself by making itself passive (potential active). Should the active server reconnect with the passive, it will become active again. Otherwise, it will remain passive. If the passive server has lost connectivity to the Public network, it will veto a failover.

Recover From a Failover

This recovery scenario is based on Neverfail Heartbeat in a configuration with the Primary server as active and the Secondary and Tertiary servers as passive.

Procedure

Note: *When failover conditions, such as a power failure, cause failures in both active and passive servers, a condition may result that causes all servers to restart in Passive mode. In this situation, manual intervention is required. See [Two Passive Servers \(In a Pair\)](#) and [Three Passive Servers \(In a Trio\)](#) for more information.*

In the following case, a failover occurred and the Secondary server is now running as the active server.

1. Review event logs on all servers to determine the cause of the failover. If you are unsure how to do this, use the *Neverfail Heartbeat Log Collector* tool to collect information and send the output to Neverfail Support.
2. If any of the following issues exist on the Primary server, performing a switchover back to the Primary server may not be possible until other important actions are carried out. Do not restart Neverfail Heartbeat until the following issues are resolved:
 - **Hard Disk Failure** – Replace the disk.
 - **Power Failure** – Restore power to the Primary server.
 - **Virus** – Clean the server of all viruses before starting Neverfail Heartbeat.
 - **Communications** – Replace or repair the physical network hardware.
 - **Blue Screen** – Determine and resolve the cause of the blue screen. This may require you to submit the Blue Screen dump file to Neverfail Support for analysis.
3. Run the **Configure Server Wizard** and verify that the server *Identity* is set to *Primary* and its *Role* is *passive*. Click **Finish** to accept the changes.
4. Disconnect the channel network cables or disable the network card.
5. Resolve the problem – list of possible failures, etc.
6. Reboot the server and reconnect or re-enable the network card.

7. After the reboot, verify 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 the Neverfail Heartbeat Management Client and confirm that the Secondary server is reporting as *active*. If the Secondary server is not displaying as *active*, follow the steps below:
 - a) If the Neverfail Heartbeat Management Client is unable to connect remotely, try running it locally. If you remain unable to connect locally then verify that the service is running via the Service Control Manager. If it is not, review the event logs to determine a cause.
 - b) Run the **Configure Server Wizard** and confirm that the server is set to Secondary and is *active*. Click **Finish** to accept the changes.

Note: If Neverfail Heartbeat is running, you can run the **Configure Server Wizard**, but you will not be able to make any changes. You must stop the Neverfail Heartbeat service before attempting to make changes via the **Configure Server Wizard**.

- c) Determine whether the protected application is accessible from clients. If it is, then start Neverfail Heartbeat on the Secondary server. If the application is not accessible, review the application logs to determine why the application is not running.

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. After Neverfail Heartbeat starts, it overwrites all protected data (configured in the File Filter list) on the Primary passive server. Contact Neverfail Support if you are not sure whether the data on the active server is 100% up to date. Go on to the next step only if you are sure that you want to overwrite the protected data on the passive server.

9. Start Neverfail Heartbeat on the Secondary active server and verify that the taskbar icon now reflects the correct status by showing **S/A** (Secondary and active).
10. Start Neverfail Heartbeat on the failed Primary server and allow the system to synchronize. When the re-synchronization is complete, you can continue running with this configuration (for example, the Secondary is the active server and the Primary is the passive server), or initiate a managed switchover.
11. Optionally, perform a managed switchover to return the Primary and Secondary servers to the same roles they had before the failover.

Split-brain Avoidance

Split-brain Avoidance ensures that only one server becomes active if the channel connection is lost, but all 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 does not failover, even if the channel connection is lost. WAN installations require different IP addresses on the Principal (Public) network for the local and remote servers.

1. To enable Split-brain Avoidance, open the **Server Monitoring** page in the Neverfail Heartbeat Management Client.
2. Click **Configure Failover**
3. Select *Prevent failover if channel heartbeat is lost but Active server is still visible to other servers (recommended)*.

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.

Note: 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.

Configuring Management IP Addressing

Management IP addresses are used to allow the passive server to communicate, because unlike the Principal (Public) IP address, they are not filtered. Management IP addresses are necessary when configuring Network Isolation Protection.

Procedure

To configure a Management IP address on the Principal (Public) network card, perform the following procedure:

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 from a subnet other than the Principal (Public) or Neverfail Channel subnet in the IP address 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 **OK**.
8. Launch the Neverfail Heartbeat Management Client and select **Server: Monitoring > Configure Pings > Ping Routing** and add the newly assigned IP addresses to the *Ping Routing* table.

Chapter 5

Network Protection

Overview

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

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

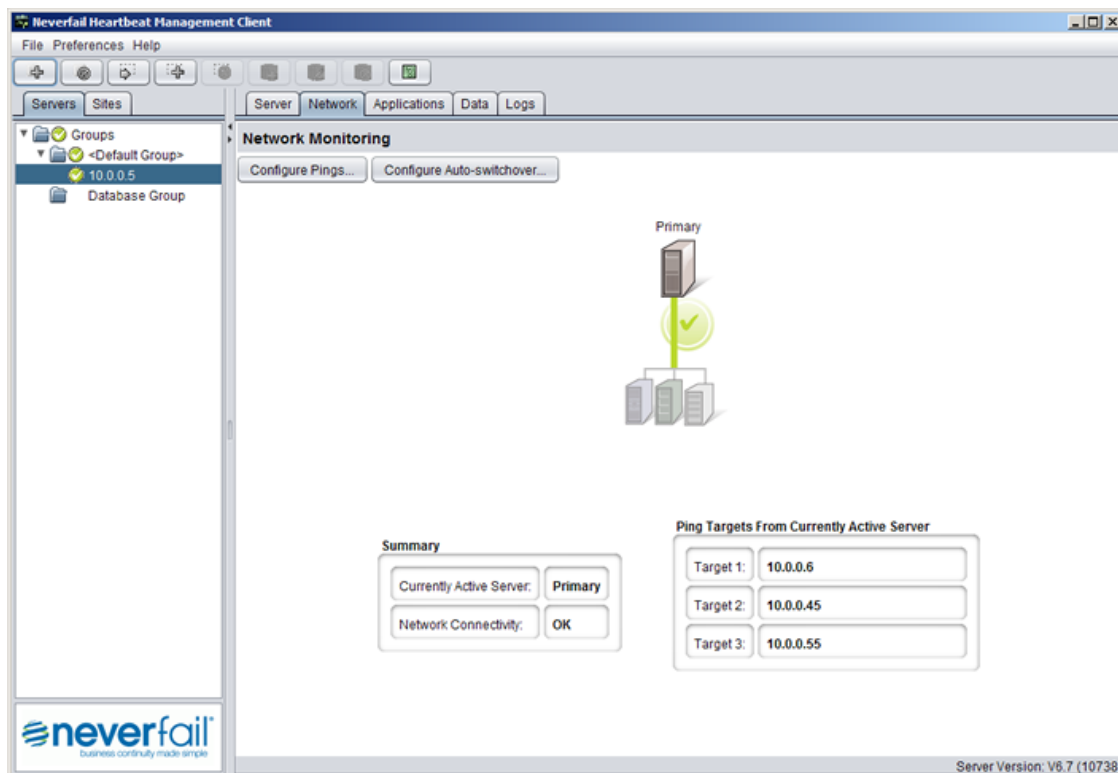


Figure 81: Network Monitoring

Configure Public Network Monitoring

The Principal (Public) network monitoring feature, previously discussed, is enabled by default during the installation of Neverfail Heartbeat. This feature integrates the polling of the particular 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

Procedure

To configure Public Network Monitoring:

1. To specify a manual target for the Principal (Public) network checking, click **Configure Pings** to invoke the **Ping Configuration** dialog.

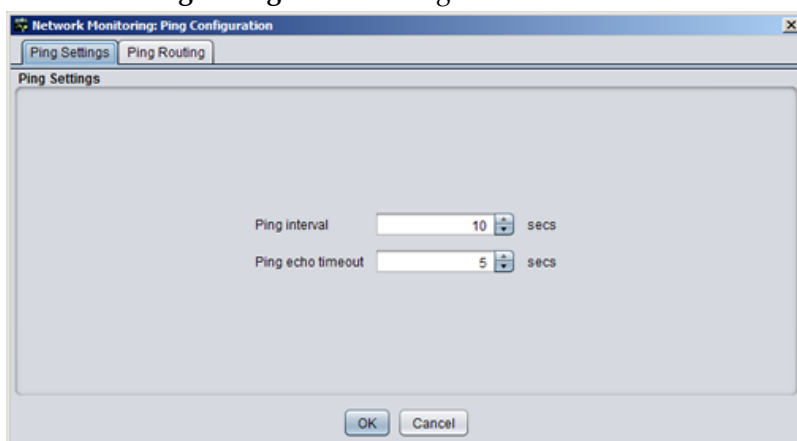


Figure 82: Network Monitoring: Ping Configuration: Ping Settings

2. Select the *Ping Routing* tab to add to or modify the existing target IP addresses for each server to ping.

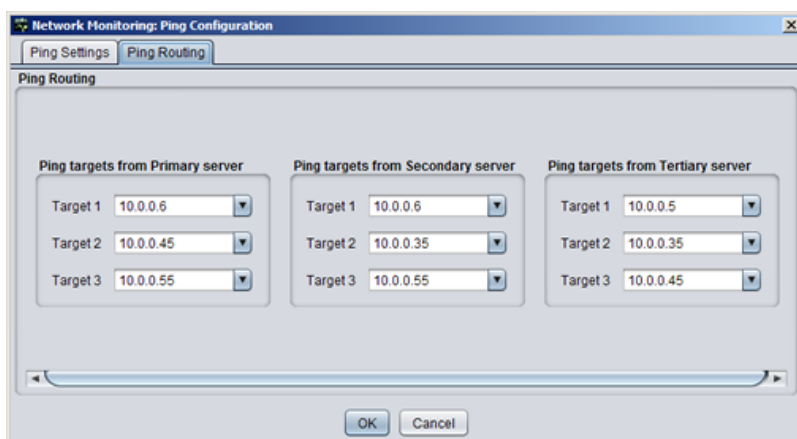


Figure 83: Network Monitoring: Ping Configuration: Ping Routing

In a WAN Pair environment, the target addresses for Principal (Public) network monitoring on the Secondary server may be different to those automatically selected on the Primary server and in a LAN-WAN Trio, the target addresses on the Tertiary server may be different from those on the

Primary and Secondary servers. 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 *Max pinged echoes missed before auto-switchover* threshold value. If the failure count of all three targets exceeds this value, Neverfail Heartbeat initiates an auto-switchover.

Enabling Automatic Switchover in a WAN

The default setting for Automatic Switchover when deployed in a WAN is *Disabled*. Should it be necessary to configure Automatic Switchover in a WAN, use the procedure below:

Procedure

To enable Automatic Switchover in a WAN:

1. In the Neverfail Heartbeat Management Client, select 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**.




Figure 84: WAN Auto-Switchover Configuration

Setting Max Server Time Difference

Neverfail Heartbeat generates a warning if the Primary and Secondary server system clocks are not synchronized. The threshold for time difference can be configured using the *Server: Summary* page.

Procedure

To set Max Server Time Difference:

1. Select the *Server: Summary* tab and click **Configure** to display the *Server: Summary Configure* dialog.
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**.




Figure 85: Server: Summary Configure

Chapter 6

Application Protection

Applications Environment

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

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

The AMFx also employs sponsorship for protected applications' 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.

Neverfail 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: Summary

The **Applications: Summary** page displays the current status of the Cluster, including the identity of the active server, the application state and health, details of application types and their corresponding running status and health. The lower portion of the page provides an *Applications Log* that allows viewing of application events as they occur.

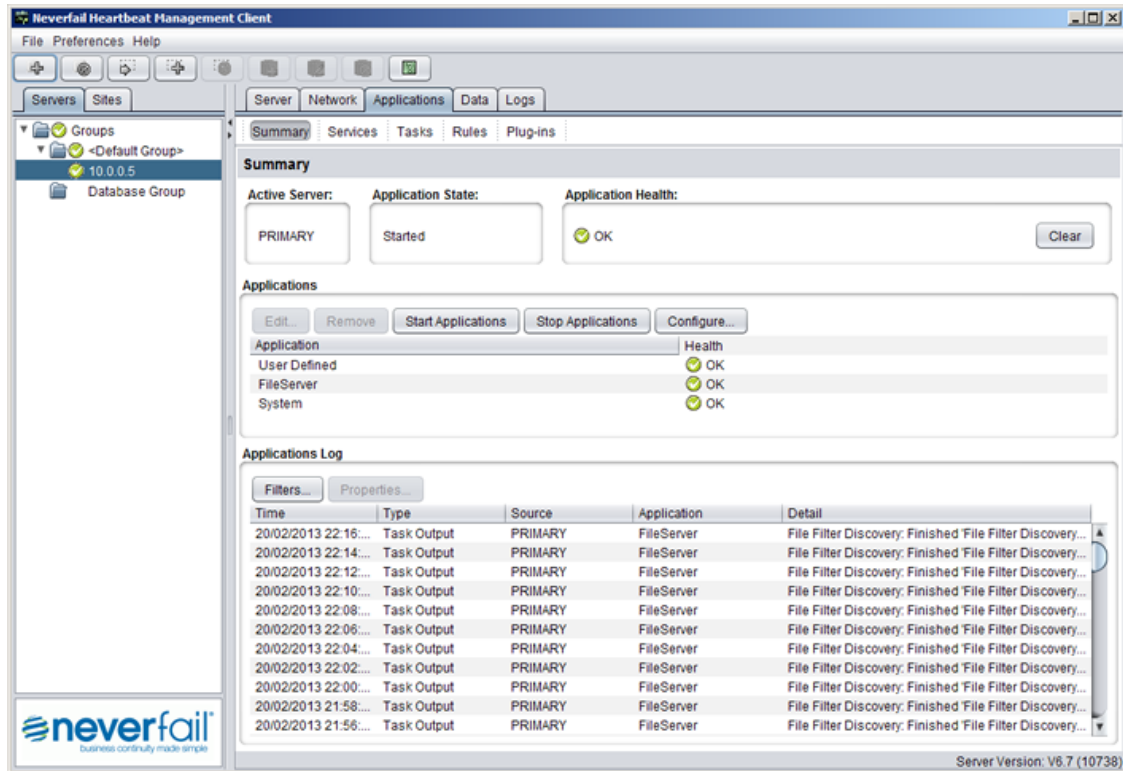


Figure 86: Applications: Summary

This page also provides controls to edit, remove, start, and stop applications, and to configure all protected applications.

View Application Status

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

Reset the Application Health Status

Procedure

The Neverfail Heartbeat Management Client monitors the health of protected applications and provides the application health status on the *Applications: Summary* page. Should the health of a protected application become degraded, the status will reflect *Degraded* and remain so, even if the health returns to *OK*, until you reset the Application Status. To reset Health Status of protected applications:

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

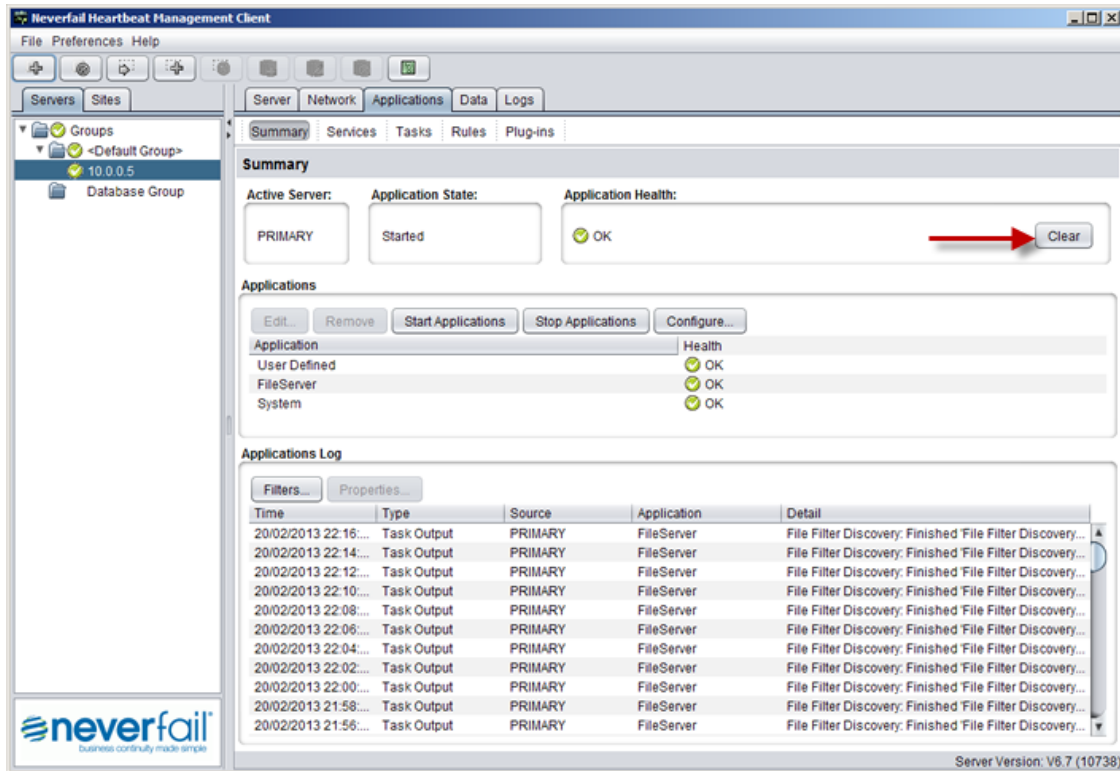


Figure 87: Clear Application Health Status

If a problem occurs (for example, a failed service or rule), the *Application Health status* becomes *Degraded*. Even if Neverfail 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* won't run.

2. 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*.

Edit Individual Applications

You can configure the amount of time to wait for applications to start or stop before taking action or reporting a failure.

Procedure

To configure these timeout settings, select the application (in the *Applications* pane) and do one of the following:

1. Right-click on the application and select **Edit** from the menu or click **Edit** at the top of the pane. The **Edit Application** dialog appears.




Figure 88: Edit Application

Note: Default application timeout settings for plug-ins is 300 sec and for user-defined applications is 180 sec.

2. Enter new values into the *Stop Timeout* and *Start Timeout* text boxes or use the arrow buttons to adjust the values (seconds).
3. Click **OK** to accept the new settings or click **Cancel** to close the dialog without making any changes.

Remove an Application

Application removal is a simple process and can be performed without having to stop Neverfail Heartbeat.

Procedure

To remove an application:

1. Select the application (in the *Applications* pane).
2. Right-click on the application and select **Remove** from the menu or click **Remove** at the top of the pane.

A confirmation message appears.

3. Click **Yes** to remove the selected application, or click **No** to dismiss the message without deleting the application.

Manually Stop and Start Applications

Neverfail Heartbeat provides the ability to Stop and Start applications from the *Applications: Summary* page of the the Neverfail Heartbeat Management Client.

Procedure

To manually Stop and Start applications:

- To stop running applications, click **Stop Applications** (at the top of the *Applications* pane). The applications stop. You can view the progress of stopping in the *Applications Log* pane.
- To start applications, click **Start Applications** (at the top of the *Applications* pane). The applications start. You can view the progress of starting in the *Applications Log* pane.

Configure Applications

You can configure protected applications and enable or disable protection and monitoring. This feature allows you to perform application maintenance without stopping Neverfail Heartbeat or taking the whole server offline. During installation, Neverfail Heartbeat creates default settings for application configurations. The Neverfail Heartbeat Management Client **Applications: Summary** page allows you to change the settings.

Procedure

To configure applications:

1. Click **Configure** (at the top of the *Applications* pane) to change these settings.

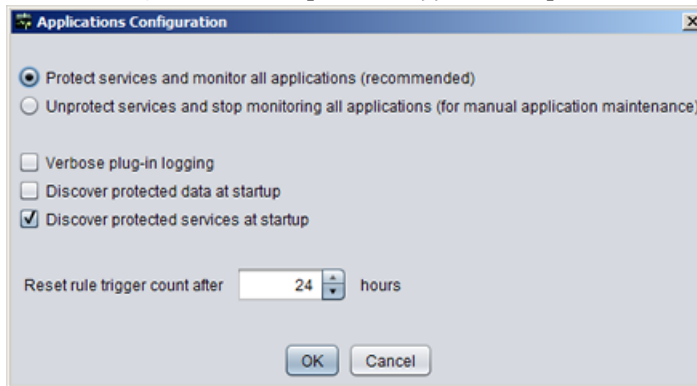


Figure 89: Applications Configuration

2. Select *Protect services and monitor all applications (recommended)* or *Unprotect services and stop monitoring all applications (for manual application maintenance)*.

Optionally select any or all of the following:

- Verbose Plug-in logging
 - Discover protected data at startup
 - Discover protected services at startup
3. Additionally, you can type a new value into the *Reset rule trigger count after* text box or use the arrow buttons to adjust the values (hours).
 4. Click **OK** to accept the new settings or click **Cancel** to close the dialog without making any changes.

View the Applications Log

The *Applications Log* is very useful in troubleshooting the protected application environment.




Figure 90: Applications Log

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

You also can filter *Applications Log* entries to reduce the number of events displayed, and use the *Applications Log* to troubleshoot application errors. For example, if an application fails, you can right-click on the associated event in the *Application Logs* and select *Properties* to open the Log and investigate the failure.

Filter Application Log Entries

By default, all events are displayed in the *Application Log* pane. To filter the events displayed, perform one of the following steps:

- Right-click on the entry and select **Filters** from the menu
- Click **Filters** at the top of the pane

The **Application Log Filters** dialog appears.

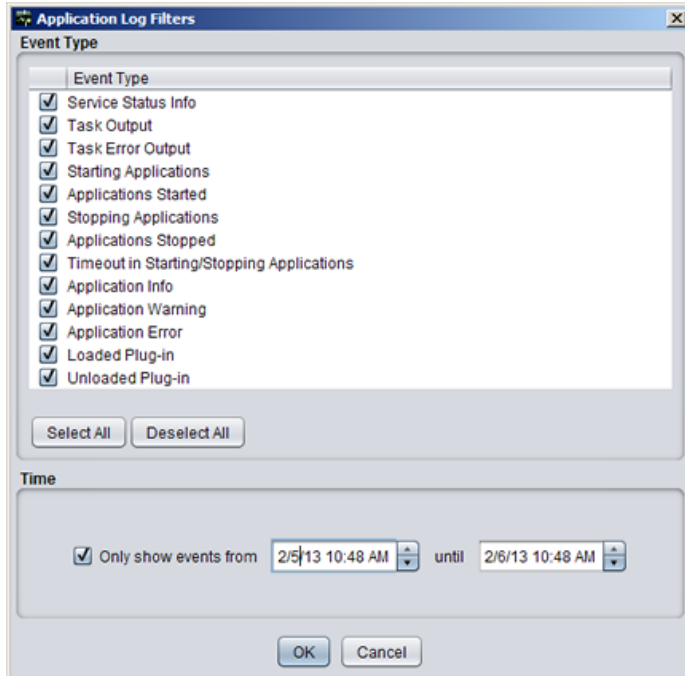


Figure 91: Application Log Filters

Use the check boxes (select to display or clear to hide) to filter *Application Log* entries by at least one *Event Type*. To display only entries within a particular time range, select the check box associated with *Only show events from* and type values into the two date/time text boxes or use the up and down arrow keys to adjust the dates and times. Click **OK** to accept the filter criteria or click **Cancel** to close the dialog without changing the filter criteria.

Applications: Services

The **Applications: Services** page shows services specified by plug-ins or by the user, and any services related by dependency.

View the Status of Services

The status of all protected services is displayed on the **Application: Services** page. The status shows both the target and actual state for the Primary, Secondary, and Tertiary (if deployed) servers and the Failure Counts for all servers.

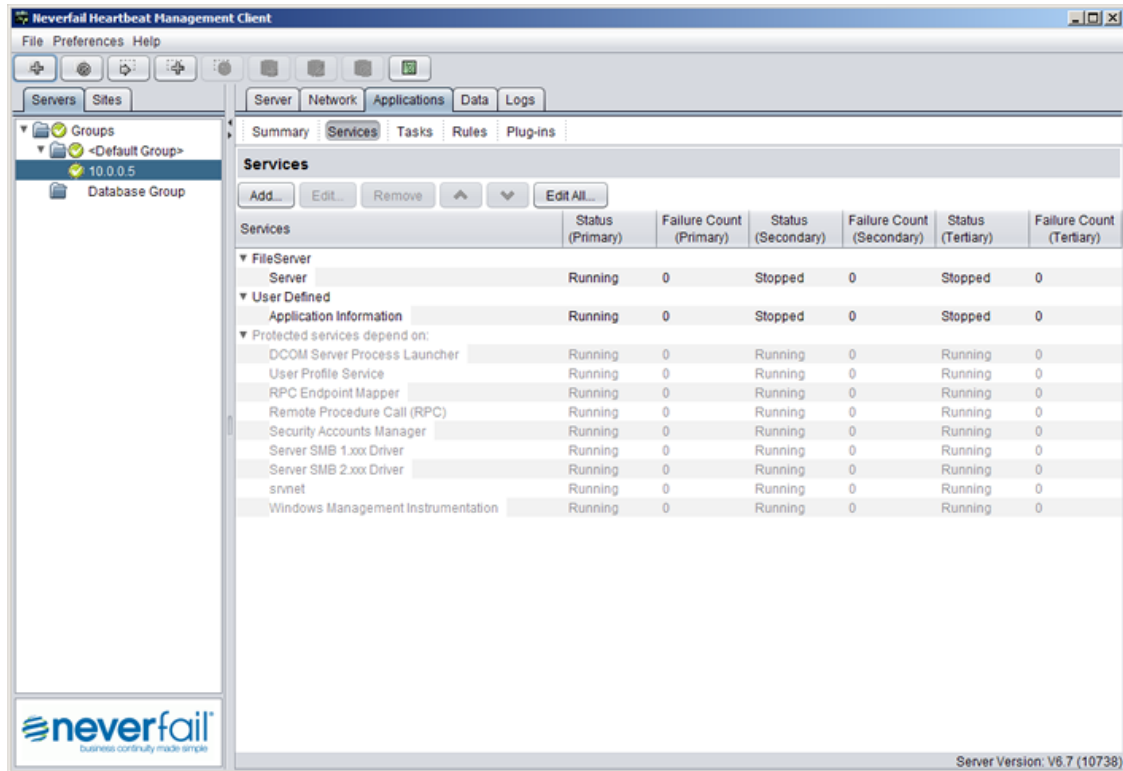


Figure 92: Applications: Services page

The target state of protected services can be specified for the active and passive server(s), and is typically *Running* on the active and *Stopped* on the passive(s). Services are protected when they are set to *Running* or *Automatic*, and otherwise are logged as unprotected. Services depending on protected services are managed (for example, started and stopped) by Neverfail Heartbeat but not monitored (for example, not restarted if stopped by some external agency). Services upon which protected services depend are monitored (for example, restarted if stopped) but not managed (for example, not stopped if protected applications are stopped).

Add a Service

To protect a service that was not automatically added by Neverfail Heartbeat during installation, the service must be added through the Neverfail Heartbeat Management Client and be in a *Running* state.

Procedure

To add a service:

1. Right-click on a service and select **Add** from the menu or click **Add** at the top of the pane.

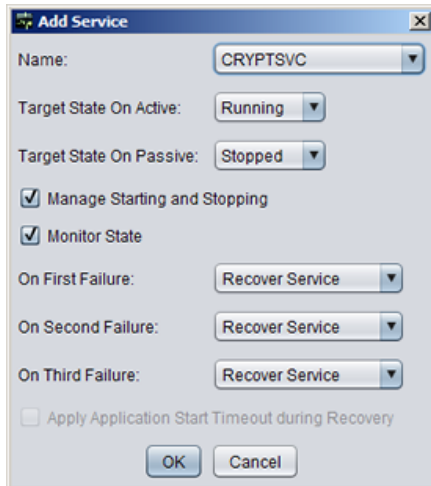


Figure 93: Add Service

2. Select the service and set the *Target State on Active* and *Target State on Passive* values. Normally, the *Target State on Active* is set to *Running* and the *Target State on Passive* is set to *Stopped*. User defined services configured with a target state of *Running* on both active and passive servers do not stop when **Stop Applications** is clicked.
3. To let Neverfail Heartbeat manage the starting and stopping of the service, select the check box. To make Neverfail Heartbeat monitor the state of the service, select the *Monitor State* check box. Neverfail Heartbeat also lets you assign three sequential tasks to perform in the event of failure. Task options include the following:
 - *Recover Service*
 - *Restart Applications*
 - *Log Warning*
 - *Switchover*
 - *Rule Action*

Note: *Rule Action tasks are additional user defined tasks previously created by the user.*

4. Assign a task to each of the three failure options and after all selections are made, click **OK** to dismiss the dialog.

Note: *When dependent services are involved, actions to take on failure should match the protected service. If a service fails and the failure option is set to *Restart Applications*, all applications are restarted.*

Edit a Service

To change the options of a protected service, select the service listed in the pane and perform the following steps:

Procedure

1. Right-click on the service and select **Edit** from the menu or click **Edit** at the top of the pane.
The **Edit Service** dialog appears, which provides a subset of same options available when a new service is added.

2. After making modifications, click **OK** to accept the changes.

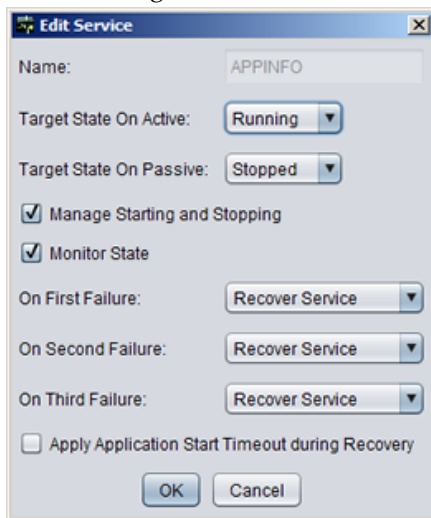


Figure 94: Edit Service

3. To unprotect a service and stop monitoring the service, click on the *Services* tab of the **Applications** page. Select the service and click **Edit**.
4. Clear the *Manage Starting and Stopping* and *Monitor State* check boxes, and then click **OK**.

Configure Service Recovery Options for All Protected Services

Neverfail Heartbeat provides the ability to configure the Service Recovery Options for all services that are protected from one central location.

Procedure

1. Navigate to the Applications: Services page.
2. Click the **Edit All** button.
Select the action to take for the 1st, 2nd, and 3rd instance of failure. The same actions selected here will be automatically configured for each protected service.



Figure 95: Applications: Services Edit All

Remove a Service

To remove a service, select the service in the pane and perform the following steps:

Procedure

1. Right-click on the service and select **Remove** from the menu or click **Remove** at the top of the pane.

A confirmation message appears.

2. Click **Yes** to remove the selected service from the protected list. Click **No** to close the message without removing anything.

Note: *If you remove a protected service that was automatically added by a plug-in (not user defined), the service will be re-added to the protected list the next time the Protected Service Discovery task runs.*

Change the Order of Services

You can change the order of services using **Up** and **Down** arrows (near the top of the page or on the right-click menu) to change the order in which they appear in the list of services. It is important to understand that the exact order in which services are started and stopped is influenced by a number of key factors:

Procedure

To change the starting and stopping order of protected services:

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

Applications: Tasks

Tasks are a generalization and extension of the start, stop, and monitor scripts in earlier versions of Neverfail Heartbeat.

Task types are determined by when the tasks are run, and include the following:

- **Network Configuration** — This is the first type of task run when applications are started, and is intended to launch `Dnscmd`, `DNSUpdate` or other network tasks. Where multiple `DNScmds` are required, these can be contained in a batch script, which is then launched by the task. Network Configuration tasks are the only types of task that can vary between Primary, Secondary, and Tertiary servers.
- **Periodic** — These tasks are run at specific configurable intervals.
- **Pre/Post Start** — These tasks are run before and after services are started on the active server.
- **Pre/Post Stop** — These tasks are run before and after services are stopped on the active server.
- **Pre/Post Shadow** — These tasks are run before and after a shadow copy is created on the active server by the Data Rollback Module.
- **Rule Action** — These tasks can be configured to run in response to a triggered rule, or when a service fails its check.

Tasks can be defined and implemented by plug-ins or by the user, or they can be built-in tasks defined by Neverfail Heartbeat. User defined tasks are implemented as command lines, which can include launching a batch script. Examples of built-in tasks include monitoring a protected service state on the active and passive servers. An example of a plug-in-defined task is the discovery of protected data and services for a particular application.

The Neverfail Heartbeat Management Client **Applications: Tasks** page provides a list of tasks and associated status information, as well as features to quickly manage tasks.

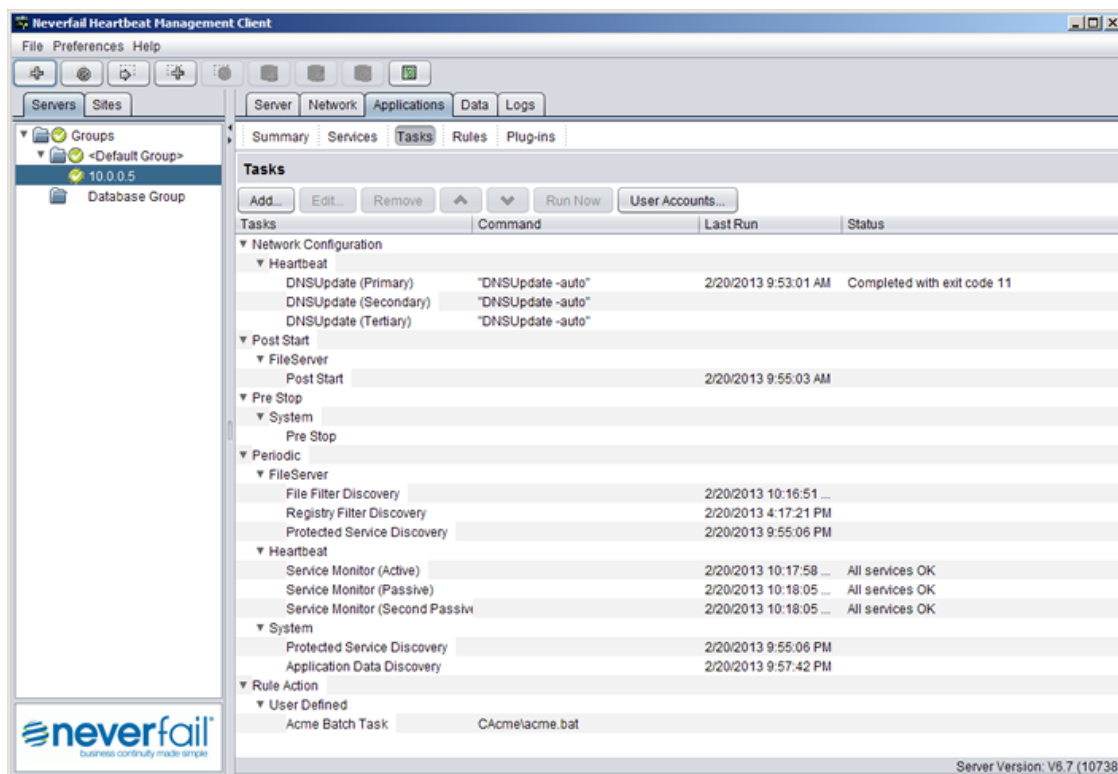


Figure 96: Applications: Tasks page

Add a Task

Tasks can be added from the *Applications: Tasks* page of the Neverfail Heartbeat Management Client. To add a User Defined task:

Procedure

1. Right-click on an existing task and select **Add** from the menu or click **Add** at the top of the pane. The **Add Task** dialog appears.

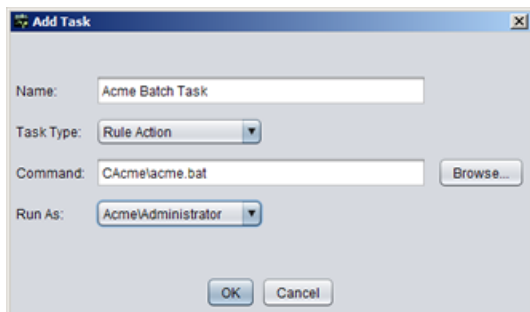


Figure 97: Add Task

2. Type a *Name* for the task into the text box.
3. Select the Task Type from the drop down list. *Task types* include: *Network Configuration*, *Periodic*, *Pre/Post Start*, *Pre/Post Stop*, *Pre/Post Shadow*, and *Rule Action*.
4. Select the identity of the server the task *Runs On* (*Primary*, *Secondary*, *Tertiary*).

5. 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](#).

6. Select from the options presented in the *Run As* drop down list (typically includes local and administrator accounts).
7. Click **OK** to add the task, or **Cancel** to exit the dialog without adding the task.

Edit a Task

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

Procedure

1. Right-click on an existing task and select **Edit** from the menu or click **Edit** at the top of the pane.

The **Edit Task** dialog appears. The parameters available to edit vary according to the task type. The following are two examples.

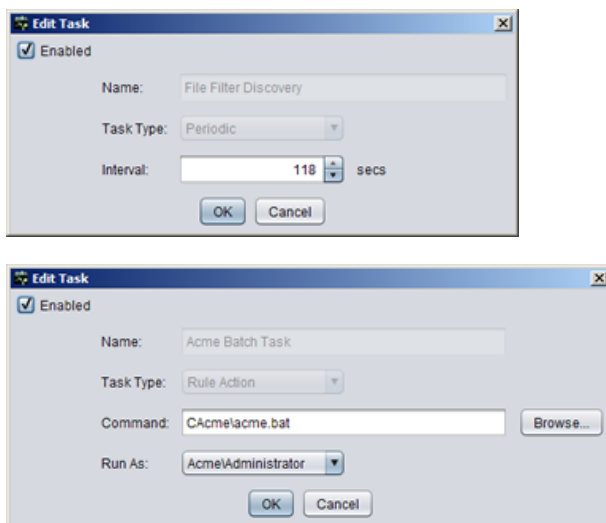


Figure 98: Edit Task

2. After completing edits of the task, click **OK** to accept the settings and dismiss the dialog.

Remove a Task

To remove a task, select the task from the list and perform the following steps:

Procedure

1. Right-click on an existing task and select **Remove** from the menu or click **Remove** at the top of the pane.
A confirmation message appears.
2. Click **Yes** to remove the task, or click **No** to close the message without removing the task.

Change the Order of Tasks

You can change the order of tasks using **Up** and **Down** arrows (near the top of the page or on the right-click menu) to change the order in which they appear in the list of tasks.

Manually Start a Task

When manually starting a task, you have the option to wait for a designated period or event to occur before launching the task, or to launch the task immediately. To launch a task immediately, select the task from the list and perform the following steps:

Procedure

- Right-click on the existing task and select **Run Now** from the menu or click **Run Now** at the top of the pane
The task runs. You can watch the *Status* column of the *Tasks* list for messages as the task runs to completion.

View, Add, and Remove User Accounts

You can view, add, and remove user accounts through the Neverfail Heartbeat Management Client.

Procedure

- Click **User Accounts** (near the top of the **Applications: Tasks** page).
The **User Accounts** dialog appears.




Figure 99: User Accounts

The **User Accounts** dialog contains a list of all currently configured user accounts, including *Username*, *Domain*, and *Checked* (username/password credential validation) status.

Add User Account

To add a user account:

Procedure

- Click **Add**.
The **Add User** dialog appears.




Figure 100: Add User Account

2. Type the name of the *User*, the associated *Domain*, and a *Password* into the corresponding text boxes.
3. Click **OK** to add the new user, or click **Cancel** to close the dialog without adding the user.

Note: *Because this information is used for executing tasks that require credentials, be sure to populate these fields with information identical to the Windows credentials.*

Remove User Account

To Remove a user, select the user account from the list in **User Accounts** dialog.

Procedure

1. Click **Remove**.
A confirmation message appears.
2. Click **Yes** to remove the user, or click **No** to close the dialog without removing the user.

Application: Rules

Rules are implemented by plug-ins. The use of rules is covered in detail in [Rules Overview](#).

Applications: Plug-ins

Plug-ins are installed to support specific applications and contain all of the components to protect the designated application.

Plug-ins are designed to start and stop the application, monitor the application, and provide all rules necessary to ensure an application stays available to users in the event of a failure by initiating an auto-switchover when configured.

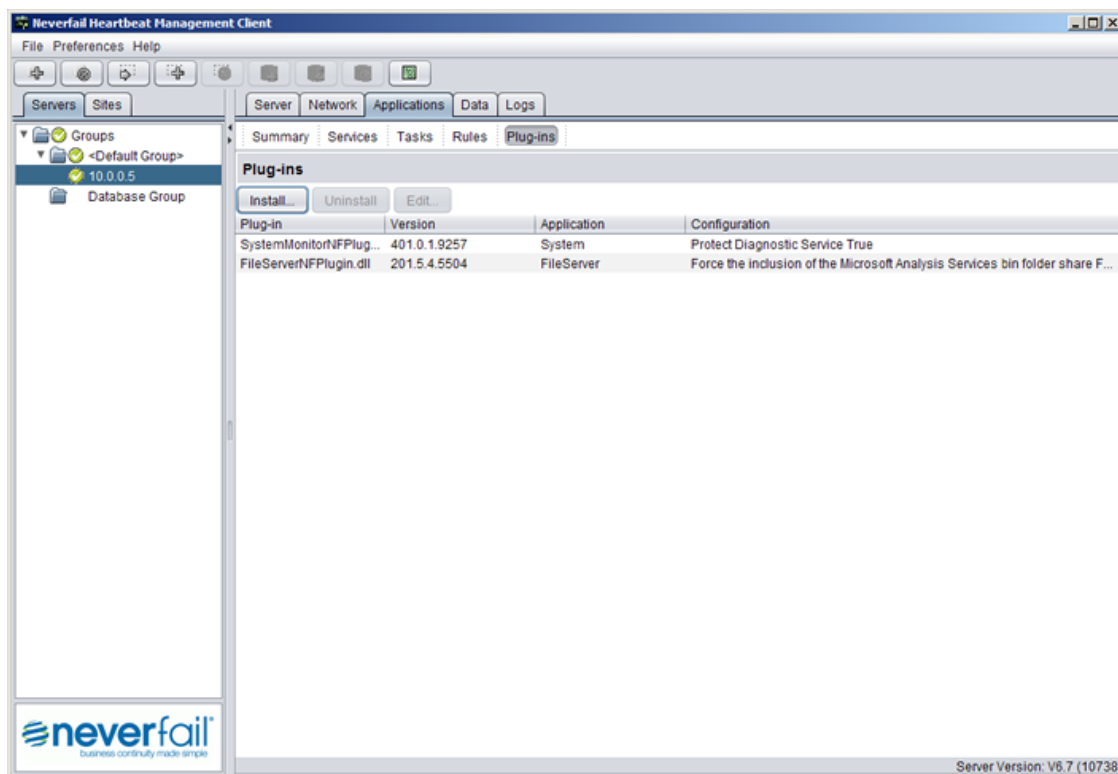


Figure 101: Applications: Plug-ins page

Install a Plug-in

To install a new plug-in, perform the following steps:

Procedure

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.

1. Right-click on an existing plug-in from the Plug-ins list and select **Install** from the menu or click **Install** at the top of the page.
The **Install Plug-in** dialog appears.




Figure 102: Install Plug-in

2. Type a path to the plug-in location (case-sensitive), or click **Browse** and navigate to the plug-in (recommended).

3. Click **OK** to install the plug-in or click **Cancel** to close the dialog without installing a plug-in.

Edit a Plug-in

You can edit the configuration of user installed plug-ins. To edit an existing plug-in, select it from the Plug-ins list and perform the following steps:

Procedure

1. Right-click on an existing plug-in from the *Plug-ins* list and select **Edit** from the menu or click **Edit** at the top of the page.
The **Edit Plug-in** dialog appears.

***Note:** Configuration options are specific to each plug-in and must be reviewed before making modifications.*




Figure 103: Edit Plug-in

2. Click **OK** to save the changes to the plug-in configuration, or click **Cancel** to close the dialog without making any changes.

Uninstall a Plug-in

Plug-ins may be uninstalled for multiple reasons, such as when upgrading or removing the application that the plug-in protects, or when directed by Neverfail Support. To uninstall a plug-in:

Procedure

1. Select the plug-in from the *Plug-ins* list.
2. Right-click on the plug-in and select **Uninstall** from the menu or click **Uninstall** at the top of the page

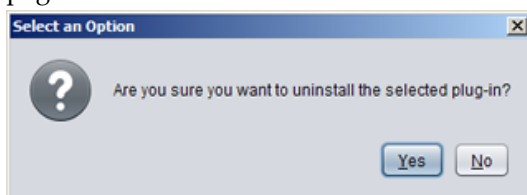


Figure 104: Remove Plug-in

A confirmation message appears.

3. Click **Yes** to uninstall the plug-in, or click **No** to close the message without uninstalling the selected plug-in.

Chapter 7

Performance Protection

Rules Overview

Rules are implemented by plug-ins (there are no user-defined rules). Rules can be either timed (they must evaluate as true continuously for the specified duration to trigger) or latched (they trigger as soon as they evaluate to true). Rules can be configured with rule actions, which are the tasks to perform when the rule triggers.

Rules use the following control and decision criteria for evaluation:

- *Name*: (the name of the rule)
- *Enabled*: (whether the rule is enabled or not)
- *Condition*: (the condition being evaluated)
- *Duration*: (the length of time the condition exists before triggering the failure action)
- *Interval*: (the length of time between failure actions)
- *First Failure*: (action to take upon first failure) The default is set to *Log Warning*
- *Second Failure*: (action to take upon second failure) The default is set to *Log Warning*
- *Third Failure*: (action to take upon third failure) The default is set to *Log Warning*.

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

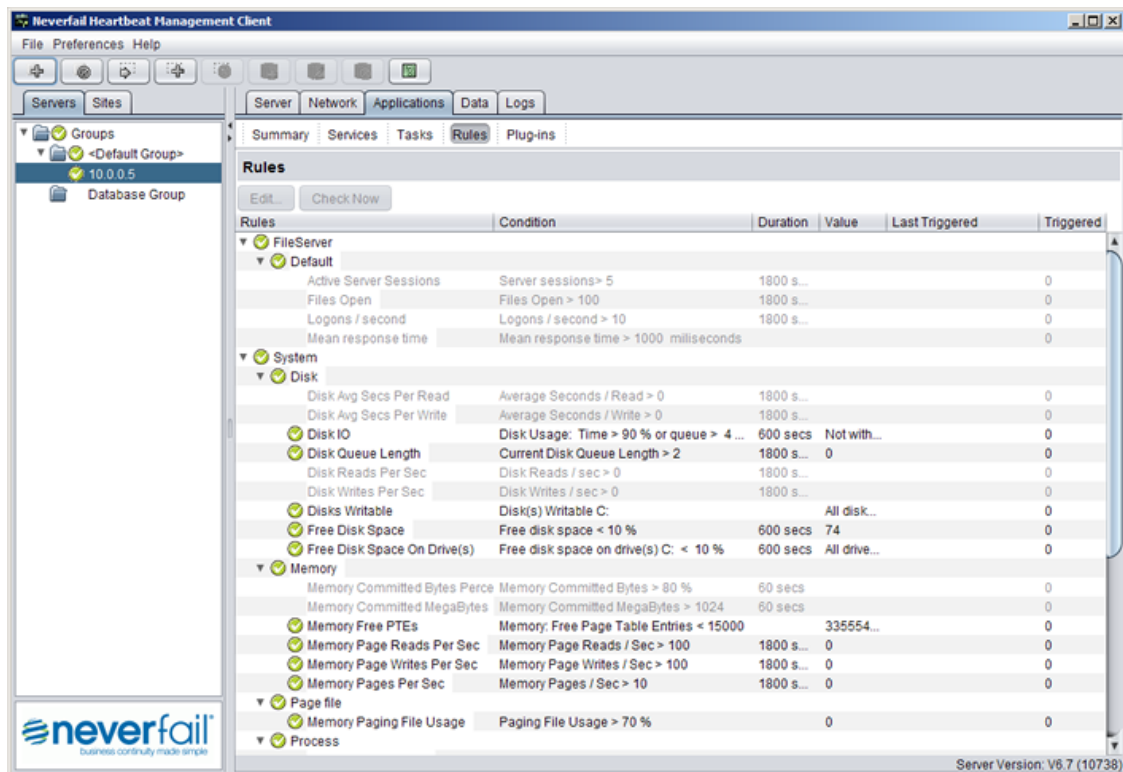


Figure 105: Applications: Rules page

Check a Rule Condition

To check a rule condition, select the rule in the *Rules* list and perform one of the following:

Procedure

To check a rule condition:

- Right-click on the rule and select **Check Now** from the menu.
- Click **Check Now** at the top of the page.
Neverfail Heartbeat immediately checks the rule conditions of the current configuration against the attributes of the system and application.

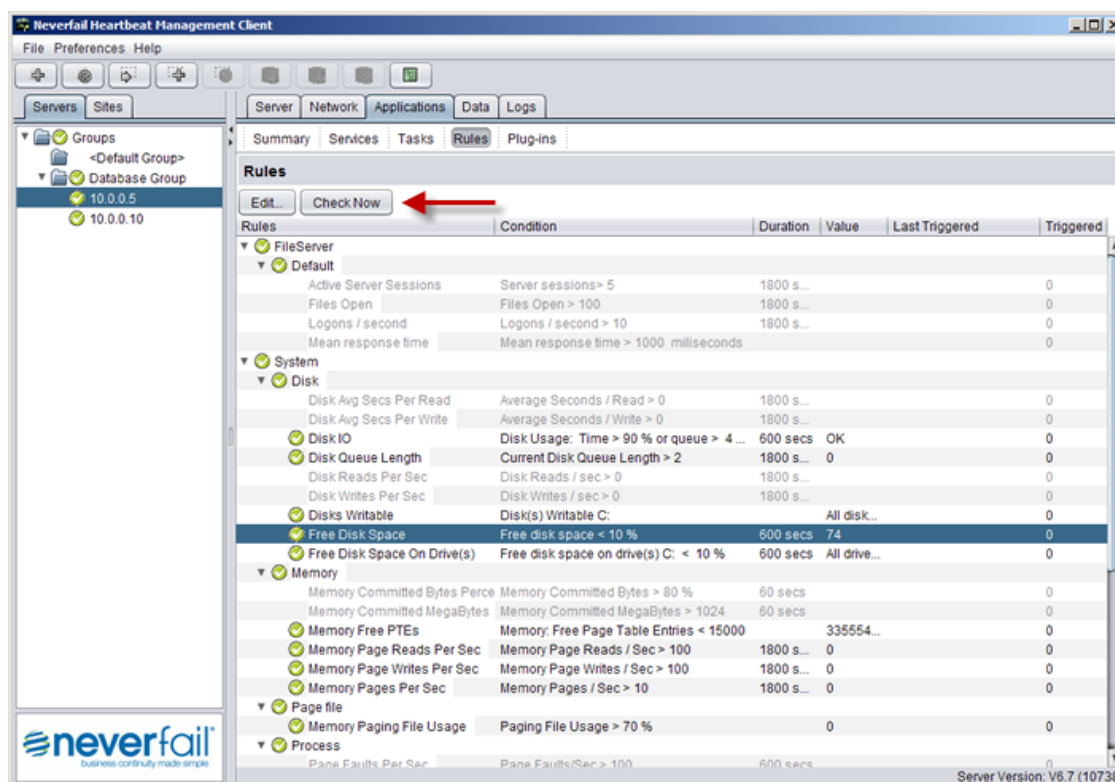


Figure 106: Applications: Rules Check Now

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.

Procedure

To Edit a rule:

1. To edit a rule, select the rule in the *Rules* list.
2. Right-click on the rule and select **Edit** from the menu or click **Edit** at the top of the page. The **Edit Rule** dialog appears. Use this dialog to *Enable* or *Disable* a Rule, set the specific options for the Rule, and to assign tasks to perform *On First Failure*, *On Second Failure*, and *On Third Failure*.



Figure 107: Edit Rule

3. When all options are selected, click **OK** to accept changes and dismiss the dialog.

Chapter 8

Data Protection

Data: Replication

Neverfail 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.

***Note:** The Neverfail Heartbeat program folder holds the send and receive queues on the active, first passive, and second passive servers, and therefore should be explicitly excluded from the set of protected files.*

You can view replication status and manage data replication through the **Data: Replication** page. Select the replication path to view in the upper right of the page, *Active to First Passive* or *Passive to Second Passive*.

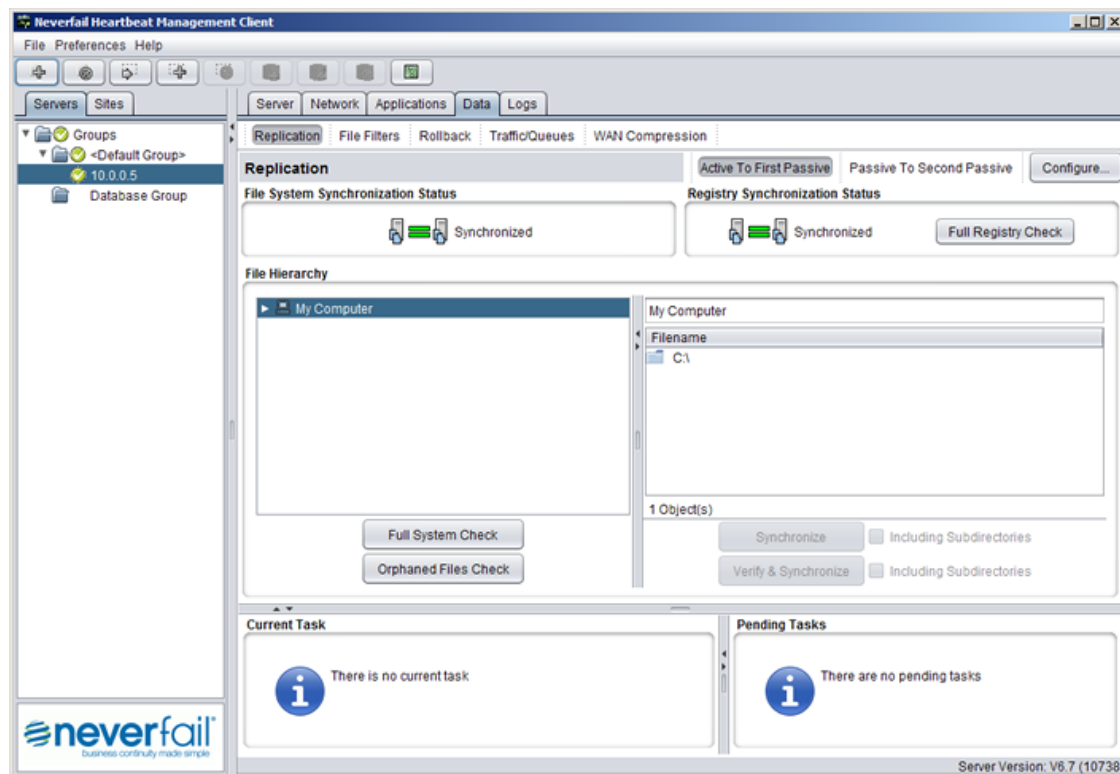


Figure 108: Data: Replication page

File System and Registry Synchronization Status

Two panes near the top of the **Replication** page in the Neverfail Heartbeat Management Client, *File System Synchronization Status* and *Registry Synchronization Status*, provide graphical status information.







Figure 109: Synchronization Status: File System and Registry

The table below lists the possible synchronization statuses that may be reported.

Table 3: Possible Synchronization Statuses: File System and Registry

Icon	Definition
	The file or folder is synchronized and verified.
	The file or folder is not synchronized. This status frequently follows a failover and indicates that manual synchronization and verification is required.
	The file or folder was not checked because a full system check was not performed or the full system check is still running and the file or folder has not yet been checked.

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.

Procedure

To initiate a Full Registry Check:

1. Select the replication path (Active to First Passive or *Passive to Second Passive*)
2. Click **Full Registry Check** in the *Registry Synchronization Status* pane.




Figure 110: Synchronization Status: File System and Registry

Initiate a Full System Check

Certain system events, such as preceding a switchover or following a failover or split-brain syndrome, may require running a full system check to ensure that the entire protected file set is synchronized and verified. A full system check performs a block-level check identical to that performed during initial synchronization and verification, and of the same files identified by the file filters.

Procedure

To initiate a full system check:

1. Click **Full System Check** in the left pane of the *File Hierarchy* pane.

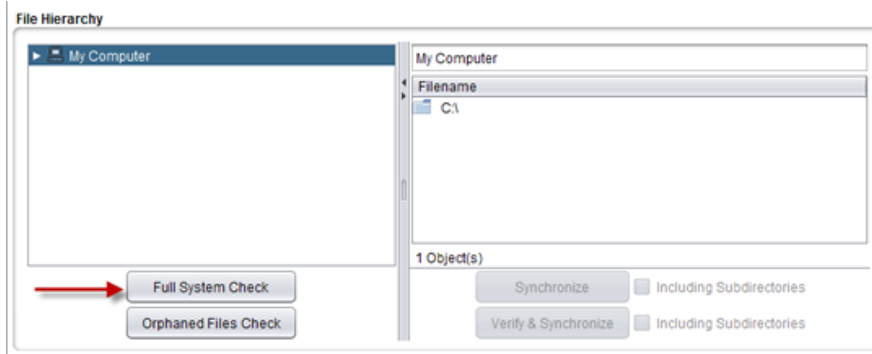


Figure 111: Data: Replication File Hierarchy pane

2. A Caution message opens and asks “Are You Sure You Want To Initiate A Full System Check?” and explains that depending on the amount of protected data, this task may take a long time to complete (a number of hours).

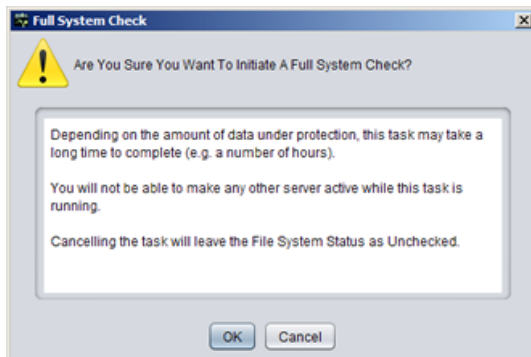


Figure 112: Full System Check Caution Message

3. Click **OK** to initiate the Full System Check, or click **Cancel** to close the message without starting the Full System Check.

Note: Once a Full System Check is initiated, allowing it to run to its conclusion is strongly recommended because canceling leaves the file system status Unchecked. Depending on the amount of data, resynchronization may take substantial time to complete. Switchover is not permitted until after the task is complete and the File System Status is Synchronized.

Fast Check

The Fast Check process is used by Neverfail Heartbeat to rapidly verify files between servers prior to starting applications. Fast Check compares file time stamps and attributes rather than the check sums of the data thereby accelerating the startup and synchronization process. If the time stamp or attribute check fails, then the normal verification and synchronization process will initiate. Additionally, you can configure the length of time to wait for Fast Check to complete before starting applications.

Fast Check is beneficial after a graceful shutdown where servers were synchronized before shutdown. Fast Check allows the server to check the file synchronization rapidly and start to service clients. If Fast Check detects files that are out-of-sync, it initiates the full verify and synchronization process to resynchronize your data.

Configure Fast Check

When combined with Controlled Shutdown, Fast Check provides the ability to perform scheduled unattended restarts of the servers.

Procedure

To enable Fast Check:

1. Navigate to **Data > Replication**.
2. Click the **Configure** button.
3. Select the *Fast Check* tab.
4. Select the manner in which Fast Check should operate using the Fast Check radio buttons.
5. Configure *Maximum Application Delay*. This is the length of time Neverfail Heartbeat will delay the startup of the application while it attempts to establish replication between active and all passive nodes.
6. Click **OK**.

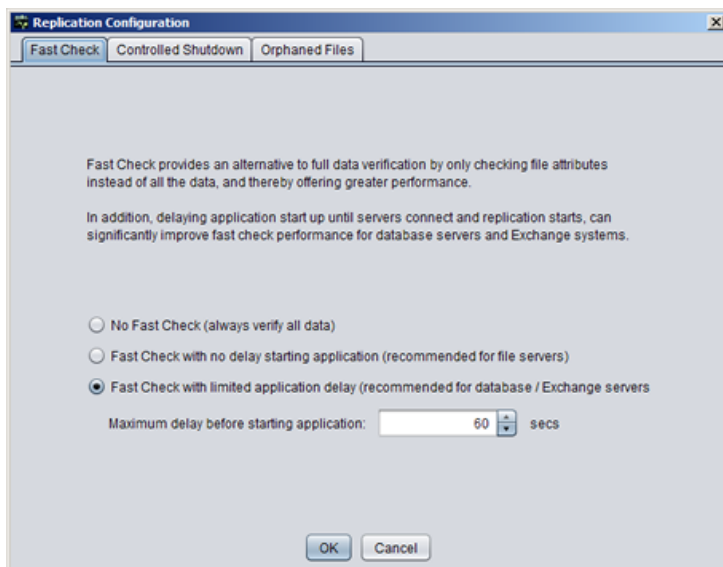


Figure 113: Configure Fast Check

Note: When Fast Check is configured in addition to Controlled Shutdown, Neverfail Heartbeat can be configured to perform an unattended restart. For more information about Controlled Shutdown, see [Controlled Shutdown](#).

Manually Initiate File Synchronization

When an out-of-sync file or folder is detected, a red icon is displayed indicating the Out-of-sync status. You can re-synchronize the out-of-sync file(s) manually using a process that is quicker and simpler than the Full System Check.

Procedure

To manually re-synchronize:

1. Select one or more files and folders from the list in the right pane of the *File Hierarchy* pane. Multiple files and folders can be selected from this file list by using the standard Windows multiple selection techniques, **Shift** + click and **Ctrl** + click.
2. When one or more folders are selected, also select the *Including Subdirectories* check box to ensure that all files within the folder(s) are also synchronized.
3. Click **Synchronize**. As the synchronization runs, you may see its progress in the *Current Task* pane at the bottom left of the **Data: Replication** page. When the synchronization process successfully completes, a green icon indicates synchronized status.

You also can right-click on a folder in the tree view (in the left pane of the *File Hierarchy* pane) to quickly select *Synchronize* or *Verify and Synchronize* from a menu. Both options automatically include subdirectories.

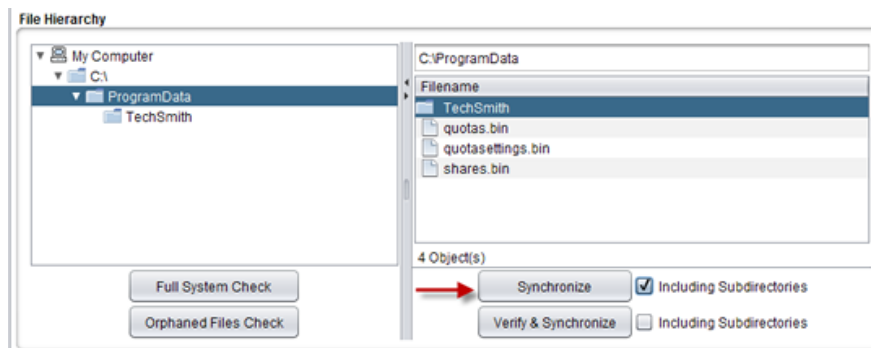


Figure 114: Manual Selection

Manually Initiate Verify and Synchronize

To perform manual verification and synchronization, the process is identical to the one described in *Manually Initiate File Synchronization* except that the process is started by clicking **Verify and Synchronize**.

Procedure

To manually verify and synchronize:

1. Select one or more files and folders from the list in the right pane of the *File Hierarchy* pane. Multiple files and folders can be selected from this file list by using the standard Windows multiple selection techniques, **Shift** + click and **Ctrl** + click.
2. When one or more folders are selected, also select the *Including Subdirectories* check box to ensure that all files within the folder(s) are also verified and synchronized.
3. Click **Verify and Synchronize**. As verify and synchronization runs, you may see its progress in the *Current Task* pane at the bottom left of the **Data: Replication** page. When the verify and synchronization process successfully completes, a green icon indicates verified and synchronized status.

You also can right-click on a folder in the tree view (in the left pane of the *File Hierarchy* pane) to quickly select *Verify and Synchronize* from a menu. This option automatically includes subdirectories.

Each verification and synchronization request (manually or automatically scheduled) is defined as a task with subsequent tasks queued for processing after the current task is completed. Each task is listed in the *Pending Tasks* list to the right of the *Current Tasks* frame.

Note: Individual tasks can be canceled, but canceling automatically triggered tasks can lead to an *Unchecked* system. A warning is presented detailing the possible consequences of canceling tasks.

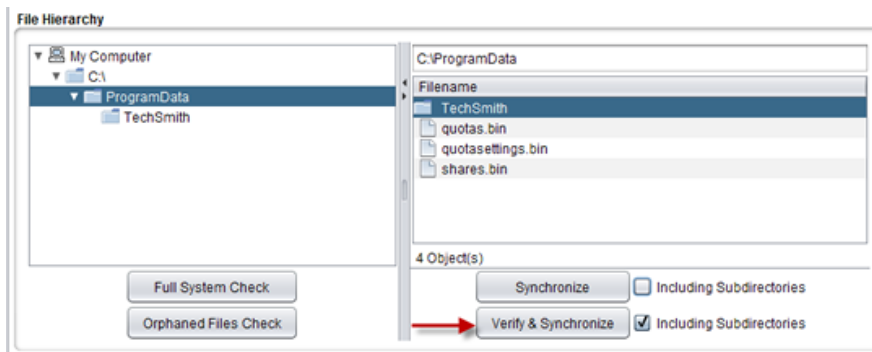


Figure 115: Manual Selection

Orphaned Files Management

Neverfail Heartbeat provides the opportunity to check the system for orphaned files and either notify the administrator or to delete the orphaned files. Orphaned files are those files in a protected set that exist on the passive server but do not exist in the protected set on the active server in a pair. In a trio configuration, where replication is from the active server to the first passive server to the secondary passive server, files in the protected set on the 2nd passive server that are not on the first passive server are considered orphaned.

Orphaned File Check can either delete or log files on the passive server that exist within the protected set; they were “orphaned” because Neverfail Heartbeat was not running when content changes were made on the active server.

Note: Orphaned File Check does not delete files on the passive server if there is no file filter to include the content as this would be unsafe.

Special Cases

Filters for files, file types, or other wildcards

Folder root filters

Orphaned File Check will manage the entire contents of that folder (for example, D:\folder**). This deletes all passive files within the folder that do not exist on the active server, and includes content created only on the passive server.

Exclusion file filters

Orphaned File Check will not delete any files excluded from the protected set by exclusion filters. This rule safeguards users and applications.

Filters for files, file types, or other wildcards

Orphaned File Check is not managing the contents of the folder (for example, D:\database*.log), only the selected files.

Orphaned File Check will only process files that match the filter and will not delete files with any other extension within the folder D:\database

Orphaned files are those files in a protected set that exist on the passive server but do not exist in the protected set on the active server in a pair.

Prior to initiating an orphaned files check, you must configure the options for actions to take in the event orphaned files are found. By default, Orphaned Files Check is configured to delete orphaned files. Should you want to log the files presence, see [Configure Orphaned Files Check](#).

Configure Orphaned Files Check

Prior to initiating an orphaned files check, you must configure the options for actions to take in the event orphaned files are found. By default, Orphaned Files Check is configured to delete orphaned files. Should you want to log the files presence, follow the steps below.

Procedure

To Configure Orphaned Files Check options:

1. Navigate to the **Data: Replication** page and click on the **Configure** button.
2. Select the *Orphaned Files* tab.
3. Select the *Detect orphaned files* check box and in the *On detection, take the following action* drop-down to automatically *Delete* the orphaned files or *Log to file* to add the files list to the log file.

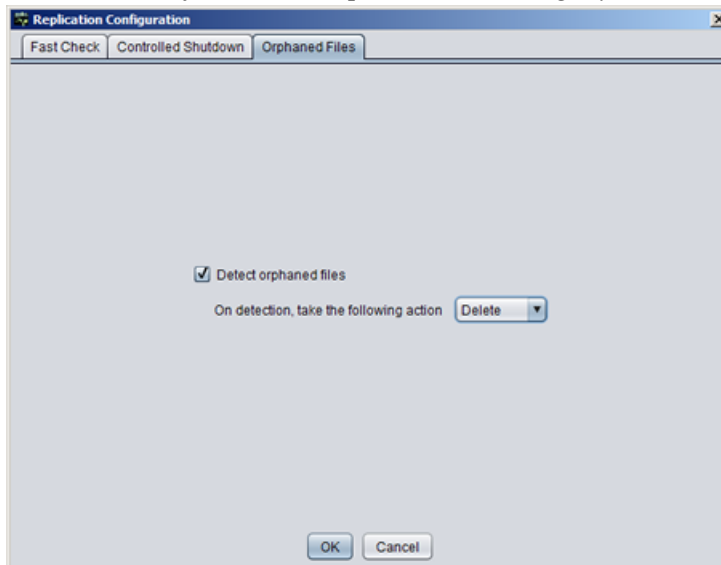


Figure 116: Orphaned Files Configuration Options

4. After selecting the options, click **OK** to close the dialog.
5. Click the **Orphaned Files Check** button.

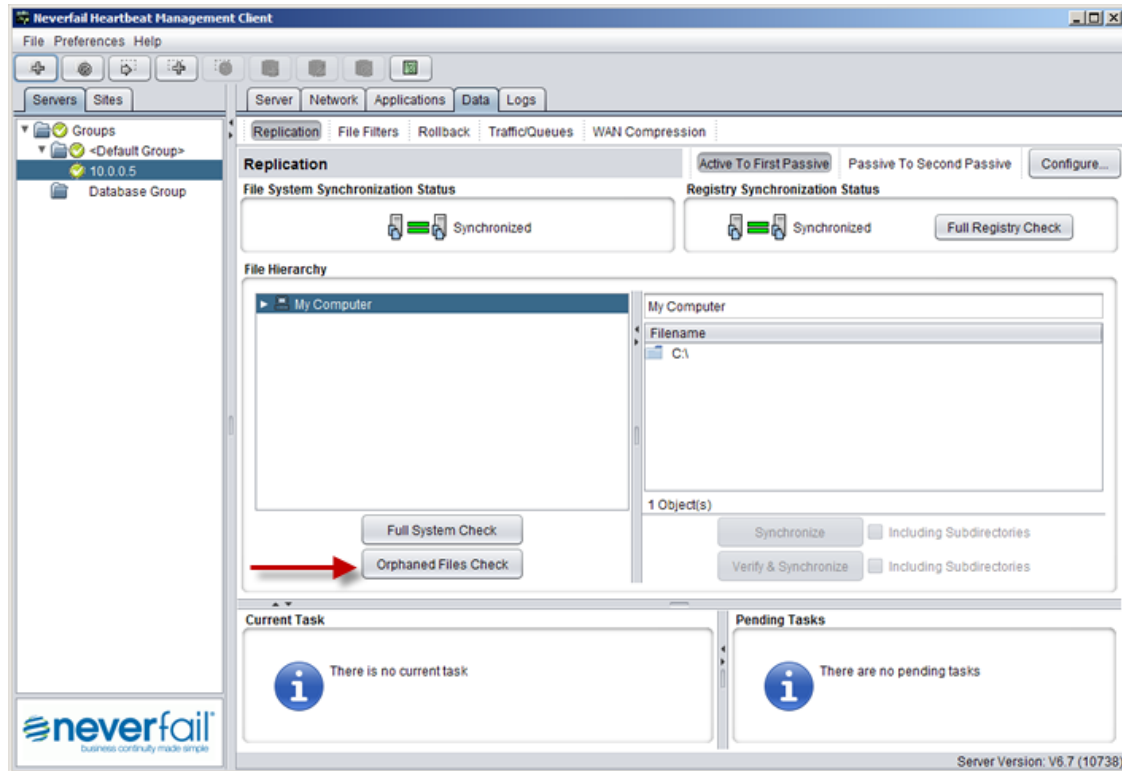


Figure 117: Initiate Orphaned Files Check

Data: File Filters

File filters dictate which files are protected, the disk I/O operations to intercept and replicate on the passive server(s) and allows you to customize the inclusion and exclusion of files from the replication process.

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

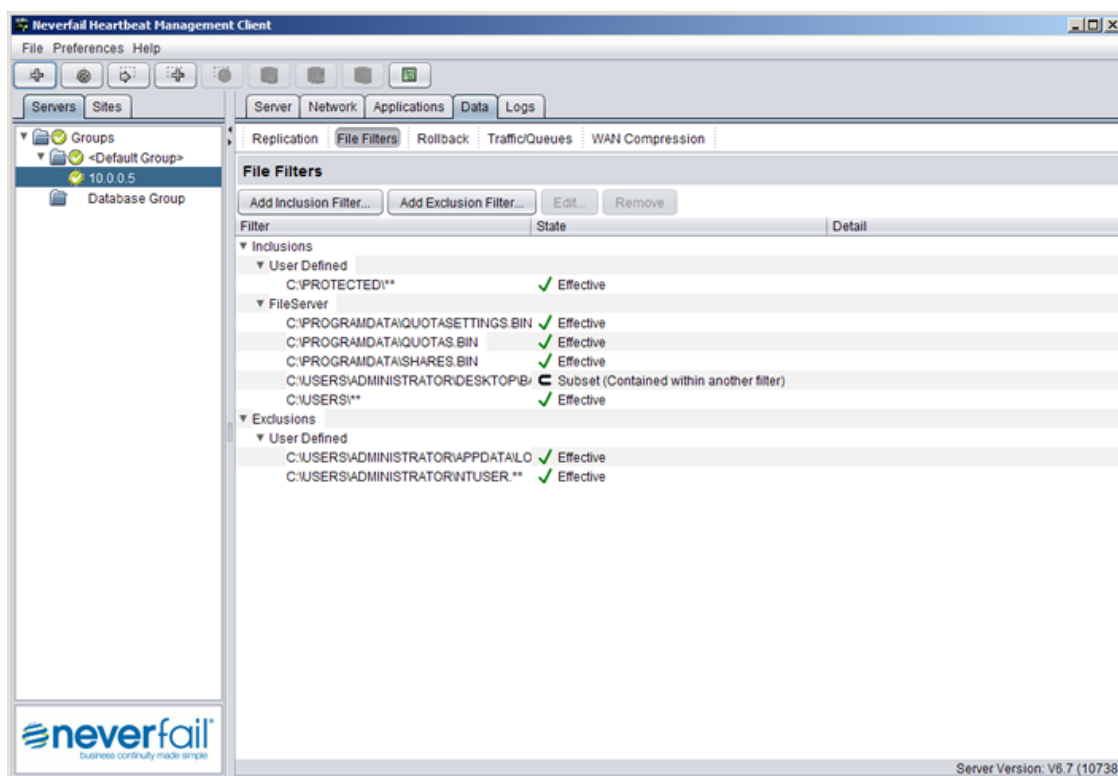


Figure 118: Data: File Filters

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.

Determine Effective Filters

An *Effective* filter is what remains of the files and folders stipulated in the inclusion filter after removing the files/folders in the Exclusion filter.

Filters are compared with each other, and if one filter is a superset of another, only the superset filter is used. This allows, for example, file servers with thousands of individual shares requested by a plug-in to use a single, more general, user configured filter. All filters are validated before they are passed to the replication components, and the *Configuration* pane of the **Data** page reports why filters are rejected (for example, because the drive does not exist).

Add a User Defined Inclusion Filter

Administrators can add filters to include additional files or folders in the protected set.

Procedure

To add a user defined Inclusion Filter, perform the following steps:

1. Click **Add Inclusion Filter** to open the **Add Inclusion Filter** dialog.




Figure 119: Add Inclusion Filter

2. Filters to protect user defined files and folders are defined by typing the complete path and pattern, by specifying a pattern containing wildcards, or browsing to the file or folder by clicking **Browse**.
3. Click **OK** to accept the changes, or **Cancel** to dismiss the dialog without making any changes.

The two forms of wildcard available are `*`, which matches all files in the current folder or `**`, which matches all files, subfolders and the files in the subfolders of the current folder. After the filter is defined, subsequent inclusion filters may be added.

Note: Neverfail Heartbeat “vetoes” replication of a few specific files and folders such as the Neverfail Heartbeat installation directory or the `System32` folder. If you create an inclusion filter that includes any of these off-limits files or folders, the entire filter is vetoed, even if you have created an exclusion filter to prevent replication of those files or folders.

Add a User Defined Exclusion Filter

Exclusion Filters are configured to create a subset of an Inclusion Filter to exclude from protection. The Exclusion Filter is created in the same way as the Inclusion Filter.

Procedure

To Add a user defined filter:

1. Filters to exclude files and folders from protection and replication are defined by clicking **Add Exclusion Filter** on the **Data: File Filters** page of the Neverfail Heartbeat Management Client.



Figure 120: Add Exclusion Filter

2. Type the complete path and pattern or specify a pattern containing wildcards, or browse to the file or folder by clicking **Browse**.
3. Click **OK** to accept the changes.
The two forms of wildcard available are `*`, which matches all files in the current folder, and `**`, which matches all files, subfolders and the files in the subfolders of the current folder.

Edit User Defined Inclusion/Exclusion Filters

User defined Inclusion/Exclusion filters can be edited using the Neverfail Heartbeat Management Client.

Procedure

To Edit a user defined Inclusion/Exclusion Filter:

1. Select the filter and click **Edit** at the top of the **File Filters** page or right-clicking the filter and selecting **Edit** from the menu.




Figure 121: Edit Inclusion Filter

2. Edit the value in the *Pattern:* text box by typing over the current file filter definition.
3. Click **OK**.
The file filter is changed and becomes active.

Remove User Defined Filters

Procedure

To Remove a user defined filter:

1. To remove an Inclusion filter or Exclusion filter, 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.
A confirmation message appears:
2. Click **Yes** to remove the filter, or click **No** to dismiss the message without deleting anything.

Data: Rollback

The Neverfail Heartbeat Data Rollback Module (DRM) provides a way to rollback data to an earlier point in time. This helps mitigate problems associated with corrupt data such as can result from virus attacks. Before configuring or using any of the DRM features accessed through this page, Neverfail recommends that you read and follow the steps described in the section immediately below, *Best Practices for Using Volume Shadow Copy Service & DRM*.

Best Practices for Using Volume Shadow Copy Service & DRM

The Volume Shadow Copy Service (VSS) component of Windows 2003 and Windows 2008 takes shadow copies and allows you to configure the location and upper limit of shadow copy storage.

1. To configure VSS, right-click on a volume in Windows Explorer, select *Properties*, and then select the *Shadow Copies* tab.

Note: VSS is also used by the *Shadow Copies of Shared Folders (SCSF)* feature of Windows 2003 and Windows 2008, and consequently, some of the following recommendations are based on Microsoft™ *Best Practices for SCSF*.

2. Decide which volume to use for storing Shadow Copies before using DRM because you must delete any existing shadow copies before you can change the storage volume.
Neverfail recommends that a separate volume be allocated for storing shadow copies. Do not use a volume to store both Neverfail Heartbeat protected data and unprotected, regularly updated data. For example: do not write backups of data (even temporarily) to a volume that contains Neverfail Heartbeat protected files, as that increases the space required for snapshots.

In accordance with the following guidelines from Microsoft:

Select a separate volume on another disk as the storage area for shadow copies. Select a storage area on a volume that is not shadow copied. Using a separate volume on another disk provides two advantages. First, it eliminates the possibility that high I/O load causes deletion of shadow copies. Second, this configuration provides better performance.

3. Be sure to allocate enough space for the retained shadow copies.
This is dependent on the typical load for your application, such as the number and size of emails received per day, or the number and size of transactions per day. The default is only 10% of the shadowed volume size and should be increased. Ideally, you should dedicate an entire volume on a separate disk to shadow storage.

Note: The schedule referred to in the **Volume Properties > Shadow Copies > Settings** dialog is for *Shadow Copies for Shared Folders*. This is not used for DRM - the DRM schedule is configured in the *Rollback Configuration* pane of the *Neverfail Heartbeat Management Client*.

4. Configure the schedule to match your clients' working patterns. Considering both the required granularity of data restoration, and the available storage.

DRM provides a means of flexibly scheduling the creation of new Shadow Copies, and the deletion of older Shadow Copies. Adjust this to suit the working-patterns of your clients and applications. For example, do clients tend to work 9am-5pm, Monday-Friday in a single time zone, or throughout the day across multiple time zones? Avoid taking Shadow Copies during an application's maintenance period, such as Exchange defragmentation, or a nightly backup.

In selecting how frequently to create new shadow copies, and how to prune older ones, you must balance the advantages of fine-granularity of restorable points-in-time versus the available disk space and the upper limit of 512 Shadow Copies across all shadowed volumes on the server.

5. Perform a trial-rollback.

After DRM is configured, Neverfail recommends that you perform a trial-rollback, to ensure that you understand how the process works, and that it works correctly.

If you do not select the option *Restart applications and replication*, then you can rollback to Shadow Copies on the passive server without losing the most recent data on the active server.

6. Start the application manually to verify that it can start successfully using the restored data.
Note the following:

- The application is stopped on the active during the period of the test.
- Following the restoration of data on the passive, it becomes active and visible to clients on the network.

After the test is complete, shut down Neverfail Heartbeat on both servers. Use the **Server Configuration Wizard** to swap the active and passive roles, and then restart. This re-synchronizes the application data from the active to the passive, and allows you to restart using the application data as it was immediately before the rollback.

7. Monitor Neverfail Heartbeat to identify any Shadow Copies that are discarded by VSS.

If DRM detects the deletion of any expected Shadow Copies, this is noted in the *Neverfail Heartbeat Event Log*.

This is an indication that VSS reached its limit of available space or number of Shadow Copies. If many Shadow Copies are automatically discarded, consider adding more storage, or reconfiguring your schedule to create and maintain fewer shadow copies.

Configure VSS on Windows 2003 Server

Procedure

To configure the Volume Shadow Copy Service on Windows 2003, use the following procedure:

1. Double-click **My Computer**.
2. Right-click a hard disk icon, and select *Properties*.
3. Select the *Shadow Copies* tab.
4. Select the volume on which to configure storage, and click **Settings**.
5. Select the volume to contain the Shadow Copy storage, and configure the size limit (if any) for the storage area.

Configure VSS on Windows 2008 Server

Procedure

To configure Volume Shadow Copy Service on a Windows 2008 Server:

1. Navigate to **Start > All Programs > Administrative Tools > Computer Management** to launch the **Computer Management** configuration tool.
2. Select **Storage > Disk Management**.
3. Right-click on a Volume and select *Properties*.
4. Select the *Shadow Copy* tab.
5. Select the Volume for which Shadow Copy is to be enabled.
6. Click **Settings**.
7. In the *Located on this volume:* field select the Volume where the shadow copies are to be stored.
8. Configure the maximum size to be made available for the shadow copies.
9. Click **OK**.

View Rollback Status and Work with Shadow Copies

The **Data: Rollback** page in the Neverfail Heartbeat Management Client provides both status of shadow copies and features to create and delete shadows manually, to rollback to a shadow, and to configure the taking of shadow copies.

T

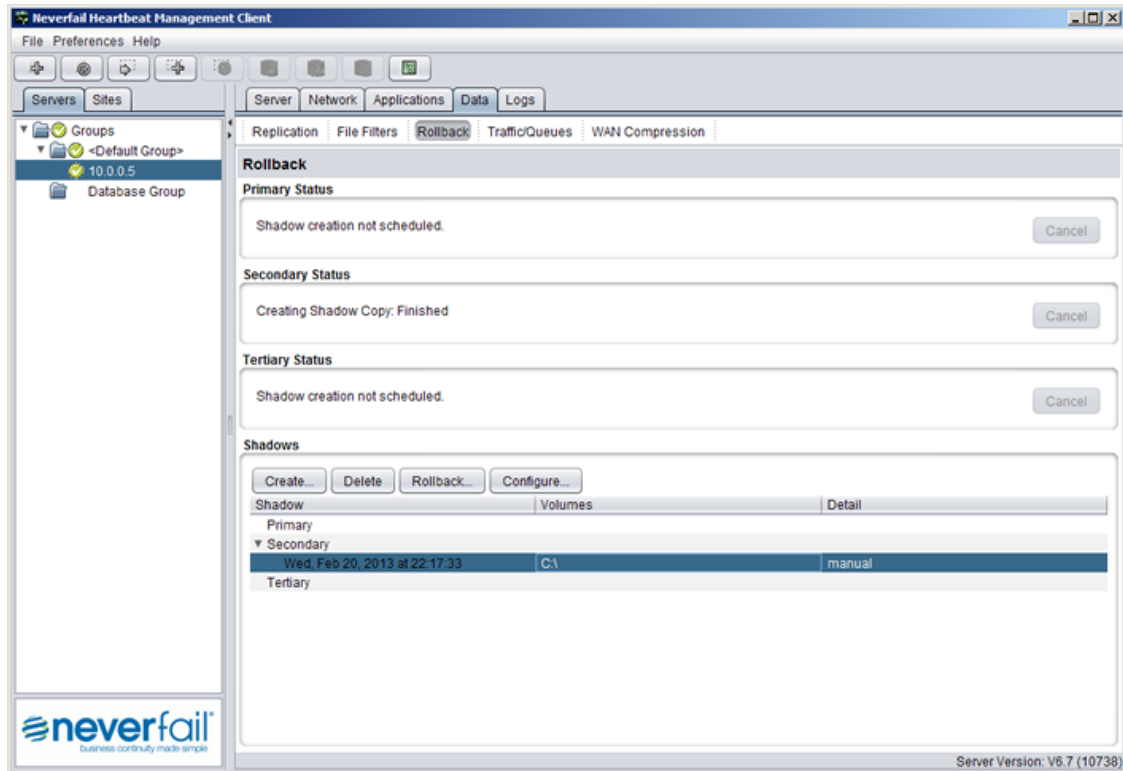


Figure 122: Data: Rollback

View Rollback Status

The Status panes (Primary Status, Secondary Status, and Tertiary Status) of the **Rollback** page provide a way to view the current status before taking any other actions.

If Shadow creation is ongoing, wait until finished before taking any action or click **Cancel** to cancel the current operation. It may not always be possible to cancel the current operation, in which case a warning appears in the Neverfail Heartbeat event log.

Configure Neverfail Heartbeat to Automatically Create Shadow Copies

DRM provides a means of flexibly scheduling the creation of new shadow copies, and the deletion of older shadow copies. Adjust this to the working patterns of your clients and applications. For example, do clients tend to work 9am- 5pm, Monday-Friday in a single time zone, or throughout the day across multiple time zones? Avoid taking shadow copies during application's maintenance periods, such as Exchange defragmentation, or a nightly backup. In selecting how frequently to create new shadow copies, and how to prune older ones, you must balance the advantages of fine-granularity of restorable points-in-time versus the available disk space and the upper limit of 512 shadow copies across all shadowed volumes on the server.

Most shadows are taken on the passive servers, but you may select to create one shadow per day on the active server. This is with the co-operation of the application, for VSS-aware applications, such as Exchange and SQL Server. As this may cause a temporary suspension of application activity, it is preferable to schedule this for periods of light load.

***Note:** When the DRM is configured to create and manage shadow copies automatically, this task runs at the selected frequency (for example, every 15 minutes) and this is reflected in the status display, which indicates the time of the next scheduled task. The task does not necessarily create a shadow every time it runs, as this is determined by the schedule configuration.*

The timestamp associated with a Shadow Copy refers to the system time of the server on which the shadow copy is taken.

Configure the Shadow Copy Schedule

DRM can create and delete shadow copies automatically according to a configurable schedule. The aim of the schedule is to provide a balance between providing a fine-granularity of rollback points-in-time on the one hand, and conserving disk space and number of shadow copies on the other. To achieve this balance, the available configuration options reflect the observation that recent events generally are of more interest and value than older ones. For example, the default schedule maintains one shadow from every day of the last week, and one shadow from every week of the last month.

Procedure

Neverfail Heartbeat can be configured to automatically create shadow copies by performing the following steps:

1. Navigate to the **Data: Rollback** page and click **Configure**. The *Configure Shadow Schedule* dialog appears.

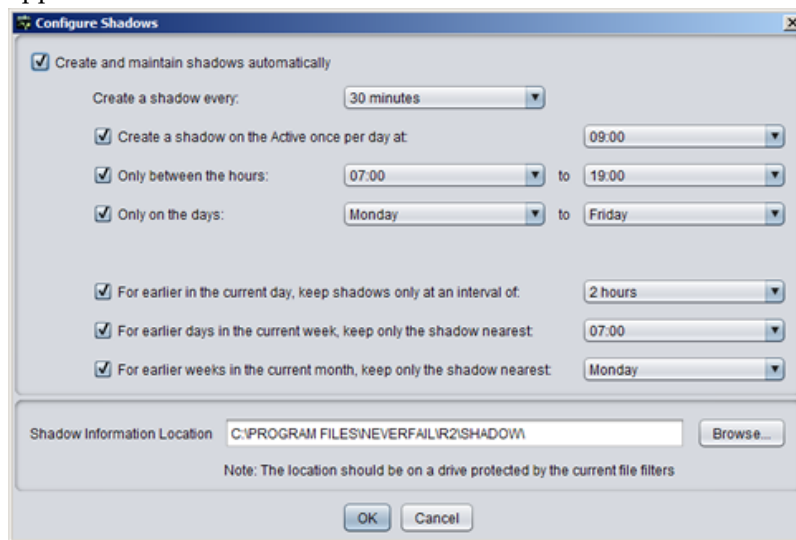


Figure 123: Configure Shadow Schedule

2. Select the *Create and maintain shadows automatically* check box.
The *Create and maintain shadows automatically* check box controls the automatic creation and deletion of Shadow copies. When selected, automatic Shadow copies are created and deleted in accordance with other user configuration settings. When cleared, you can still manually create, delete, and rollback shadow copies from the *Rollback* pane.

Note: *Configure the schedule to suit your clients' working patterns; the required granularity of data restoration, and the available storage.*

3. Select the frequency and time periods for creating shadows. (See *Configure Shadow Creation Options*, below.)
4. Select the shadows to keep or remove from earlier time periods. (See *Configure Shadow Keep Options*.)

Note: *The Volume Shadow Copy Service (VSS) component of Windows 2003/2008, may automatically delete old shadows because of lack of disk space even when the Create and maintain shadows automatically check box is not selected.*

Configure Shadow Creation Options

These options set the frequency for shadow creation on the passive and active servers respectively.

Procedure

Note: No shadows are created when the system status is *Out-of-sync* or *Not Replicating*.

- **Create a shadow every:**

This drop-down list controls how frequently a shadow copy is taken on the passive servers, the default setting is every 30 minutes. When the shadow is actually taken is also controlled by *Only between the hours:* and *Only on the days:*, if either of these are set then shadows are taken at the frequency defined by this drop down list but only within the days/hours defined by them.

- **Create a shadow on the Active once per day at:**

If the check box is cleared, then no shadows are automatically created on the active. If it is selected, then a Shadow is taken each day at the time selected from the drop down list. The Shadow is taken with "application co-operation", which means that if the application protected by Neverfail Heartbeat is integrated with VSS, it is informed before the shadow is taken and given the opportunity to perform whatever tidying up it is designed to do when a VSS Shadow is taken.

Note: It is possible to select a time outside of the *Only between the hours:* range. This prevents creation of the shadow.

Whether a shadow is actually taken is also controlled by *Only between the hours:* and *Only on the days:*, if either of these are configured, then a shadow is taken only within the days/hours defined by them. The following two options limit the number of shadows taken during periods when the data is not changing.

- **Only between the hours:**

If this check box is selected, then the range defined by the two drop down lists are applied to the automatic creation of shadows on either on the passive server(s) (as controlled by *Create a shadow every:*), or on the active server (as controlled by *Create a shadow on the Active once per day at:*).

For example, to limit shadow captures to night time hours, you can define a range of 20:00 to 06:00.

- **Only on the days:**

When the check box is selected, the range defined by the two drop down lists is applied to the automatic creation of shadows either on the passive server(s) (as controlled by *Create a shadow every:*) or active server (as controlled by *Create a shadow on the Active once per day at:*).

For example, to limit shadow captures to weekend days, you can define a range of Saturday to Sunday.

Note: The shadow copy information location is configurable. The default location ensures that the information location includes a copy of the necessary file filters to be used in a rollback. Neverfail recommends that the default setting be used for shadow copy information location.

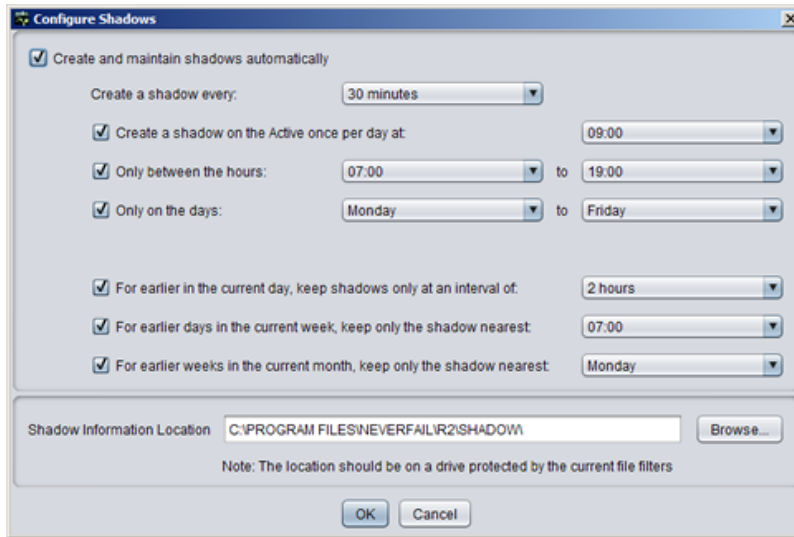


Figure 124: Shadow Creation Options

Configure Shadow Keep Options

The purpose of the following three options is to reduce the number of older shadows while preserving a series, which spans the previous 35 days.

Procedure

Manually created shadows are not deleted automatically, but VSS deletes old shadows (whether manually created or not) whenever it requires additional disk space for the creation of a new shadow. When manually created shadows match the criteria for keeping a shadow from a particular time period, automatic shadows in close proximity are deleted. For example, a manually created shadow is not deleted, but can be used for the “keep algorithm”.

- **For earlier in the current day, keep shadows only at an interval of:**

If the check box is selected, then only the first shadow is kept for each interval as defined by the value (hours) selected from the drop-down list. Earlier in the current day means since Midnight and older than an hour. The intervals are calculated from either at Midnight or if *Only between the hours:* is selected, then from the start hour. For shadows taken before the start time (as the start time may change), the interval is calculated backwards again starting at the start time.

- **For earlier days in the current week, keep only the shadow nearest:**

If the check box is selected, then only the shadow nearest to the time (24 hour clock) selected from the drop-down list is kept for each day. Earlier days in the current week means the previous seven days not including today (as today is covered by the above option). A day is defined as Midnight to Midnight.

If a shadow was taken at 5 minutes to midnight on the previous day it is not considered when calculating the nearest.

- **For earlier weeks in the current month, keep only the shadows nearest:**

If the check box is selected, then only the shadow nearest to the selected day is kept for each week. Earlier weeks in the current month means the previous four weeks not including either today or the previous 7 days (as they are covered by the above two options).

To calculate the “nearest”, an hour is required. The calculation attempts to use the selected time from *For earlier days in the current week, keep only the shadow nearest:* if it is selected, otherwise the

Only between the hours start time is used if it is selected, finally, when neither of these options are configured, Midnight is used.

All automatic shadows taken more than 35 days ago are deleted. The intervening 35 days are covered by the above three options.

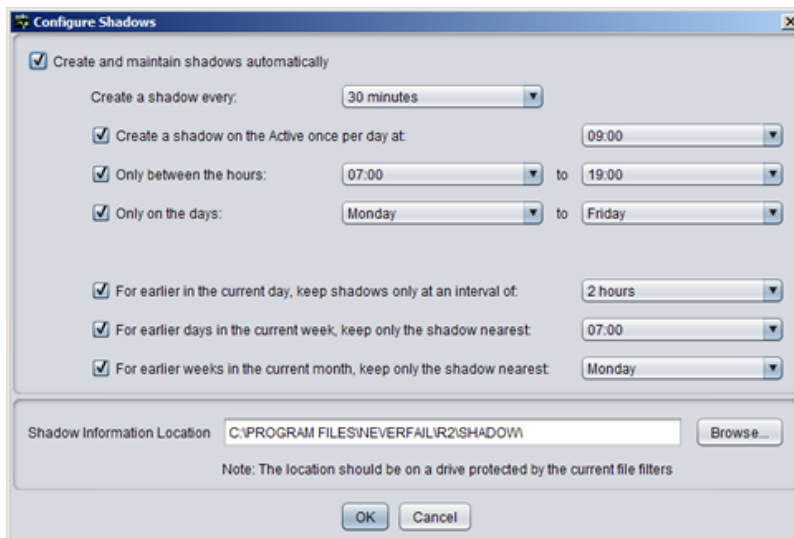


Figure 125: Shadow Keep Options

Roll Back Protected Data to a Previous Shadow Copy

Should the need arise to roll data back to a previous point in time, perform the following:

Procedure

1. Go to the *Shadows* pane of the **Rollback** page and select an existing Shadow from the Primary, Secondary, or the Tertiary server list.
2. A dialog is presented allowing you to create a shadow immediately before the rollback, and select whether to restart applications and replication after the rollback.

Note: Electing to create a shadow before the rollback means that if you change your mind, you can restore to the most recent data.

Choosing to restart applications and replication simplifies the restore procedure, but eliminates the chance to examine the data before it is replicated to the other server.

3. Click **Rollback** in the *Shadows* pane.
A confirmation dialog is presented.
4. Click **Yes**.

Neverfail Heartbeat stops the applications and replication, and then restores protected files and the registry from the Shadow Copy. Neverfail Heartbeat then sets the file and registry filters to those persisted in the Shadow Copy. If the Shadow Copy is on a currently passive server, then this server will become active after the rollback.

If the rollback fails, the reason for the failure is shown in the status display. This may be because a particular file set of files or registry key cannot be accessed. For example, a file may be locked because the application is inadvertently running on the server performing the rollback, or

permissions may prevent the SYSTEM account from updating. Rectify the problem and try performing the rollback again.

5. If selected, applications and replication are restarted and the Cluster re-synchronizes with the restored data.
 - If you selected not to restart applications and replication automatically, you can now start the application manually. This allows you to check the restored data.
 - If you decide to continue using the restored data, click **Start** on the *Neverfail Heartbeat System Overview* pane to re-synchronize using this data.
 - If you decide you want to revert to the pre-rollback data, which is still on the other (now passive) server, you can shut down Neverfail Heartbeat, use the **Configure Server Wizard** to swap the active and passive roles, and then restart. This re-synchronizes the servers with the pre-rollback data.

As a result of the rollback, the file and registry filters are set to the configuration, which was in use when the shadow copy was taken.

Manually Create Shadow Copies

Shadow Copies can be created manually using the steps below:

Procedure

1. In the **Data: Rollback** page *Shadows* pane, click **Create**.
The **Create Shadow** dialog opens.




Figure 126: Create Shadow Dialog

2. Select *Primary*, *Secondary*, or *Tertiary* as the server location where the shadow will be taken.
3. Click **OK** to create the manual shadow copy, or click **Cancel** to close the dialog without creating a shadow copy.

Delete a Shadow Copy

Procedure

Should the need arise to delete shadow copies, follow the procedure below:

- To delete a shadow copy, select it in the *Shadows* pane of the **Rollback** page. Click **Delete**.
The selected shadow copy is deleted.

Data: Traffic/Queues

The **Data: Traffic/Queues** page presents information on the communications status of Neverfail Heartbeat. Select the graphic of the Neverfail Channel to view the status of a specific channel. To view the status of the send and receive queues on a specific server, select the server icon.

The queue delays for each server (indicated by the value of the *Age of Oldest Entry*) represent the amount of time a specific piece of information remains in the send queue and the receive queue respectively before it is moved to the passive server receive queue or committed to disk.

Configure Send and Receive Queue Size

The send and receive queues are pre-configured to utilize a maximum of 10 GB of disk space. Queued data may require more space than the default setting provides when slow connectivity is provided for the Neverfail Channel. You can adjust the maximum disk space for queued data on the **Data: Traffic/Queues** page.

Procedure

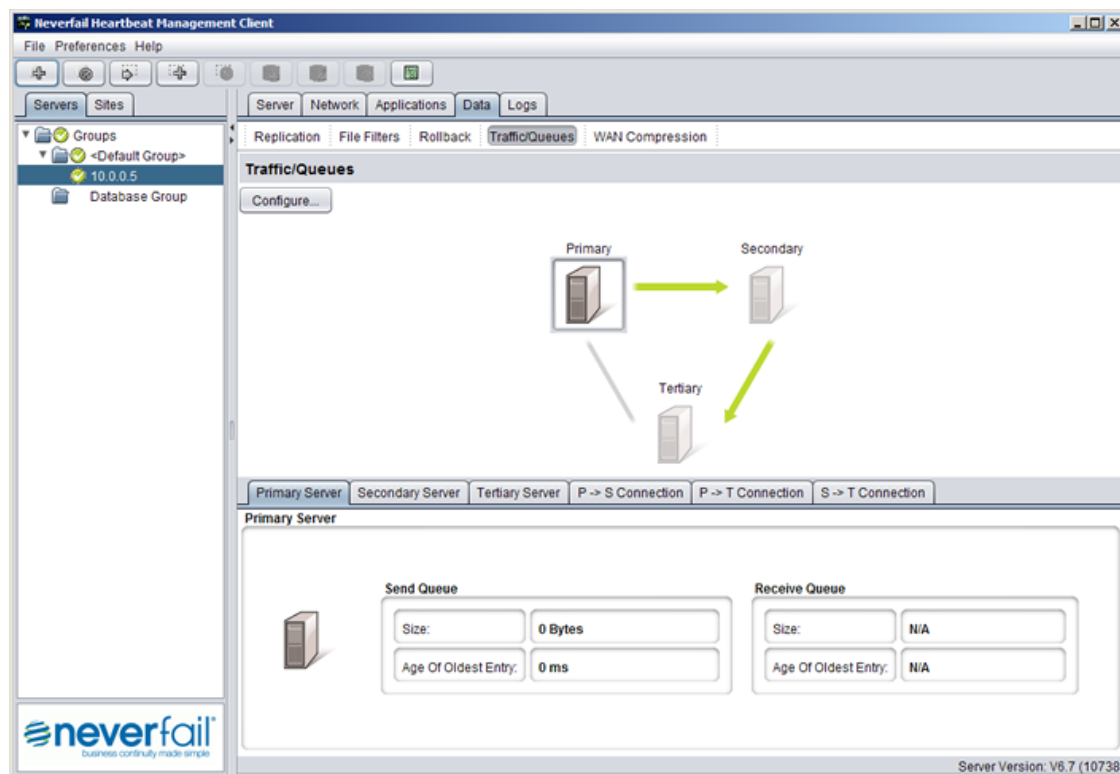


Figure 127: Data: Traffic/Queues

1. Click **Configure** to open the **Traffic/Queues Configuration** dialog.



Figure 128: Traffic/Queues Configuration

2. Type a numeric value (GB) into the text box or use the up and down arrow buttons to adjust the value.
3. Click **OK** to apply the new setting and dismiss the dialog, or click **Cancel** to close the dialog without making any changes.

Review the Neverfail Channel Status

In the **Data: Traffic/Queues** page, the Neverfail Channel status is displayed:

- Connected - A green solid arrow icon
- Waiting - An orange solid icon is displayed when the channel has just disconnected. Neverfail Heartbeat will wait for a configured amount of time before deciding the channel is disconnected.
- Not connected - A red broken line 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 are all displayed. The *Channel Connection* tab in the lower pane displays the IP addresses used by the Neverfail Channel for the Primary to Secondary, Secondary to Tertiary, and Tertiary to Primary (if Tertiary is deployed) connections and the port that the communications are using.

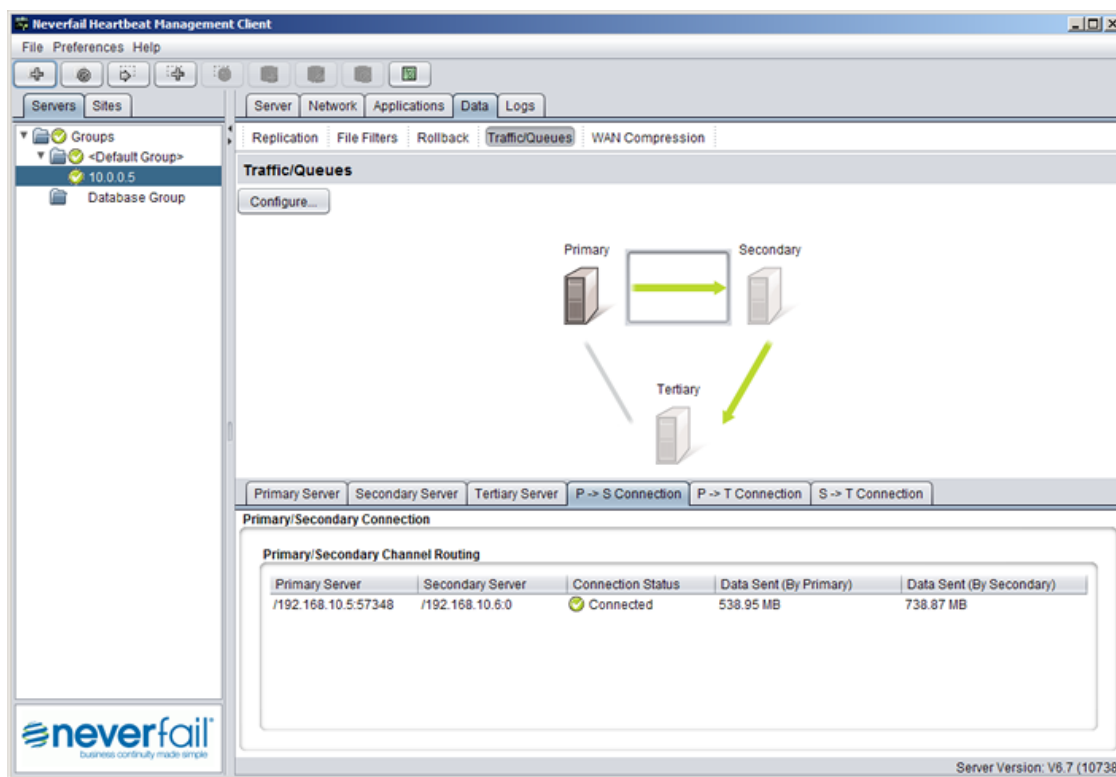


Figure 129: Neverfail Channel Status

Data: WAN Compression

Neverfail Heartbeat offers WAN Compression as an optional feature to assist in transferring data fast over a WAN. When included in your Neverfail Heartbeat license, WAN Compression can be configured through the **Data: WAN Compression** page. The **Data: WAN Compression** page provides a quickly accessible status on the current state of WAN operations.

The WAN Compression feature allows the administrator to select from the following options:

Note: Allow Neverfail Heartbeat to select the compression type automatically - is the recommended setting.

- *Allow Neverfail Heartbeat to select the compression type automatically* — Neverfail Heartbeat selects the level of WAN compression based upon current configuration without user intervention.
- *Advanced* — Neverfail Heartbeat uses the WAN Deduplication feature in addition to compression to remove redundant data before transmitting across the WAN thereby increasing critical data throughput.
- *Standard* — Neverfail Heartbeat uses compression on data before it is sent across the WAN to improve WAN data throughput speed.
- *None* — Selected when deployed in a LAN or where WAN Compression is not required.

When Neverfail Heartbeat is deployed for Disaster Recovery (in a WAN) or in a Tertiary configuration, WAN Compression is by default configured to Automatic. Neverfail recommends that this setting not be changed unless specifically instructed to do so by Neverfail Support.

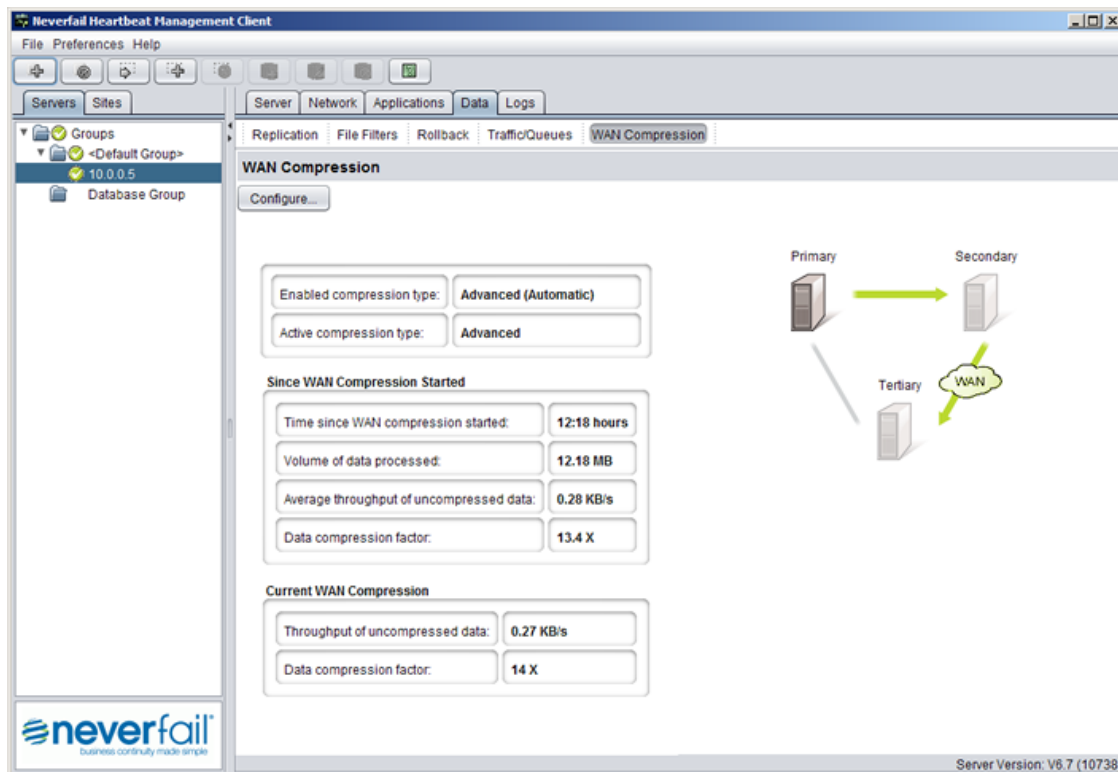


Figure 130: Data: WAN Compression page

Configure WAN Compression

1. To configure WAN Compression, click **Configure**. The **WAN Compression Configuration** dialog appears.

The WAN Acceleration feature allows administrators to configure the compression level used when data is transmitted via the Neverfail Channel. Options include:

- *Allow Neverfail Heartbeat to select the compression type automatically* - Neverfail Heartbeat will automatically configure the compression based upon the cluster configuration
- *Advanced* - Consists of compression and deduplication
- *Standard* - Compression only
- *None* - Neither compression or deduplication




Figure 131: WAN Compression Configuration

2. Select the *Compression Type* (*Allow Neverfail Heartbeat to select the compression type automatically, Advanced, Standard, None*).
3. Select a location for the WAN Compression data log. There are three ways to do this:
 - Type a path into the *WAN Compression data log location* text box
 - Click **Browse** and navigate to the location
 - Select the check box to *Allow Neverfail Heartbeat to set the WAN Compression data log location automatically*
4. Click **OK** to save the configuration or **Cancel** to close the dialog without saving the configuration.

Reference

Appendix

A

Other Administrative Tasks

Configure Alerts

The **Event Log** page lets you configure the Neverfail Heartbeat server to send predefined alerts to remote Neverfail Heartbeat administrators via email.

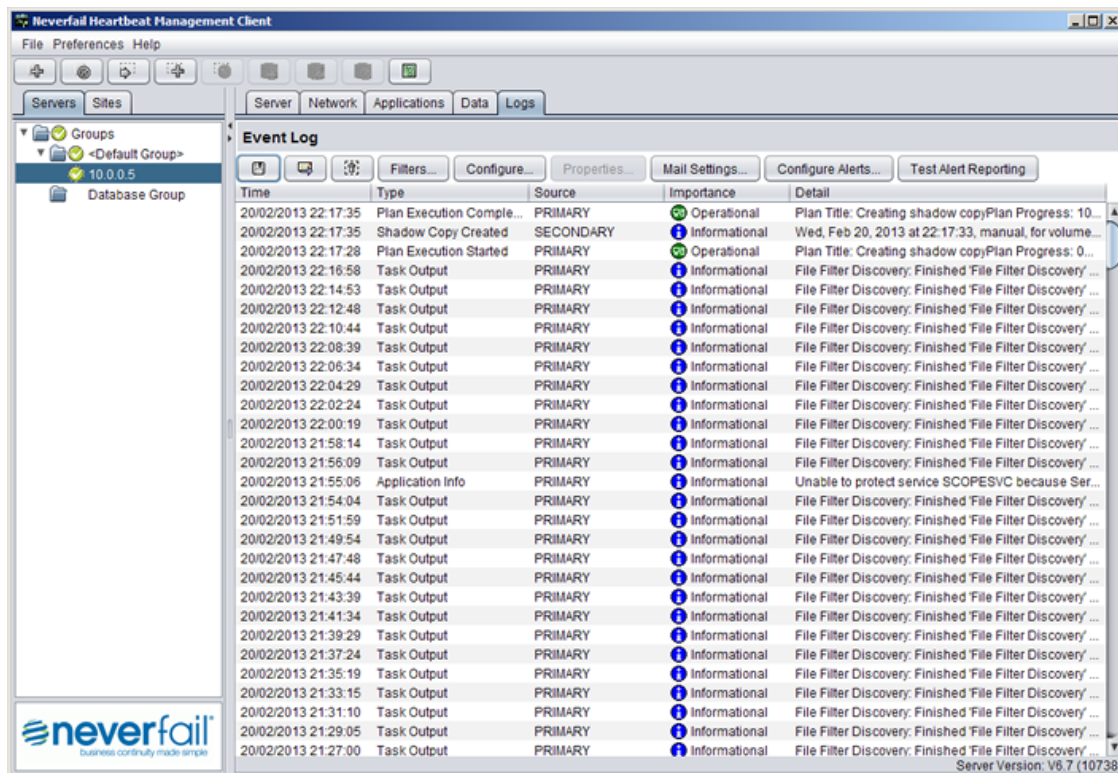


Figure 132: Logs: Event Logs

1. Click **Configure Alerts** (near the top of the **Event Log** page) to open the **Configure Alerts** dialog.

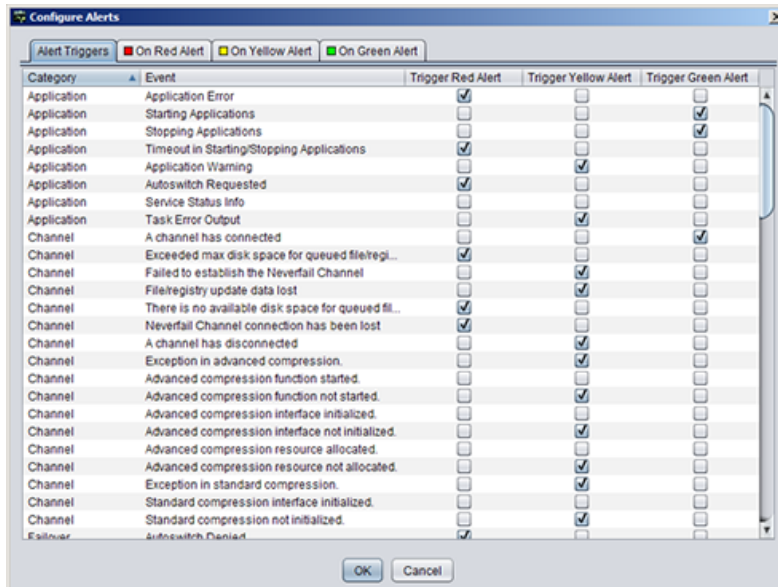


Figure 133: Configure Alerts: Alert Triggers

There are three alert states that can be configured: *Red* alerts, which are critical alerts, *Yellow* alerts, which are less serious, and *Green* alerts which are informational in nature and can be used for notification of status changes (for example, a service that was previously stopped now is started). The alerts are preconfigured with the recommended alerting levels.

- Each alert can be re-configured to trigger as a red, yellow, or green alert or no alert by selecting or clearing the appropriate check boxes. After the alert trigger levels are defined, click **OK** to save the configuration.

Configure Alert Reporting

Neverfail Heartbeat can alert the administrator or other personnel and route logs via email when an Alert condition exists. To configure this capability, in the *Event Log* pane, click **Mail Settings**. In the **Mail Settings** dialog, enter the *Outgoing mail server (SMTP)* of each server in the Cluster. Enter the mail server name using its fully qualified domain name. Next, configure the default *Send mail as* email address. This can be customized but the email address used must be an email account authorized to send mail through the SMTP server.

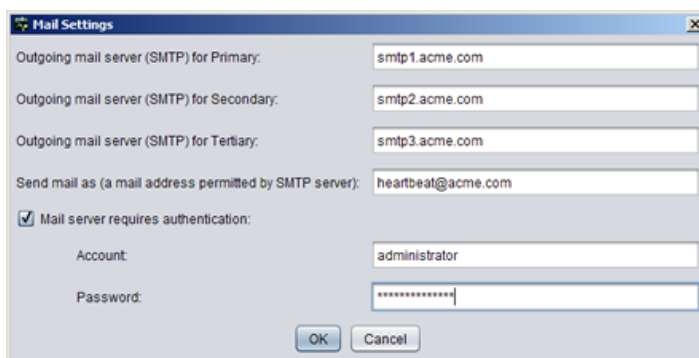


Figure 134: Mail Settings

Note: Where Neverfail Heartbeat is protecting an Exchange Server, it is not recommended to configure the alerts to use the protected Exchange server and is advisable if at all possible to use a different Exchange server somewhere else within the organization.

Where SMTP servers require authentication to accept and forward SMTP messages, select the *Mail Server requires authentication* check box and specify the credentials for an appropriate authenticated user account. Click **OK** to save the changes or click **Cancel** to close the dialog without making any changes.

After the trigger levels are configured and the email server defined in the **Event Log** page **Mail Settings** dialog, configure the recipients of email alerts in the **Configure Alerts** dialog. Email alerts for Red, Yellow, and Green alert triggers can be sent to the same recipient, or configured separately to be sent to different recipients depending on the level of alert.

Add and Remove Email Alert Recipients

The process for adding recipients is the same for all three trigger levels.

1. Click on the tab for the type of alert (Red, Yellow, and Green) and select the *Send mail* check box.

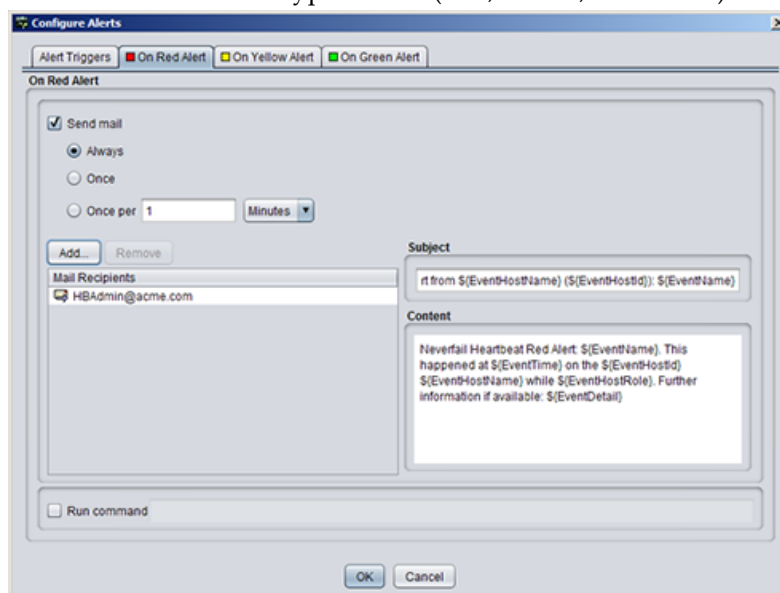


Figure 135: Adding and Removing Email Alert Recipients

2. Select how many times to send the email (*Always*, *Once*, or *Once per [user configurable time period]*)
3. Click **Add**.
The **Add Mail Address** dialog opens.
4. Enter a recipient's fully qualified email address into the text box and click **OK** to add the recipient. Each recipient is entered into a row in the *Mail Recipients* list for the trigger level alert.
5. Repeat steps 3 and 4 to add more recipients.
To remove a recipient from the *Mail Recipients* list, select the recipient in the list and click **Remove**.
The *Subject* and *Content* of the alert emails for all three alerts can be adjusted to suit the environment. Neverfail recommends using the pre-configured content and adding customized content as needed.

Note:

On the *Configure Alerts* pane, the *On Red Alert*, *On Yellow Alert* and *On Green Alert* tabs allow the frequency, recipient, and text of emails to be configured.

When *Send mail* is selected, there are three alternatives:

- *Always* – this will always send an email if this alert type is triggered.
- *Once* – this will send an email once for each triggered alert. An email will not be sent again for the same triggered alert, until *Neverfail Heartbeat* is re-started
- *Once per* – within the time period selected, an email will only be sent once for the same triggered alert, subsequent emails for that trigger will be suppressed. Once the time period has expired, an email will be sent if the same alert is triggered

Note that the *Once* and *Once Per* options will operate per individual triggered Alert. Multiple emails will be sent, but only if the event being alerted is different.

Using WScript to Issue Alert Notifications

An alternative way of issuing notifications for alerts is to run a command by selecting the *Run Command* check box under the relevant alert tab and typing a command into the associated text box. This command can be a script or a command line argument to run on the alert trigger and requires manual entry of the path to the script or command.

The pre-configured WScript command creates an event in the *Application Event Log* and can be customized to include the *Neverfail Heartbeat* specific informational variables listed in the following table.

Table 4: Neverfail Heartbeat Variables

Variables	Values
\$EventHostID	Host ID
\$EventHostName	Host name
\$EventHostRole	Role of the host at the time of the event
\$EventId	ID of event as listed above
\$EventName	Human-readable name of event
\$EventDetail	Detail message for event
\$EventTime	Time at which event occurred

For example, the following command line argument creates an event in the *Application Event Log* that includes the machine that caused the alert, the time the alert happened, the name and details of the alert:

```
Wscript //T:10 $(installdir)\bin\alert.vbs "Neverfail Heartbeat alert on
$EventHost at $EventTime because $EventName ($EventDetail). Event Id is
$EventId"
```

After the alert recipients and/or actions to run are defined, click **OK** to save the changes and enforce the defined notification rules or click **Cancel** to close the dialog without making any changes.

Test Alert Reporting

Click **Test Alert Reporting** to perform a test of the alert reporting feature without triggering an alert and risking the operation of the active server.

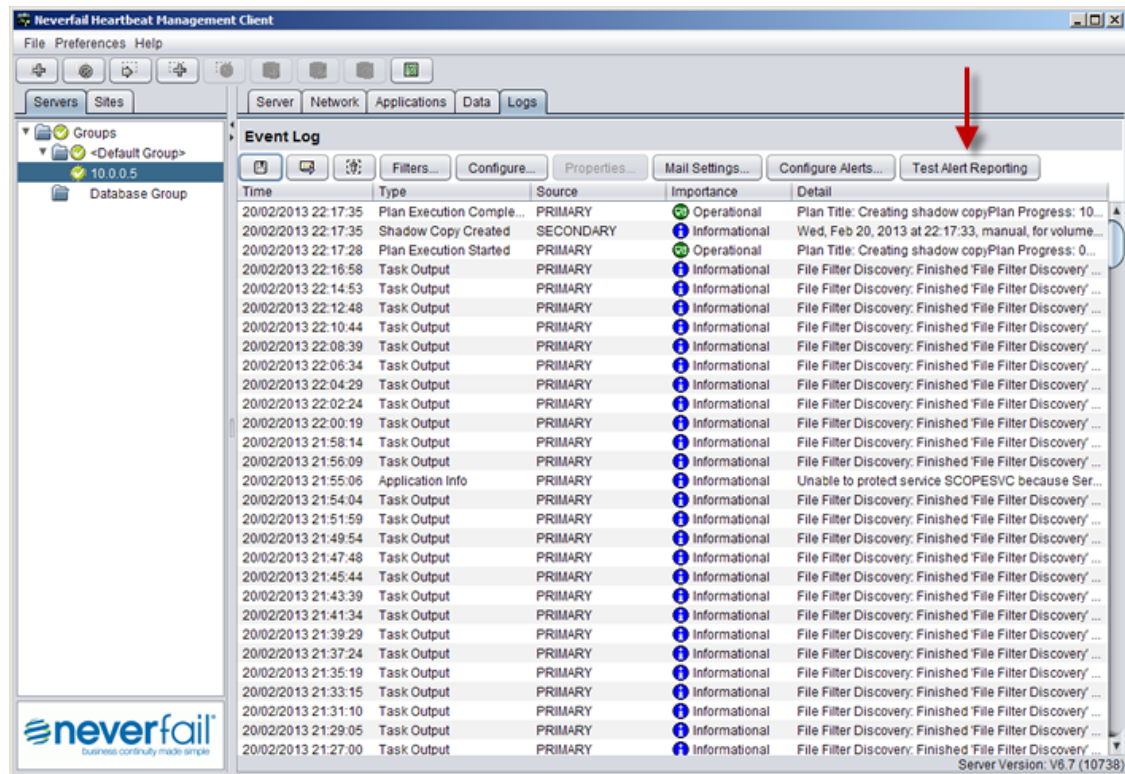


Figure 136: Test Alert Reporting

Configure Event Log Files

To configure default settings for log files, click **Configure** to invoke the **Event Log Configuration** dialog. Select the *General* tab to configure the log file. This dialog allows you to define where the exported comma separated variable file is stored and the name of the file by entering the path and filename manually or browsing to a location using the browse feature. Click **Browse** to open an Explorer type interface and navigate to the appropriate location.




Figure 137: Event Log Configurations: General

The length of the event list can also be adjusted using the *Record At Most* option. The default is to record 300 events but changing the value increases or decreases the length of the log list accordingly. After the logs are configured, click **OK** to commit the changes.

Review Event Logs

The events that Neverfail Heartbeat logs are listed chronologically (by default) in the *Event Log* pane, the first log appears at the top and subsequent logs below it. The display order for the events can be sorted either descending or ascending by clicking on the column heading.

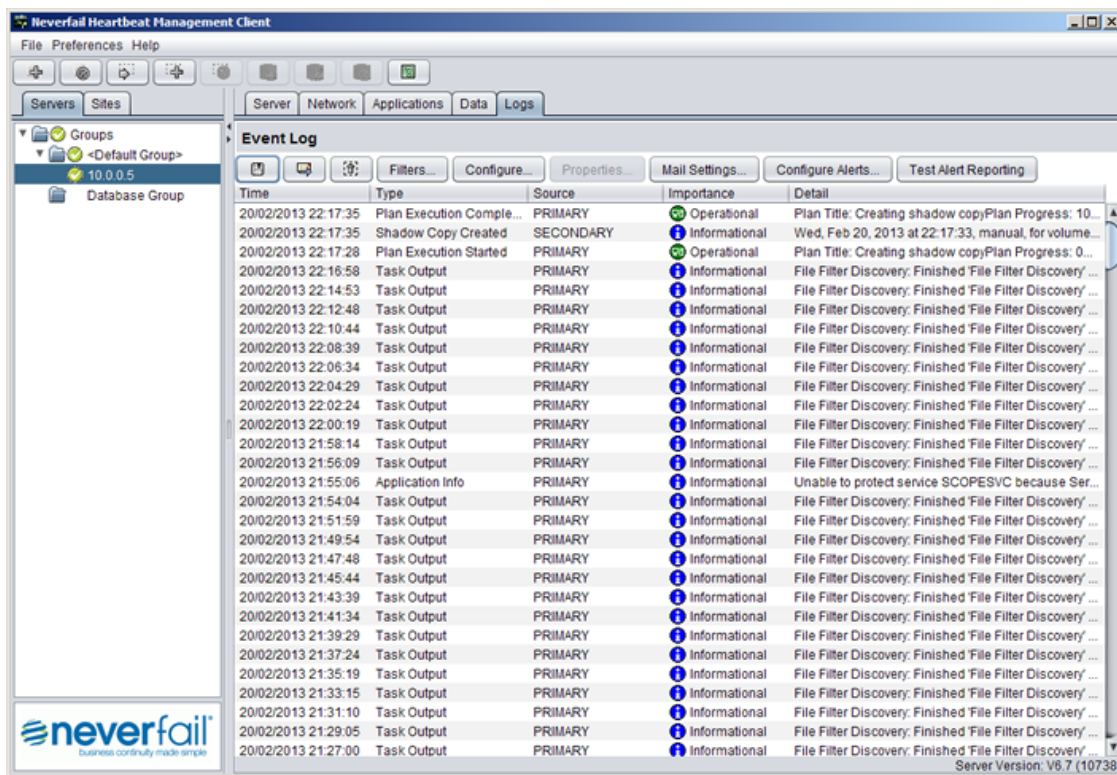


Figure 138: Event Log page

The events listed in the *Event Log* pane show the time the event happened, its importance, the type of event that triggered the log, and its detail.

Since the detail in the data grid is truncated, it may be necessary to review the log in more detail by double-clicking its entry in the pane.

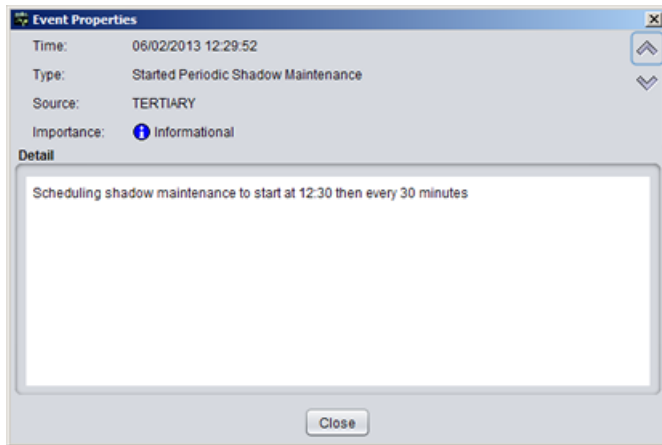


Figure 139: Event Log Properties

The **Event Properties** dialog gives the full detail and trace of the log that caused the event along with the source of the error aiding in troubleshooting. Further logs can be reviewed without having to close this window by using the **Up** and **Down** arrows of the dialog box to scroll through the list of logs. This can help identify the source of the problem when many simultaneous events occur. The **Event properties** dialog may be closed by clicking **Close**.

There are four categories of importance of events that Neverfail Heartbeat by default is configured to log:

Table 5: Neverfail Heartbeat Event Categories

Icon	Definition
	These are critical errors within the underlying operation of Neverfail Heartbeat and can be considered critical to the operation of the system.
	Warnings are generated where the system finds discrepancies within the Neverfail Heartbeat operational environment that are not deemed critical to the operation of the system.
	System logs are generated following normal Neverfail Heartbeat operations. Review these to verify the success of Neverfail Heartbeat processes such as file synchronization.
	Information events are similar to system logs but reflect operations carried out within the graphical user interface rather than operations carried out on the Neverfail Heartbeat Server service itself such as logging on etc.

The list of logs that Neverfail Heartbeat records may be filtered to hide less important logs by clicking **Filters** to invoke the **Event Log Filters** dialog, selecting the *Show Events of at Least* check box in the *Importance group*, selecting the importance level from the drop down list, and clicking **OK**. Only logs equal to or above the selected severity are displayed.




Figure 140: Event Log Filters

You can filter logs to display a subset of entries between a specific date and time range by selecting the *Only Show Events From* check box and adjusting the start and end date, time, and clicking **OK**.

Table 6: Event Log Buttons

Icon	Definition
	Remove all entries from the event log — Click to clear the list.
	Export event log as comma-separated text — Click to export the list to a comma separated variable file. Configure the data export file name and path through the Event Log Configuration dialog (click Configure).
	Mail event log to recipients immediately — Click to email the list to recipients immediately.

Appendix

B

Troubleshooting

Two Active Servers

The occurrence of two active servers is not by design and when detected, must be resolved immediately. When there are two identical active servers live on the same network, Neverfail refers to the condition as Split-brain syndrome.

Symptoms

Split-brain syndrome can be identified by the following symptoms:

1. Two servers in the Cluster are running and in an active state. This is displayed on the task bar icon as P/A (Primary and active), S/A (Secondary and active), or T/A (Tertiary and active).
2. An IP address conflict may be detected in a Cluster running Neverfail Heartbeat on the Principal (Public) IP address.
3. A name conflict may be detected in a Cluster running Neverfail Heartbeat. In a typical WAN environment, the Primary and Secondary (for a Pair), or Primary/Secondary and Tertiary servers (in a Trio) connect to the network using different IP addresses and no IP address conflict occurs. If the servers are running with the same name, then a name conflict may result. This happens only when both servers are visible to each other across the WAN.
4. Clients (for example, Outlook) cannot connect to the server running Neverfail Heartbeat.

Causes

Two active servers (Split-brain syndrome) can be caused by a number of issues. It is important to determine the cause of the Split-brain syndrome and resolve the issue to prevent reoccurrences of the issue. The most common causes of two active servers are:

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

Resolutions

After split-brain syndrome has occurred, the server with the most up-to-date data must be identified.

Note: Identifying the wrong server at this point can result in data loss. Be sure to reinstate the correct server.

The following can help identify the server with the most up-to-date data:

1. Review the date and time of files on both servers. The most up-to-date server should be made 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 your server. This can help identify the MAC address of the server currently visible to clients.

Note: If the two active servers have both been servicing clients, perhaps at different WAN locations, one and only one server can be made active. Both servers contain recent data, which cannot be merged using Neverfail Heartbeat. One server must be made active and one server passive before restarting replication. After replication is restarted, ALL data on the passive server is overwritten by the data on the active server. It may be possible to extract the up-to-date data manually from the passive server prior to restarting replication. Consult the Microsoft knowledge base for information regarding various tools that may be used for this purpose. For further information, contact your Neverfail Support representative.

To Resolve Two Active Servers (Split-Brain Syndrome), perform the following steps:

1. Identify the server with the most up-to-date data or the server you prefer to make active.
2. Shutdown Neverfail Heartbeat on all servers (if it is running).
3. On the server you select to make passive, right-click the task bar icon, and select **Configure Server Wizard**.
4. Click the *Machine* tab and set the server role to passive.

Note: Do not change the Identity of the server (Primary, Secondary, or Tertiary).

5. Click **Finish** to accept the changes. Reboot this server.
6. Start Neverfail Heartbeat (if required) and verify that the task bar icon now reflects the changes by showing **P/-** (Primary and passive) or **S/-** (Secondary and passive) or **T/-** (Tertiary and passive).
7. On the active server, right-click the task bar icon and select **Server Configuration Wizard**.
8. Click the *Machine* tab and verify that the server role is set to active.
9. Click **Finish** to accept the changes. Reboot this server.

Important: As the server restarts, it connects to the passive server and starts replication. When this happens data on the passive server is overwritten by the data on the active server.

10. Start Neverfail Heartbeat (if required) and verify that the task bar icon now reflects the changes by showing **P/A** (Primary and active), **S/A** (Secondary and active), or **T/A** (Tertiary and active).
11. Log into the Neverfail Heartbeat Management Client.
12. Verify that the servers have connected and replication has started.

Two Passive Servers (In a Pair)

The Primary and Secondary servers are both passive at the same time.

Symptoms

The first indication that Neverfail Heartbeat may be experiencing two passive servers is when users are unable to connect to protected applications. This situation can prove serious to your business, and must be addressed immediately. If you have already configured alerts, you are notified that replication is not functioning properly.

Causes

- Two passive servers generally results from some kind of sudden failure on the active server — for example, unexpected termination of the Neverfail Heartbeat R2 Service, a transient power failure, a server reset triggered from hardware power or reset buttons, or any other type of unclean shutdown. Following an unclean shutdown, an active server automatically assumes the passive role to isolate itself from the network until the failure can be investigated.
- The active server suffers a failure before completion of the handshake, which establishes the Neverfail Channel. In this situation, the passive server has no way of detecting that the active server is not responding when the failure occurs - no channel connection was established, so it is impossible for the passive server to determine the condition of the active server. The active server may suffer a transient failure as described above; and the passive server cannot respond by failing over into the active role. This leaves both servers in the passive role.
- Both Primary and Secondary server experience a power outage simultaneously (for example, because they are using the same power source and neither is attached to a UPS). In this situation, a failover is not possible. 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 an attempt is made to start Neverfail Heartbeat without reconfiguring one server as active, Neverfail Heartbeat responds with the warning: No active server amongst [PRIMARY, SECONDARY]*

Resolution

To resolve two passive servers, perform the following steps:

1. Determine which server to make active.
2. If Neverfail Heartbeat is running on either server, shut it down. Leave any protected applications running on the server you selected to make active.
3. On the server you selected to make active, open the **Configure Server Wizard**, and select the active role. Do NOT change the Identity (Primary / Secondary). Save the changes and exit the wizard.
4. On the server you selected to make passive, open the **Configure Server Wizard**, and confirm that the role is passive. Do NOT change the Identity (Primary / Secondary). Exit the wizard.
5. Reboot the passive server. This ensures that all protected application services are stopped.
6. Start Neverfail Heartbeat on both servers.

Three Passive Servers (In a Trio)

The Primary, Secondary, and Tertiary servers are all passive at the same time.

Symptoms

The first indication that Neverfail Heartbeat may be experiencing three passive servers is when users are unable to connect to protected applications. This situation can prove serious to your business, and

must be addressed immediately. If you have already configured alerts, you are notified that replication is not functioning properly.

Causes

- Three passive servers generally results from some kind of sudden failure on the active server — for example, unexpected termination of the Neverfail Heartbeat R2 Service, a transient power failure, a server reset triggered from the **Power** or **Reset** button, or any other type of unclean shutdown. Following an unclean shutdown, an active server automatically assumes the passive role to isolate itself from the network until the failure can be investigated.
- The active server suffers a failure before completion of the handshake, which establishes the Neverfail Channel. In this situation, the passive servers have no way of detecting that the active server is not responding when the failure occurs - no channel connections were established, so it is impossible for the passive servers to determine the condition of the active server. The active server may suffer a transient failure as described above; and the passive servers cannot respond by failing over into the active role. This leaves all three servers in the passive role.
- Primary, Secondary, and Tertiary server experience a power outage simultaneously (for example, because they are using the same power source and are not attached to a UPS). In this situation, a failover is not possible. When the servers are rebooted, each displays the following error message: Cannot start replication because previous run did not shutdown properly. Check configuration.

Note: *If an attempt is made to start Neverfail Heartbeat without reconfiguring one server as active, Neverfail Heartbeat responds with the warning:*

No active server amongst [PRIMARY, SECONDARY, TERTIARY]

Resolution

To resolve three passive servers, perform the following steps:

1. Determine which server to make active.
2. If Neverfail Heartbeat is running on any of the servers, shut it down. Leave any protected applications running on the server you selected to make active.
3. On the server you selected to make active, open the **Configure Server Wizard**, and select the active role. Do NOT change the identity (Primary / Secondary / Tertiary). Save the changes and exit the wizard.
4. On the other two servers you selected to make passive, open the **Configure Server Wizard**, and confirm that the role is passive on each server. Do NOT change the identity (Primary / Secondary / Tertiary). Exit the wizard.
5. Reboot the passive servers. This ensures that all protected application services are stopped.
6. Start Neverfail Heartbeat on all three servers.

Invalid Neverfail Heartbeat License

The Neverfail Heartbeat License is generated from the HBSIG of the host machine. This unique key is generated by examining the Fully Qualified Domain Name (FQDN), Machine SID, and software installed on the server. A valid license key must match the HBSIG.

Symptoms

During normal operations, you receive an error message stating your Neverfail Heartbeat License key has expired or Neverfail Heartbeat fails to start after rebooting the server or stopping Neverfail Heartbeat.

Causes

A license key can become invalid for any of the following reasons:

- Taking a server out of a domain and adding it to another domain.
- The Neverfail Heartbeat License has expired - If a licensing problem arises during an implementation, Neverfail Heartbeat may provide a temporary or time-limited license so that the implementation can proceed. Temporary or time-limited licenses have a defined expiration date, and prevents Neverfail Heartbeat from starting when the date is exceeded.
- Windows Management Instrumentation (WMI) hung or not running. Neverfail Heartbeat uses WMI to validate the license on the Primary server and if WMI is hung or not running validation cannot complete.

Resolutions

1. If the invalid license error is due to changes in the domain status of the Primary server, or expiration of a temporary or time-limited Neverfail Heartbeat License key, simply generate a new license key for the Primary server. To learn more about generating a new Neverfail License key see Knowledgebase article #507 — How to request a license key for Neverfail Heartbeat.
2. If the invalid license error is not due to expiration of a temporary or time-limited Neverfail Heartbeat License key, review the Windows Services and ensure that WMI is running. If WMI is running, stop the WMI Service, restart it, and then attempt to start Neverfail Heartbeat.

Synchronization Failures

When Neverfail Heartbeat is started, a Full System Check runs to ensure that:

- All protected Registry Keys and values from the active server are present on the passive servers.
- All protected File/Folder structures from the active server are present on the passive servers.

After the Full System Check finishes, the File System Status and the Registry Status should be in a *Synchronized* status. There may be cases when the File System Status or the Registry Status is shown as *Out-of-sync* or *Synchronized and busy processing*. Some of the cases are described below, with possible reasons and workarounds.

Services Running on the Passive Server

Symptoms

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

Causes

A service that is running on the passive server may open a protected file for exclusive access. If Neverfail Heartbeat attempts to update a file which has been opened in this way, the following error is logged by the Apply component: [N29]The passive Neverfail 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.

Services that keep files locked on the passive server might be:

- Protected application services
- File-level anti-virus tool services
- The NNTP service in an Neverfail Heartbeat for IIS deployment (if the `\Inetpub` folder is shown as *Out-of-sync*)
- IISAdmin service in an Neverfail Heartbeat for IIS deployment (if `C:\WINDOWS\system32\inetsrv\MetaBase.xml` is shown as *Out-of-sync*). IISAdmin service starts on the passive after a reboot of the server and must be stopped manually.

Resolutions

Until the file is closed on the passive server, Neverfail Heartbeat reports that the file's status, and hence the *File System Status*, is *Out-of-sync*.

To resolve an *Out-of-sync* system status, take the actions below:

1. Ensure Protected Application services are set to *Manual* on both servers and that they are not running on the passive server(s).
2. Ensure that the *Recovery Actions* set from the Service Control Manager (SCM) for the Protected Application services are *Take No Action* (otherwise, the Protected Application services are restarted by the SCM).
3. Ensure that file-level anti-virus is not part of the protected set as the file-level anti-virus and the corresponding services are running on both servers.
4. Ensure the NNTP service is not running on the passive server in an Neverfail Heartbeat for IIS deployment (if `\Inetpub` folder is shown as *Out-of-sync*). This is valid for some of the Exchange implementations as well, where IIS Admin is protected.
5. Ensure that IISAdmin is not running on the passive server in an Neverfail Heartbeat for IIS deployment (if `C:\WINDOWS\system32\inetsrv\MetaBase.xml` is *Out-of-sync*) if IISAdmin service is started on the passive.

Neverfail Channel Incorrectly Configured

Symptoms

If the Neverfail Channels are not properly configured, they cannot initiate the handshake to establish communications through the channel connection. Failure to establish the channel connection prevents a Full System Check and leaves the File System Status and Registry Status as *Out-of-sync*.

Causes

The most common Neverfail Channel configuration errors are:

- Neverfail Heartbeat Packet Filter enabled on one or more channel NICs
- Channel IP addresses configured in different subnets (in LAN configurations)
- In a WAN configuration, no static routes between the channel NICs

Resolutions

To resolve a Neverfail Channel configuration error:

1. Verify that the Neverfail Heartbeat Packet Filter is disabled on channel NICs.

2. Verify that channel IP addresses are properly configured.
3. In a WAN configuration, verify that static routes between channel NICs are properly configured.
4. Ensure that NetBIOS settings on the channel NICs have been disabled.

Incorrect or Mismatched Disk Configuration

Common disk configuration errors which may affect a Cluster:

Symptoms

When Neverfail Heartbeat starts, the complete set of File Filters is checked for consistency. If any of the entries points to a non-existent drive letter or to a non-NTFS partition, the list of File Filters is reset to the default value of `C:\Protected**`. This is a safety measure; Neverfail Heartbeat requires the same drive letter configuration on the Primary, Secondary, and Tertiary servers, and only supports protection of NTFS partitions.

Causes

Different partition structures on the Primary, Secondary, and Tertiary servers, resulting in one or more file filters pointing to drives which cannot be protected on all servers. For example:

- The Primary server has drive `G:`, which is a valid NTFS partition; but there is no corresponding drive on the Secondary or Tertiary server
- The Primary server has drive `G:`, which is a valid NTFS partition; but the equivalent drive on the Secondary or Tertiary server is a CD / DVD drive or a FAT / FAT32 partition, which cannot be protected by Neverfail Heartbeat.

In either case, if a file filter is configured 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**`

Resolutions

If this occurs, follow the steps documented in KB-500 — The set of File Filters is reset to `C:\Protected**`. What should I do next?

The First Passive Server has Less Available Space than the Active Server

Free up some additional disk space on the first passive server. Make sure you are not deleting data from the protected set as you might lose data in the event of a switchover. This may require you to update the disk subsystem on the first passive server. When complete, you must manually start replication.

Symptoms

Replication stops and the following error is reported:

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

Causes

The first passive server has less available disk space than the active server and this prevents replication of updates to the first passive server because the quantity of updates from the active server exceeds the available disk space on the first passive server.

Resolution

Free up some additional disk space on the first passive server. Make sure you are not deleting data from the protected set as you might lose data in the event of a switchover. This may require you to update the disk subsystem on the first passive server. When complete, you must manually start replication.

Unprotected File System Features

Symptoms

Another possible reason why Neverfail Heartbeat cannot synchronize certain files or directories is the presence in the replication set of so-called “unprotected” file system features.

The default behavior for Neverfail Heartbeat in the presence of Unprotected Features from category 2 (Extended Attributes and file encryption) is to log an error and set the File System Status to *Out-of-sync*. If these types of files are present in the replication set, replication continues, but the system remains *Out-of-sync*.

Causes

Neverfail Heartbeat does not synchronize if the replication set contains files with unprotected file system features. Unprotected file system features are described by category in KB-321 — Neverfail for File Server: Unprotected Features of the Windows 2003/2008 File System.

Resolutions

Two methods of dealing with these Unprotected Features are described in KB-321 — Neverfail for File Server: Unprotected Features of the Windows 2003/2008 File System. If these features are not essential for the normal operation of your file system, zipping and unzipping the affected files within their parent directory removes the Unprotected Features, allowing the Neverfail Heartbeat to synchronize the file system.

Registry Status is Out-of-Sync

The Registry may be reported as Out-of-Sync when one or more Registry keys fail to synchronize. There are at least two possible reasons.

Resource Issues

Symptoms

Neverfail Heartbeat logs the following error message:

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

Causes

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

Resolutions

This is usually a sign that the server does not have enough virtual memory left. Restart the server to correct this problem.

Registry Security Issues

Symptoms

Neverfail Heartbeat is unable to read/sync/replicate the registry.

Causes

If a protected registry key has permissions that deny Write access to the System account, Neverfail Heartbeat may be unable to synchronize or replicate it.

Resolutions

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

Channel Drops

Performance Issues

Symptoms

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 channel connection between the servers is lost.

Causes

This condition is unusual and generally points to an application, or Windows itself, experiencing a fault on one of the passive servers. The most likely issue here is a sudden reboot / restart of the passive server and may be due to one of the following causes:

- The server is configured for automatic software update management and some updates force a server reboot.
- There is a software or Operating System issue which occasionally results in a BSOD and system restart.
- The Neverfail Heartbeat R2 service itself experiences problems and may hang or terminate unexpectedly.

Resolutions

- Determine the likely source of the hang or reboot by examining the Windows event logs.
- Alternatively, if the server does not show any evidence of a system restart or application hang, the issue may be due to one or both of the channel NICs forcing a channel disconnection.

Passive Server Does Not Meet Minimum Hardware Requirements

Symptoms

The data rate between the servers is very high during a Full System Check and the channel drops.

Causes

A passive server does not meet the recommended hardware requirements for Neverfail Heartbeat or it meets them but is much less powerful than the other server(s) in the Cluster. The underpowered server cannot apply the received replication data from the active or other passive server at the rate that the data is sent to this passive server

Resolutions

To avoid reinstalling your Neverfail Heartbeat solution, it is best to tackle this issue by upgrading the hardware (for example, memory and or CPU) on the passive server. It is important to establish the identity (Primary, Secondary, or Tertiary) of the affected server before you perform the upgrade.

Hardware or Driver Issues on Channel NICs

Symptoms

The Neverfail Channel drops or disconnects and reconnects intermittently.

Causes

- Old/wrong drivers on the channel NICs
- If the physical connection used for the Neverfail Channel connection uses a hub or Ethernet switch, a hardware fault may cause the channel to drop
- Defective Ethernet patch or crossover cables
- Improper configuration of the NICs used for the channel connection
- ISP problems in a WAN environment

Resolutions

1. Verify that channel NIC drivers are the correct/latest versions. This is a known issue with HP/Compaq ProLiant NC67xx/NC77xx Gigabit Ethernet NICs but may affect other NIC types as well. See KB-116 — vCenter Server Heartbeat and Gigabit Ethernet NIC drivers. (NC77XX).
2. Verify hubs and Ethernet switches are operating properly. Identify and replace any defective components.
3. Test for defective Ethernet patch or crossover cables and replace if defective.
4. Correctly configure the NICs used for the channel connection.
5. Verify the physical link to identify any ISP problems.

Firewall Connection

In both a LAN or WAN deployment of Neverfail Heartbeat, the channel may be connected via one or more Internet firewalls. Since firewalls are intended to block unauthorized network traffic, it is important to ensure that any firewalls along the route of the channel are configured to allow channel traffic.

Symptoms

The Neverfail Channel cannot connect or connects and disconnects continuously.

Causes

In a WAN deployment, port #57348 (or any other port configured for the Neverfail Channel) is closed on one or more firewalls on the route between the channel NIC on the active server and its counterpart on the passive server.

Resolutions

Open port #57348 (and any other port configured for the Neverfail Channel) on all firewalls on the route between the channel NIC on the active server and its counterpart on the passive server.

Incorrect Neverfail Channel Configuration

Symptoms

IP conflicts are encountered on one of the channel IP addresses. The Neverfail Channel does not connect or connects and disconnects.

Causes

Identical IP addresses at each end of the channel, IP addresses in different subnets without static routing at each end of the channel, or a channel NIC configured for DHCP when a DHCP server is not available.

During installation, Neverfail Heartbeat configures the channel NICs with user provided information. Providing incorrect information or incorrectly modifying the channel NIC configuration after installation can cause the Neverfail Channel to fail communicating.

On rare occasions, if the servers in a Cluster have NICs of the same type in a different order, both the name and IP address of a channel NIC on the Primary server may be transferred to the Principal (Public) NIC on the Secondary and or Tertiary server; or the name and IP address of the Principal (Public) NIC may be transferred to a channel NIC. Similarly, the names of the channel NICs may be reversed on the Secondary and or Tertiary server under these circumstances. If this happens, it can be hard to reconcile the names of the NICs with their physical identities, making it difficult to assign the correct IP address to each NIC on the Secondary and Tertiary servers.

Resolution

It is part of the normal Neverfail Heartbeat installation process to manually assign the correct IP addresses to each NIC on the Secondary and Tertiary server. If there is no channel connection between the servers, verify that the IP addresses on the Secondary and Tertiary server's channel NICs are correctly configured. Verify the settings for the Principal (Public) NIC, since any configuration error here may not be apparent until a switchover is performed or a failover occurs.

It is possible to capture the identities of all of the NICs on the Secondary and Tertiary server prior to installing Neverfail Heartbeat, by opening a Windows Command Prompt on that server and executing the following command:

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

This saves the current name, TCP/IP configuration, and MAC address of each NIC on the Secondary and Tertiary server to a file called `ipconfig.txt`, which is present on that server after the Plug and Play phase of the Neverfail Heartbeat install is complete. At this point, it is possible to compare the pre-install and post-install state of each NIC by running `ipconfig /all` from a Windows command prompt and comparing the output of this command with the content of the file `ipconfig.txt`. The MAC address of each NIC is tied to the physical identity of each card, and never changes - so it is possible to identify each NIC by its MAC address and determine its original name and network configuration, even if these have been updated by the Plug and Play process.

Neverfail Heartbeat Packet Filter is Enabled on the Channel NIC(s)

Symptoms

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

Causes

During the installation of Neverfail Heartbeat, the Neverfail Heartbeat Packet Filter is automatically installed and enabled on all NICs on all servers in the Cluster. Disabling the Packet Filter on the channel NIC(s) on each server is a documented part of the installation of Neverfail Heartbeat. If the Neverfail Heartbeat Packet Filter is left enabled on one or more channel NICs after installation is complete, it may interfere with network traffic across the channel.

Resolutions

In the Properties tab for each channel NIC on both servers, confirm that the check box for the Neverfail Heartbeat Packet Filter is unchecked, meaning that the Neverfail Heartbeat Packet Filter is disabled on that NIC.

Subnet/Routing Issues — In a LAN

Symptoms

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

Causes

The Neverfail Channel may disconnect or fail to connect due to the Principal (Public) NIC and/or one or more channels sharing the same subnet.

Resolutions

If Neverfail Heartbeat is deployed in a LAN environment, the Principal (Public) IP address and the channel IP address on a server should be in separate subnets. If there are multiple redundant channels, each should have its own subnet. Verify the network configuration for each NIC and correct any issues.

Subnet/Routing Issues — In a WAN

Symptoms

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

Causes

When the Neverfail Channel disconnects or fails to connect in a WAN deployment it may be because the static route is not configured or is configured incorrectly.

When Neverfail Heartbeat is deployed in a WAN, it is generally not possible for the Principal (Public) IP address and the channel IP addresses to be in different subnets, since there is usually a single network path between the two servers. To ensure that channel traffic is routed only between the endpoints of the channel, it is necessary to configure a static route between these endpoints.

Resolutions

Refer to KB-466 — How to Create a Static Route for the VMware Channel Connection where the channel and Principal Public IP addresses are on the same subnet in a WAN environment, for a detailed discussion about WAN channel routing issues, and for instructions on how to configure a static route for the Neverfail Channel.

MaxDiskUsage Errors

Disk Usage and Disk Quota Issues

Neverfail Heartbeat uses queues to buffer the flow of replication data from the active server to the first passive server and from the first passive to the second passive server. This configuration provides resilience in the event of user activity spikes, channel bandwidth restrictions, or channel drops (which may be encountered when operating in a WAN deployment). Some types of file write activity may also require buffering as they may cause a sharp increase in the amount of channel traffic. The queues used by Neverfail Heartbeat are referred to as either the send queue or the receive queue with each server in the Cluster maintaining both a send queue and receive queue for each channel connection.

Send Queue

Neverfail Heartbeat considers the send as 'unsafe' because the data in this queue is awaiting replication across the channel to a passive server and is vulnerable to loss in the event of a failover. As a result of failover, some data loss is inevitable, with the exact amount depending upon the relationship between current channel bandwidth and the required data transmission rate. If the required data transmission rate exceeds current channel bandwidth, the send queue fills; if the current channel bandwidth exceeds the required data transmission rate, the send queue empties. This situation is most commonly seen in a WAN environment, where channel bandwidth may be restricted. In a LAN with normally high bandwidth on a dedicated channel, the size of the send queue is zero or near zero most of the time.

Note: On a server that is not protected with Neverfail Heartbeat, all data is technically 'unsafe' because it is possible to lose all data if the server fails.

Receive Queue

The target queue on the passive server is called the receive queue and is considered safe. Neverfail Heartbeat considers the receive queue safe because the data in this queue has already been transmitted across the channel from the active or first passive server, and is not lost in the event of a failover, since all updates to the passive server are applied as part of the failover process.

The queues (on all servers) are stored on-disk, by default in the <Neverfail Heartbeat Install Directory>\R2\log, with a quota configured for the maximum permitted queue size (by default, 10 GB on each server). Both the queue location and the quota are configurable.

There are two ways to set the queue size:

- With Neverfail Heartbeat started, open the Neverfail Heartbeat Management Client and select **Data: Traffic/Queues**. Set the value of *Max Disk Usage* and click **OK**. It is necessary to shut down and restart Neverfail Heartbeat (specify that the stopping of protected applications is not necessary) for the change to take effect.
- With Neverfail Heartbeat shut down on the active server, open the **Configure Server Wizard** and select the **Logs** tab. Set the value of *Maximum Disk Usage* and click **Finish**.

***Note:** Neverfail Heartbeat is a symmetrical system, and can operate with any server in the active role. For this reason, the queue size is always set to the same value for all servers.*

MaxDiskUsage Errors

If Neverfail Heartbeat exceeds its pre-configured queue size, it reports an error message. There are several possible reasons for this, with the most common ones shown below.

When Neverfail Heartbeat reports `[L9] Exceeded the maximum disk usage (NFChannelExceededMaxDiskUsageException)`, the following conditions exist:

- On the active server, it indicates that the size of the send queue has exceeded the disk quota allocated for it.
- On a passive server, it indicates that the size of the receive queue or send queue has exceeded the disk quota allocated for it.

Neither of these conditions is necessarily fatal, or even harmful; but it is important to try to determine the sequence of events, which led to the condition appearing in the first place.

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

Symptoms

Replication stops, and the Neverfail Heartbeat *Event Log* displays the error `[L9]Exceeded the maximum disk usage`, originating from the ACTIVE server.

Causes

As stated previously, if there is a temporary interruption in the Neverfail Channel, or there is insufficient channel bandwidth to cope with the current volume of replication traffic, the send queue may begin to fill. If the situation persists, the size of the queue may eventually exceed the configured disk quota.

Resolutions

Assuming there are no other channel connection issues (see KB-992 — Neverfail Channel Drops) you can increase the amount of disk space allotted to the queues to prevent this situation recurring. The default setting is 10 GB, which may be insufficient on servers with a large volume of replication traffic and/or limited channel bandwidth. If you have sufficient disk space, set the queue size to zero (unlimited). This allows Neverfail Heartbeat to utilize any free disk space to store the queues.

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

Symptoms

Replication stops, and the Neverfail Heartbeat *Event Log* displays the error `[L9]Exceeded the maximum disk usage`, originating from a PASSIVE server.

Causes

- In this situation, the bottleneck lies between the Neverfail Channel NIC and the disk subsystem on a passive server. When replication traffic passes across the channel faster than it can be written to disk on the passive server, it is buffered temporarily in the passive server's receive queue. As before, if this situation persists, the size of the queue may eventually exceed the disk quota allotted.

- If the passive server is much less powerful than the active server, in terms of processor speed, RAM or disk performance, it may lag behind the active server during periods of high replication activity. If you suspect this is the case, it may be useful to monitor one or more Windows performance counters to determine which component is experiencing sustained high activity. Intensive page file use or persistently large disk queue length may indicate a problem, which can be solved by upgrading one or more physical components of the server.
- Note that any server can be active or passive. If the Secondary or Tertiary server is more powerful than the Primary server, hardware-related issues might only occur while the Secondary or Tertiary server is in the active role.

Resolutions

If you have multiple physical disks on each server, it may be worth locating the Neverfail Heartbeat send and receive queues on a separate physical disk, away from the Windows directory, the Windows page file, and any protected files to help alleviate disk performance issues. To do this:

1. Shut down Neverfail Heartbeat.
2. Open the **Server Configuration Wizard** and select the *Logs* tab.
3. Set the path for *Message Queue Logs Location* and click **Finish**.
4. Start Neverfail Heartbeat on all servers.

Note: *The selected path is applicable only to the specific server where the change was performed.*

5. You may alleviate the symptoms of this problem by simply increasing the amount of disk space allotted to the queues. If you have reason to suspect that a hardware issue is the root of the problem, it is better to correct that problem at the source if possible.
6. It is also possible for the size of the receive queue to increase sharply in response to certain types of file write activity on the active server. This is most obvious when Neverfail Heartbeat is replicating a large number of very small updates (typically a few bytes each) - the volume of update traffic may be far greater than the physical size of the files on the disk, and so the receive queue in particular may become disproportionately large. This pattern of disk activity is often seen during the population of Full-Text Catalogs in Microsoft SQL Server.
7. Increase the amount of disk space available for the queues, as described above; it may be also help to alleviate the issue by moving the queues to their own physical disk, or upgrading memory or the disk subsystem.
8. Neverfail Heartbeat requires a certain amount of system resources 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 Neverfail Heartbeat). It is always a good idea to ensure that there are sufficient resources for all of the applications and services running on such a server to provide maximum performance, stability, and resilience in the face of changing client, server, and network activity.

[L20]Out of disk space (NFChannelOutOfDiskSpaceException)

Symptoms

Replication stops and the Neverfail Heartbeat *Event Log* displays the error [L20]Out of disk space, originating from either server.

Causes

This is similar to the [L9]Exceeded the maximum disk usage scenario, with one important difference - one of the queues has exceeded the amount of physical disk space available for it, without reaching its quota limit. So, for example, if the maximum queue size is set to 10 GB, but only 3 GB of physical disk space remains, this message is reported if one of the queues exceeds 3 GB in size.

Resolutions

The strategy for dealing with this is simple - it is necessary either to free up more disk space, or to 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

Any piece of 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, which are running at the same time. If the total resource requirement for the applications exceeds the available physical resources, the operating system gracefully attempts to provide resources but some applications may be under-resourced. This may mean that an application cannot obtain enough memory to operate normally, or that a process is required to wait to access the hard disk.

In a situation where applications are competing for resources, it is likely that one or more applications suffer from poor performance. Operations performed by the application may take longer than usual to complete, and in turn, may 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 Neverfail Heartbeat and for servers running any other application. Neverfail Heartbeat is able to monitor system performance counters and provide warnings if 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.

It is very important to ensure that the machines hosting Neverfail Heartbeat meet recommended hardware requirements and are powerful enough to cope with the load imposed by Neverfail Heartbeat, the protected application, and any other critical applications. Neverfail SCOPE Data Collector Service provides users with the information to make this decision at install time, and can monitor server performance while Neverfail Heartbeat is running.

Poor Application Performance

Symptoms

The servers are unable to accommodate the load placed upon them during normal operation.

Causes

This may be due to the Primary server's resource usage in one or more areas being close to the maximum possible before Neverfail Heartbeat was installed.

Resolutions

Neverfail SCOPE Data Collector Service is designed to report on these types of conditions, and can provide warnings if CPU usage or memory usage exceeds a certain percentage of the available resource. The information provided by Neverfail SCOPE Data Collector Service means that the risk of application

slowdown could be minimized by performing any recommended hardware upgrades on the Primary server before Neverfail Heartbeat is installed.

Servers Could Accommodate the Initial Load but the Load has Increased

Symptoms

Application response times have slowed in response to increased user activity.

Causes

It is also possible that the servers may be able to operate normally when Neverfail Heartbeat is first installed, with performance decreasing because of an increase in user activity - for example, the number of users on your Exchange system may increase, or the typical usage pattern for a user may become more intense. This may be a gradual and sustained increase over time; or it may be transient if some specific event triggers a temporary surge in user activity.

Resolution

If the situation is sporadic, it may correct itself when the load decreases. If the increase is sustained and permanent, it may be necessary to upgrade the server hardware to compensate.

One Server is Able to Cope, but the Other Cannot

Symptoms

Applications operate normally when the Primary server is active but slow when the Secondary or Tertiary server is active (or vice versa).

Causes

If there is a large discrepancy in the processing power between the servers, it may be that one of the servers can handle the operational load, and the other cannot. The load on a server is generally higher when it is in the active role and the protected application(s) started, so it is possible that applications run successfully when the Primary server is active, but may experience performance issues when the Secondary or Tertiary is active (or vice-versa). Problems may arise even when the more powerful server is active, such as when resource intensive tasks are running.

Resolutions

It is good practice to ensure that all servers have approximately equivalent processing power, RAM and disk performance. It may be necessary to upgrade the hardware so that servers have roughly the same performance.

Scheduled Resource Intensive Tasks

Symptoms

Resource-intensive scheduled tasks impact performance at certain times.

Causes

System performance may be fine until two or more resource-hungry processes run simultaneously; or, one process may perform actions, which increase the load on Neverfail Heartbeat by triggering

additional (and sometimes unnecessary) replication traffic. Typical examples might be processes such as backups, database maintenance tasks, disk defragmentation or scheduled virus scans.

Resolution

As far as possible, it is good practice to schedule such operations so that they do not overlap, and to schedule them outside regular working hours, when the load imposed on the server by users accessing the protected application is likely to be smaller.

Glossary

Active

The functional state or role of a server when it is visible to clients through the network, running protected applications, and servicing client requests.

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 between proprietary services. Neverfail Heartbeat switchovers and failovers require no changes to AD resulting in switchover/failover times typically measured in seconds.

Active–Passive

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

Advanced Configuration and Power Interface (ACPI)

A specification that dictates how the operating system can interact with the hardware especially where power saving schemes are used. The Primary, Secondary, and Tertiary servers must have identical ACPI compliance.

Alert

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

Asynchronous

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

Basic Input/Output System (BIOS)

The program a personal computer's microprocessor uses to get the computer system started 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.

Cached Credentials

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

Channel Drop

An event in which the dedicated communications link between servers 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 Neverfail Channel.

Checked

The status reported for user account credential (username/password) validation.

Cloned Servers

Servers that have identical configuration settings, names, applications, Security Identifiers (SIDs) and IP addresses, following the installation of Neverfail Heartbeat.

Cloning Process

The Neverfail Heartbeat process whereby all installed programs, configuration settings, and the machine name, Security Identifier (SID), and IP address are copied to another server.

Cluster

A generic term for a Neverfail Heartbeat Pair or Trio and the set of machines (physical or virtual) involved in supporting a single protected server. A Neverfail Heartbeat Cluster can include the machines used in a VMware or Microsoft cluster.

Connection

Also referred to as Cluster Connection. Allows the Neverfail Heartbeat Management Client to communicate with a Neverfail Heartbeat Cluster, either on the same machine or remotely.

Crossover Cable

A network cable that crosses the transmit and receive lines.

Data Replication

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

Data Rollback Module

A Neverfail Heartbeat module that allows administrators to rollback the entire state of a protected application, including files and registry settings, to an earlier point-in-time. Typically used after some form of data loss or corruption.

Degraded

The status reported for an application or service that has experienced an issue that triggered a Rule.

Device Driver

A program that controls a hardware device and links it to the operating system.

Disaster Recovery (DR)

A term indicating how you maintain and recover data with Neverfail Heartbeat in event of a disaster such as a hurricane or fire. DR protection can be achieved by placing the Secondary server (in a Pair) or the Tertiary server (in a Trio) at an offsite facility, and replicating the data through a WAN link.

DNS (Domain Name System) Server

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

Domain

A logical grouping of client server based machines where the administration of 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 2003 and above domains.

Dualed

A way to provide higher reliability by dedicating more than one NIC for the Neverfail Channel on each server.

Failover

Failover is the process by which the first 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 crash of a server.

First Passive

The passive server in a Neverfail Heartbeat Pair or Trio communicating with and receiving replicated data directly from the active server.

Full System Check (FSC)

The internal process automatically started at the initial connection or manually triggered through the Manage Server GUI to perform verification on the files and registry keys and then synchronize 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.

Global Catalog

A global catalog is a domain controller that stores a copy of all Active Directory objects in a forest. The global catalog stores a full copy of all objects in the directory for its host domain and a partial copy of all objects for all other domains in the forest.

Graceful (Clean) Shutdown

A shutdown of Neverfail Heartbeat based upon completion of replication by use of the Neverfail Heartbeat Neverfail Heartbeat Management Client, resulting in no data loss.

Group

An arbitrary collection of Neverfail Heartbeat Clusters used for organization.

Hardware Agnostic

A key Neverfail Heartbeat feature allowing for the use of servers with different manufacturers, models, and processing power in a single Neverfail Heartbeat Cluster.

Heartbeat

The packet of information issued by the passive server across the channel, which the active server responds to indicating its presence.

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 that are used to address a problem in a product.

Identity

The position of a given server in the Neverfail Heartbeat Cluster: Primary, Secondary, or Tertiary.

Install Clone

The installation technique used by Neverfail Heartbeat to create a replica of the Primary server using NTBackup or Wbadmin and to restore the replica to the Secondary and/or Tertiary servers.

Low Bandwidth Module (LBM)

A Neverfail Heartbeat module that compresses and optimizes data replicated between servers over a WAN connection. This delivers maximum data throughput and improves application response time on congested WAN links.

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

The ability of one physical server (hosting more than one virtual server) to protect multiple physical servers.

Network Monitoring

Monitoring the ability of the active server to communicate with the rest of the network by polling defined nodes across the network at regular intervals.

Neverfail Channel

The IP communications link used by the Neverfail system for the heartbeat and replication traffic.

Neverfail Extranet

The Neverfail web site dedicated to supporting partners and customers by providing technical information, software updates, and license key generation.

Neverfail Heartbeat

The core replication and system monitoring component of the Neverfail solution.

Neverfail Heartbeat Packet Filter

The network component, installed on all servers, that controls network visibility.

Neverfail License Key

The key obtained from the Neverfail extranet that allows the use of components in the Neverfail suite; entered at install time, or through the Configure Server Wizard.

Neverfail Pair

Describes the coupling of the Primary and Secondary server in a Neverfail solution.

Neverfail Plug-ins

Optional modules installed into a Neverfail Heartbeat server to provide additional protection for specific applications.

Neverfail SCOPE

The umbrella name for the Neverfail process and tools used to verify the production servers health and suitability for the implementation of a Neverfail solution.

Neverfail SCOPE Report

A report provided upon the completion of the Neverfail SCOPE process that provides information about the server, system environment, and bandwidth.

Neverfail Switchover/Failover Process

A process unique to Neverfail in which the passive server gracefully (switchover) or unexpectedly (failover) assumes the role of the active server providing application services to connected clients.

Neverfail Trio

Describes a set of three coupled servers (Primary, Secondary, and Tertiary) in a Neverfail solution.

Pair

See Neverfail Heartbeat Pair above.

Passive

The functional state or role of a server when it is not delivering service to clients and is hidden from the rest of the network. For a Neverfail Heartbeat Trio, see also First Passive and Second Passive.

Pathping

A route-tracing tool that works by sending 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. On starting the computer, PnP automatically configures the necessary IRQ, DMA and I/O address settings for the attached peripheral devices.

Plug-in

An application specific module that adds Neverfail Heartbeat protection for the specific application.

Pre-Clone

An installation technique whereby the user creates an exact replica of the Primary server using VMware vCenter Converter or other 3rd party utility prior to the initiation of installation and uses the replica as a Secondary and or Tertiary server.

Pre-Installation Checks

A set of system and environmental checks performed as a prerequisite to the installation of Neverfail Heartbeat.

Primary

An identity assigned to a server during the Neverfail Heartbeat installation process that normally does not change during the life of the server and usually represents the production server prior to installation of Neverfail Heartbeat.

Principal (Public) IP Address

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

Principal (Public) Network

The network used by clients to connect to server applications protected by Neverfail Heartbeat.

Principal NIC

The network card which hosts the Principal IP address.

Protected Application

An application protected by the Neverfail Heartbeat solution.

Quality of Service (QoS)

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

Receive Queue

The staging area on a server used to store changes received from another server in the replication chain before they are applied to the disk/registry on the passive server.

Remote Desktop Protocol (RDP)

A multi-channel protocol that allows a user to connect 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 channel, and applying them to the passive server(s) so the servers are maintained in a synchronized state.

Role

The functional state of a server in the Neverfail Heartbeat Cluster: active or passive.

Rule

A set of actions performed by Neverfail Heartbeat when defined conditions are met.

Second Passive

The passive server in a Neverfail Heartbeat Trio communicating with and receiving replicated data directly from the first passive server.

Secondary

An identity assigned to a server during the Neverfail Heartbeat installation process that normally does not change during the life of the server and usually represents the standby server prior to installation of Neverfail Heartbeat.

Security Identifier (SID)

A unique alphanumeric character string that identifies each operating system and each user in a network of 2003/2008/2012 systems.

Send Queue

The staging area on a server used to store intercepted data changes before being transported across to a passive server in the replication chain.

Server Monitoring

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

Shared Nothing

A key feature of Neverfail Heartbeat in which no hardware is shared between the Primary, Secondary, and Tertiary servers. This prevents a single point of failure.

SMTP

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

SNMP

Simple Network Management Protocol (SNMP) is an Internet-standard protocol for managing devices on IP networks.

Split-Brain Avoidance

A unique feature of Neverfail Heartbeat that prevents a scenario in which 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 in which more than one server in a Neverfail Heartbeat Cluster are operating in the active mode and attempting to service clients, resulting in the independent application of different data updates to each server.

Storage Area Network (SAN)

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

Subnet

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

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 Neverfail Channel, from the active server to the first passive server or from the first passive server to the second passive server to ensure the data on the passive server is a mirror image of the protected data on the active server.

System Center Operations Manager (SCOM)

System Center Operations Manager is a cross-platform data center management server for operating systems and hypervisors.

System State

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

Task

An action performed by Neverfail Heartbeat when defined conditions are met.

Tertiary

An identity assigned to a server during the Neverfail Heartbeat installation process that normally does not change during the life of the server and usually represents the disaster recovery server prior to installation of Neverfail Heartbeat.

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 your computer and a specified destination computer.

Trio

See Neverfail Heartbeat Trio above.

Ungraceful (Unclean) Shutdown

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

Unprotected Application

An application not monitored nor its data replicated by Neverfail Heartbeat.

Virtual Private Network (VPN)

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

Windows Management Instrumentation (WMI)

A management technology allowing 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, clusters, and groups.