



Administrator's Guide

Neverfail Engine

Notice

Neverfail, LLC 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, Neverfail, LLC has relied on the best available information published by such parties. Neverfail, LLC is continually developing its products and services, therefore the functionality and technical specifications of Neverfail's products can change at any time. For the latest information on Neverfail's products and services, please contact us by email (info@neverfail.com) or visit our Web site (neverfail.com).

Neverfail is a registered trademark of Neverfail, LLC. All third party product names referred to in this document are acknowledged as the trade marks for their respective owner entities.

Copyright (c) 2026 Neverfail, LLC. All rights reserved.

Contents

Neverfail Engine Concepts

Architecture

Protection

Neverfail Engine Networking Configuration

Neverfail Engine Communications

Neverfail Engine Switchover and Failover Processes

Recovery from a Failover

Managing Neverfail Engine Clusters

Review the Status of Neverfail Engine Clusters and Groups

Exit Neverfail Advanced Management Client

Shutdown Windows with Neverfail Engine Installed

Controlled Shutdown

Configuring Neverfail Engine

Configure Server Wizard

Configure Machine Identity

Configure Server Role

Change the Client Connection Port

Configure Service Stop Timeout

Configure Channel IP Routing

Configure the Default Channel Port

Configure Low Bandwidth Optimization

Configure Public IP Addressing

Management IP Addressing

Considerations for Passive Node Management Using Third Party Technology

Add Remove a Neverfail Engine License Key

Configure the Message Queue Logs

Configure Maximum Disk Usage

Server Protection

[Monitoring the Status of Servers](#)

[Configure Neverfail Engine Settings](#)

[Forcing a Switchover](#)

[Failover versus Switchover](#)

[Split Brain Avoidance](#)

Network Protection

[Configure Public Network Monitoring](#)

[Enabling Automatic Switchover in a WAN](#)

[Setting Max Server Time Difference](#)

Application Protection

[Applications Environment](#)

[Applications Summary](#)

[Applications Services](#)

[Applications Tasks](#)

Data Protection

[Data Replication](#)

Other Administrative Tasks

[Post Installation Configuration](#)

[Business Application Groups](#)

[Configure Event Log Files](#)

[Review Event Logs](#)

[Recloning Secondary or Tertiary Server](#)

Troubleshooting

[Two Active Servers](#)

[Two Passive Servers](#)

[Invalid Neverfail Engine License](#)

[Synchronization Failures](#)

[Channel Drops](#)

[MaxDiskUsage Errors](#)

[Application Slowdown](#)

About This Book

The Administrator's Guide provides information about configuring and performing the day-to-day management of Neverfail Engine (Neverfail Engine) when deployed in a Pair over a Local Area Network (LAN) or Wide Area Network (WAN), or a Trio deployed over both a LAN for High Availability and a WAN for Disaster Recovery. 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 Engine and the actions that Neverfail Engine 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 TM 2008 R2, 2012, 2012 R2, 2016, 2019 and 2022 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 Engine Cluster after successful installation. To help you protect your applications, this guide provides an overview of the protection offered by Neverfail Engine and the actions that Neverfail Engine 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.

Document Feedback

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

Abbreviations Used in Figures

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

Abbreviation	Description
Channel	Neverfail Channel
EMS	Engine Management Service
CE	Neverfail Engine
NIC	Network Interface Card
P2V	Physical to Virtual
V2V	Virtual to Virtual
P2P	Physical to Physical
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 <https://www.neverfail.com/services-and-support/>.

Online and Telephone Support

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

Support Offerings

To find out how Neverfail Support offerings can help meet your business needs, go to <https://www.neverfail.com/services-and-support/>.

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 Engine, Neverfail Professional Services provides offerings to help you optimize and manage your Neverfail

Engine servers. To access information about education classes, certification programs, and consulting services, go to <https://www.neverfail.com/services-and-support/>.

Neverfail Engine Documentation Library

The following documents are included in the Neverfail Engine documentation library:

Document	Purpose
Installation Guide	Provides detailed setup information.
Using Neverfail EMS	Provides detailed usage instructions for Engine Management Service.
Administrator's Guide	Provides detailed configuration and conceptual information.
Deploying to AWS Cloud Environment	Deploying Neverfail Engine in Amazon Web Services Cloud Environment.
SCOPE Data Collector	SCOPE Data Collector Service Overview.
Release Notes	Provides late-breaking information, known issues, and updates. The latest Release Notes can be found at https://www.neverfail.com/services-and-support/ .

Conventions

The documentation uses consistent conventions to help you identify items throughout the printed and online library.

Convention	Specifying
Bold	Window items including buttons.
<i>Italics</i>	Book and CD titles, variable names, new terms, and field names.
Fixed font	File and directory names, commands and code examples, text typed by you.
Straight brackets, as in [value]	Optional command parameters.
Curly braces, as in {value}	Required command parameters.

Convention	Specifying
Logical OR, as in value1 value2	Exclusive command parameters where only one of the options can be specified.

Neverfail Engine Concepts

Neverfail Engine 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 Engine you must understand the basic concepts under which Neverfail Engine operates and the terminology used.

Note: In this document, the term "Cluster" refers to a Neverfail Engine Cluster. Refer to the **Glossary** for more information about Neverfail Engine Clusters.

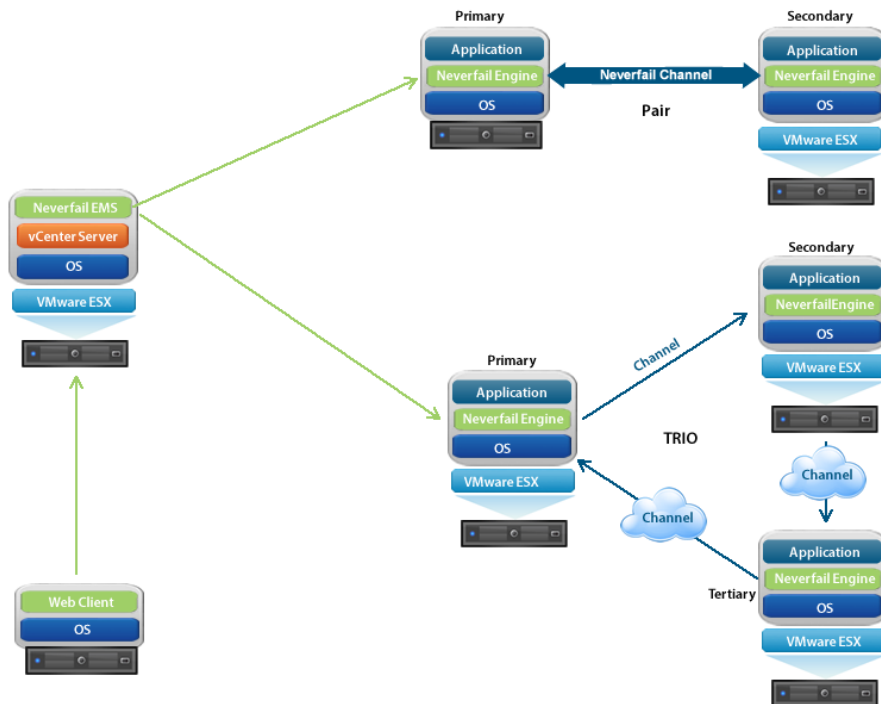
- **Architecture**
- **Protection**
- **Neverfail Engine Networking Configuration**
- **Neverfail Engine Communications**
- **Neverfail Engine Switchover and Failover Processes**
- **Recovery from a Failover**

Architecture

Neverfail Engine provides a flexible solution that can be adapted to meet most business requirements for deployment and management of critical business systems. Capitalizing on VMware vCenter Server's ability to manage virtual infrastructure assets combined with Neverfail's application-aware continuous availability technology, Neverfail Engine brings a best in class solution for protecting critical business systems.

Neverfail Engine consists of the Engine Management Service that is used to deploy and manage the Neverfail Engine service that provides for application-aware continuous availability used for protecting critical business systems.

Using Engine Management Service, users can deploy and manage Neverfail Engine with the ability to view Neverfail Engine status and perform most routine Neverfail Engine operations from a single pane of glass.



Neverfail describes the organization of Neverfail Engine servers based upon Clusters, Cluster status, and relationships between Clusters. Neverfail refers to a Cluster of two servers as a Neverfail Engine Pair or three servers as a Neverfail Engine Trio. Installing Neverfail Engine on the servers and assigning an identity to the servers results in a Neverfail Engine Pair or Trio.

Each server is assigned both an Identity (Primary , Secondary , or Tertiary if installed) 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 (except in the special case described below) whereas the role of the server is subject to change as a result of the operations the server is performing. When Neverfail Engine is deployed on a Pair or Trio of servers, Neverfail Engine 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) or both High Availability and Disaster Recovery.

Note: The identity of an existing Disaster Recovery (DR) Secondary server can change under certain circumstances. This is when a DR pair is extended to become a Trio. In this case, the Secondary will be re-labeled as the Tertiary, so that the Tertiary is always the DR stand-by in any Trio.

In its simplest form, Neverfail Engine operates as a Neverfail Engine 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 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 Engine can assist replication by utilizing WAN Compression with the built-in WAN Acceleration feature.

Additionally, Neverfail Engine can be deployed as a Trio incorporating both High Availability (HA) and Disaster Recovery (DR) or can be extended from an HA or DR pair to a Trio resulting in the following scenarios:

- Primary-Secondary (HA) + Tertiary (DR)
- Primary-Secondary (HA) > extending Pair to Trio resulting in: Primary-Secondary (HA) + Tertiary (DR)
- Primary-Secondary (DR) > extending Pair to Trio resulting in: Primary-Secondary (HA) + Tertiary (DR)

Protection

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

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

A failover occurs when the 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 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:** Engine 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, Engine can gracefully switch the roles of the active and passive servers (referred to as a switchover) allowing the previously 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:** Engine 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, Engine will first try to restart the application on the active server. If a restart of the application fails, then Engine can initiate a switchover.

A switchover gracefully closes down any protected applications that are running on the active server and restarts them on the 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:** Engine 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.

Engine Plug-ins provide these monitoring and preemptive repair capabilities. Engine 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, Engine 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:** Engine ensures the data files that applications or users require in the application environment are made available should a failure occur. Once installed, Engine 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 the passive server. 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 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 passive server, which places all the requests in an array of log files termed the receive queue. The passive server then confirms the changes have been logged by sending the active server an acknowledgment.

The 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 passive server.

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

Neverfail Engine Networking Configuration

The server IP address used by a client to connect to the active server, the Public IP address, must be a static IP address (not DHCP-enabled). In the example below, the Public IP address is configured as 192.168.1.127.

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

Neverfail Engine uses a proprietary filtering system that works with the native Windows Filter Platform to expose a set of Application Programming Interfaces (APIs) to permit, block, modify, and/or secure inbound and outbound traffic while providing enhanced performance over previous versions of the Neverfail Packet Filter Driver.

In a High Availability configuration, the Public NIC on the passive server uses the same IP address as the active server but is prevented from communicating with the live network through a filtering system installed with Neverfail Engine. This filter prevents traffic using the 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 Engine must be configured to use a different 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 Engine uses DNS Update task to update Microsoft Windows 2003, 2003 R2, 2008, 2008 R2, 2012, 2012 R2, and 2016 DNS servers with the new Public IP address. DNS Update runs the `DNSUpdate.exe` to perform the following actions:

- First, `DNSUpdate` 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, `DNSUpdate` 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, DNSUpdate must instruct them to force a replication with the already updated Primary servers.

The NICs on the Primary and Secondary servers intended for use by the Neverfail Channel must be configured so that they use IP addresses outside of the 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 and Secondary servers use the same NetBIOS name. When Neverfail Engine 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 100Mbits per second or more across standard Cat-5 cabling.

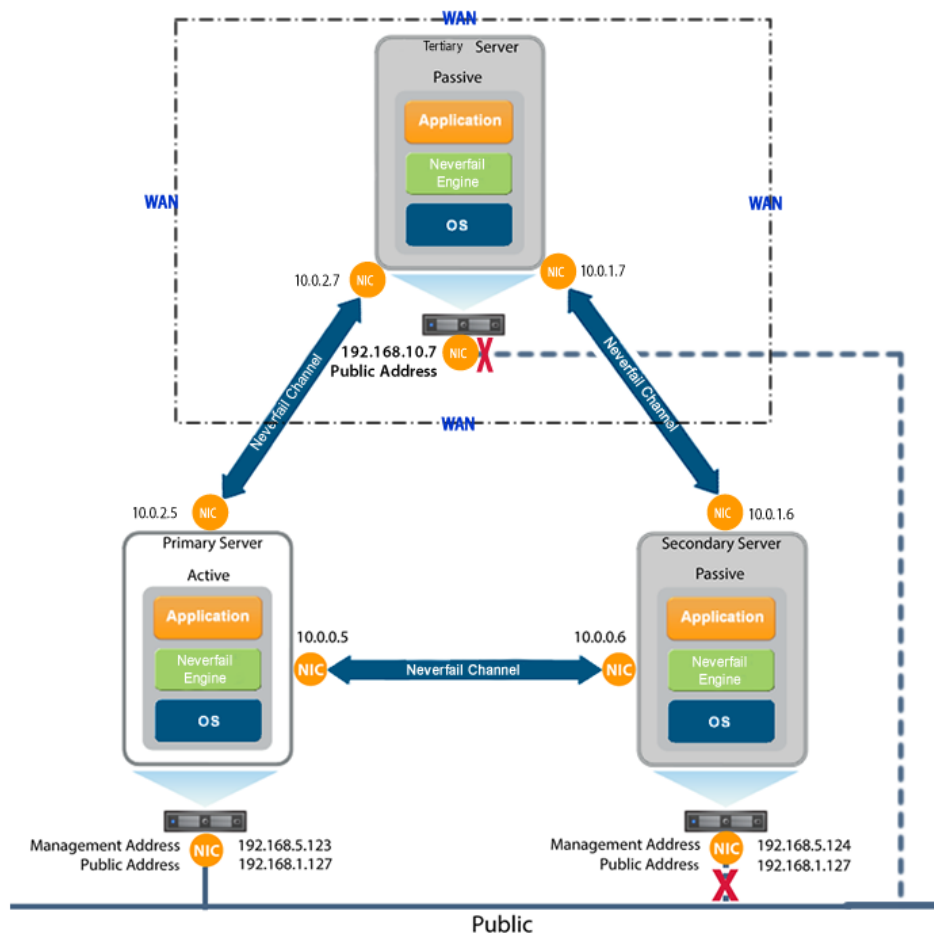
Note: A dedicated channel requires no hubs or routers, but any direct connection requires crossover cabling.

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 Public Network.

Neverfail Engine 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 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 Dualled if more than one dedicated NIC is provided for the Neverfail Channel on each server.

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.



Neverfail Engine Switchover and Failover Processes

Neverfail Engine 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 Engine Pair.

The Managed Switchover Process

A managed switchover can be initiated manually from the Engine Management Service or the Advanced 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.

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 Engine 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.

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 Public IP address. This 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.

1. Correct the incident that caused the failover.
2. Verify the integrity of the disk data on the failed server.
3. Restart the failed server.
4. Neverfail Engine will detect that it has not shut down correctly, and enter a Pending Active mode. In this mode, applications are not started, and the server is not visible on public network.
5. The server will attempt to connect to its peers, to determine if there is an active server. If it connects to its peers, and another server is active, it will connect as passive. If it connects to its peers and no other server is active, it will connect as active. If it doesn't connect with its peers within 2 minutes, it becomes passive.
6. At this stage, the instances of Neverfail Engine running on the servers are connected. Manually start replication which will also resynchronize the data on the passive server.
7. Allow Neverfail Engine 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 Engine Pair (for example, giving the Primary and Secondary the same roles that they had before the failover).
8. Perform a managed switchover.

Managing Neverfail Engine Clusters


Neverfail Engine operates in Clusters of two or three servers with each Cluster administered as a single entity using the Engine Management Service or Neverfail Advanced Management Client. The Neverfail Advanced 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 Engine Management Service or Neverfail Advanced Management Client. The Advanced Client can be downloaded from Engine Management Service UI.

- **Review the Status of Neverfail Engine Clusters and Groups**
- **Exit Neverfail Advanced Management Client**
- **Shutdown Windows with Neverfail Engine Installed**
- **Controlled Shutdown**

Review the Status of Neverfail Engine Clusters and Groups

1. Click on the top level of the Neverfail Advanced 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.



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

2. Click on either:

Document	Purpose
Server	To view the <i>Server: Summary</i> page
Applications	To view the applications status on the <i>Applications: Summary</i> page
Network	To view the network status on the <i>Network Monitoring</i> page
File System	To view the File System status on the <i>Data: Replication</i> page
Registry	To view the Registry status on the <i>Data: Replication</i> page

Exit Neverfail Advanced Management Client

1. Click **Exit** on the File menu.

The *Confirm Exit* message appears.

2. Click **Yes** to close the *Neverfail Advanced Management Client* window or **No** to dismiss the message without exiting the *Neverfail Advanced Management Client*.

Shutdown Windows with Neverfail Engine Installed

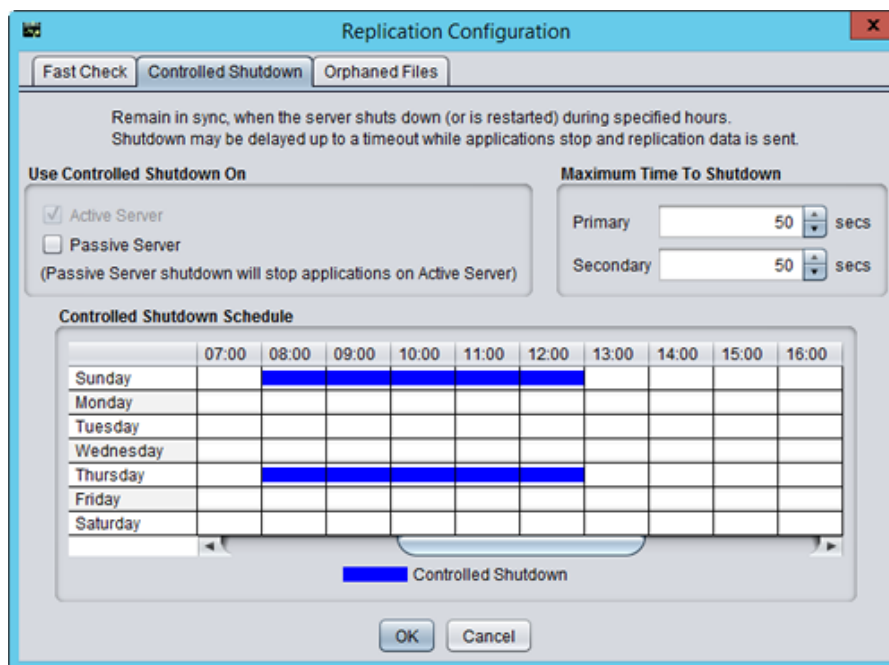
Always stop Neverfail Engine before attempting to shut down Microsoft Windows. If an attempt is made to shut down Windows without stopping Neverfail Engine, Neverfail Engine will not stop in a graceful manner.

Controlled Shutdown

A Controlled Shutdown is a process where the Neverfail Engine 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 Advanced Management Client *Data: Replication* page, this feature allows Neverfail Engine to gracefully shutdown in the absence of the administrator.

1. Navigate to the *Data: Replication* page of the Neverfail Advanced 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.

The ability to configure the length of time for the server(s) to wait for the Controlled Shutdown is configurable starting Windows Server 2008.



7. Click **OK**.

When the *Fast Check* process is enabled in addition to the Controlled Shutdown process, Neverfail Engine 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](#).

Configuring Neverfail Engine

- **Configure Server Wizard**
- **Configure Machine Identity**
- **Configure Server Role**
- **Change the Client Connection Port**
- **Configure Service Stop Timeout**
- **Configure Channel IP Routing**
- **Configure the Default Channel Port**
- **Configure Low Bandwidth Optimization**
- **Configure Public IP Addressing**
- **Management IP Addressing**
- **Considerations for Passive Node Management Using Third Party Technology**
- **Add/Remove a Neverfail Engine License Key**
- **Configure the Message Queue Logs**
- **Configure Maximum Disk Usage**

Configure Server Wizard

Prior to making changes using the Neverfail Engine's Configure Server Wizard, you must stop Neverfail Engine (both Neverfail Engine Service and Neverfail Engine Web Services).

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

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



Configure Machine Identity

Prior to making changes using the *Neverfail Engine's Configure Server Wizard*, you must stop Neverfail Engine (both Neverfail Engine Service and Neverfail Engine Web Services).

The identity of a server (Primary and Secondary) describes the physical hardware of the machine and should not be confused with what the server is doing (the role).

1. 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**.

The screenshot shows the 'Server Configuration' wizard window with the 'Machine' tab selected. The 'Physical Hardware Identity' section is active, showing radio buttons for 'Primary', 'Secondary', and 'Tertiary'. The 'Active Server' section also has radio buttons for 'Primary', 'Secondary', and 'Tertiary'. The 'Client Connection Port' is set to 52267. The 'Service kill timeout during reboots' is set to 5 seconds. The 'Current Role' section is also visible, showing radio buttons for 'Active' and 'Passive'. The 'Finish' and 'Cancel' buttons are at the bottom right.

Server Configuration

Introduction **Machine** Channel Public Management License Logs

Machine Configuration

Physical Hardware Identity

A Neverfail Engine Cluster consists of two or three servers.

The terms *Primary* and *Secondary Tertiary* are used to describe the *Physical Hardware Identity* of each server in the group. A *Server Pair* consists of one *Primary* server and one *Secondary* server. A *Trio* consists of one *Primary*, one *Secondary* and one *Tertiary* server. **Once set, the Physical Hardware Identity of a server should not normally change.**

Current Role

The terms *Active* and *Passive* are used to describe the *Current Role* of each server in the group. A *Server Pair* consists of one *Active* server and one *Passive* server. A *Trio* consists of one *Active* server and two *Passive* servers.

Physical Hardware Identity

Primary Secondary Tertiary

Active Server

Primary Secondary Tertiary

Client Connection Port

52267

Service kill timeout during reboots

During Windows reboots, services are forcibly terminated after X seconds. If the Engine service is killed this way, it records an unclean stop and may not automatically resume the Active role on the next restart.

5 seconds

Set this value to at least 20 seconds (Microsoft allows 5-60 seconds) to ensure a graceful Stop of Engine during any system reboot.

Reset << Back Next >>

Finish Cancel

Configure Server Role

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

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

1. 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 Engine Management Service) connect to Neverfail Engine.

1. 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.

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

Configure Service Stop Timeout

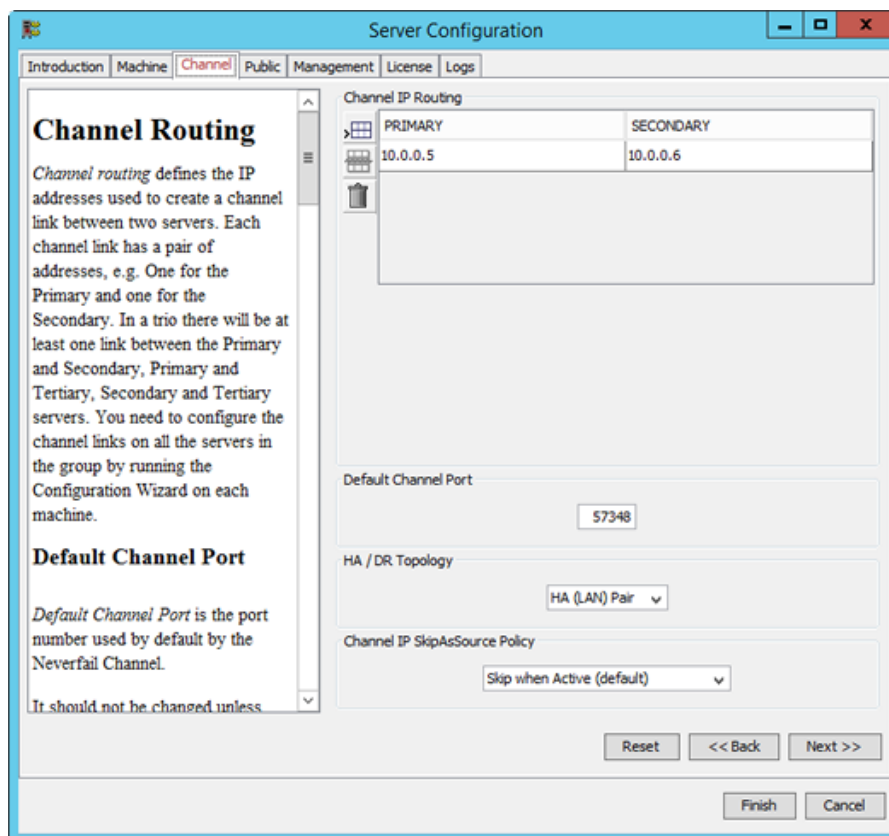
The Service kill timeout during reboot specifies the time Windows OS waits till terminating the service.

1. To configure the *Service Stop Timeout*, select the **Machine** tab of the *Configure Server Wizard* and type a new value (in seconds) in the **Service Stop Timeout** text box. Click **Next** or **Finish** to accept changes.

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). Each link uses two addresses, one for the local server and one for the remote server.

1. 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



2. Additionally, you can specify a SkipAsSource policy for channel addresses to ensure that they are not used for public traffic. SkipAsSource prevents an IP address from being selected by the operating system as a source IP address for out-going network connections.
 - Never Skip - channel IP addresses will never have SkipAsSource set.
 - Always Skip - channel IP addresses will always have SkipAsSource set.

- Skip when Active - channel IP addresses will have SkipAsSource set when the server is active but not when passive.
 - Skip when Active and Public Subnet - channel IP addresses will have SkipAsSource set if the server is active and the channel IP address is in the same subnet as a public IP address. When the server is passive the SkipAsSource setting is removed from the channel IP addresses.
3. 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 servers. Do not change this port unless required by another application.

1. 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.

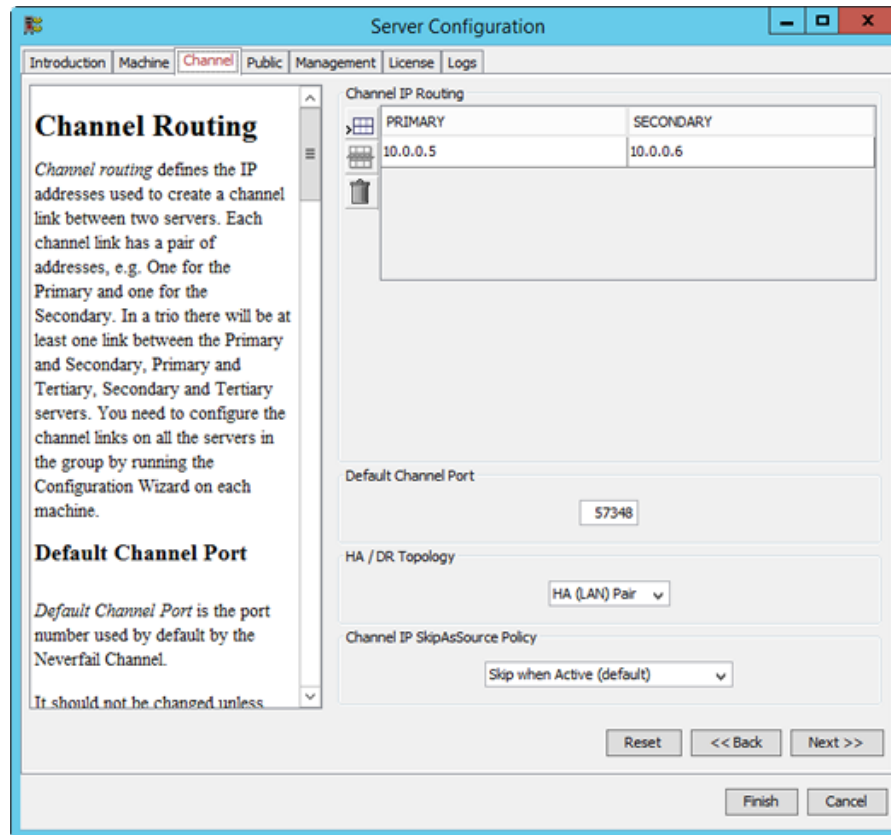
Configure Low Bandwidth Optimization

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

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 server pair you can choose HA or 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.

1. 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.

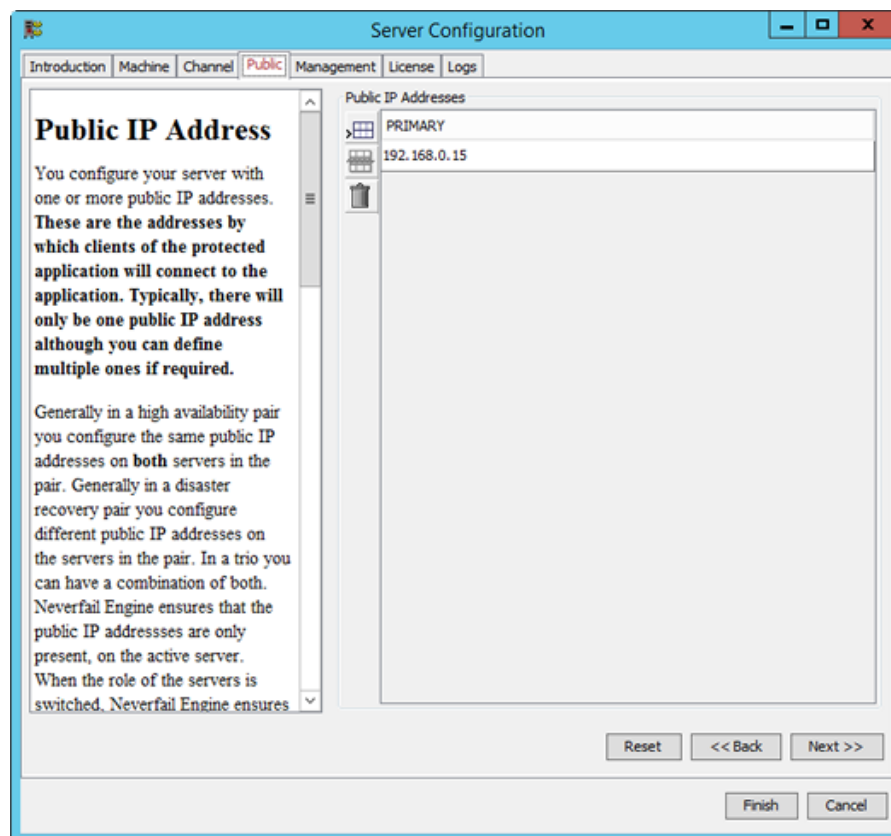


Configure Public IP Addressing

A typically configured Neverfail Engine Cluster uses only one Public IP address when deployed as a pair or on a LAN, but can be configured with more than one Public IP address. These are the addresses by which clients of the protected application connect to the application. Typical installations configure the same Public IP address on the Primary and Secondary servers. All traffic to and from these 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 Public IP addresses on whichever server is currently active. When the Neverfail Engine service is shut down, the filtering remains in place to prevent IP address conflicts between servers.

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

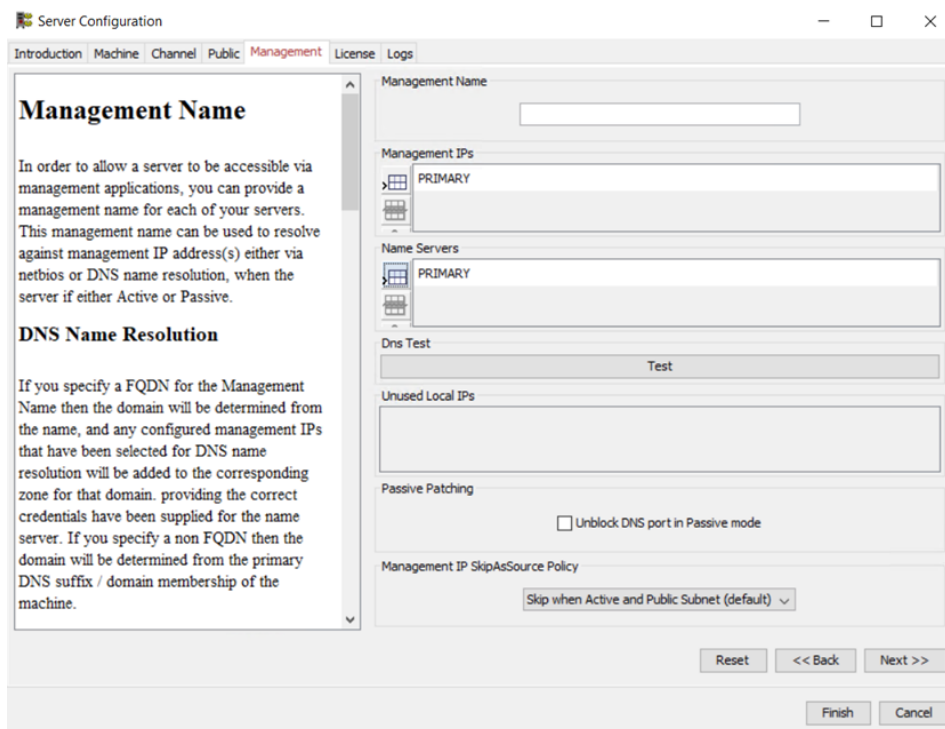
An address must not appear more than once, and no Public IP address may appear in the list of IP addresses on the Channel tab.



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.

Management IP Addressing

The Management page of the Server Configuration Wizard allows you to set up management access for the configured server. This can be done by assigning a management name, IPs and name servers.



The **Management Name** is the name of the machine (used for management purposes only) when the server is in the Passive role. For example, this machine name can be used for applying updates to the operating system. This name can be declared using NetBios or FQDN formats, depending on the configured management IPs.

The management name is resolved to the configured management IP addresses and can be accessed via DNS or NetBios, when the server is in the Passive role (if the server is in the Active role, the machine name is always the cluster name).

The **Name Servers** option allows you to either specify an explicit IP or name for the name server, or set it to Auto. When this option is set to Auto, the name server(s) are deduced from the server's domain membership.

Each name server can be defined as dynamic, by either using the machine account or by specifying a different account, or static, using appropriate credentials.

Management IP addresses are additional IP addresses that you manually configure on a server; they are IP addresses that are neither public or channel IP addresses. Management IP addresses are typically used to access a server for management purposes and can be used to access a server when it is passive. Management IP addresses are displayed here so that you can see the management IP addresses on your local server.

The **DNS Test** button allows you to test the adding/checking/removing of DNS entries to the DNS server.

Additionally, you can specify a SkipAsSource policy for Management IP addresses to ensure that they are not used for public traffic. SkipAsSource prevents an IP address from being selected by the operating system as a source IP address for out-going network connections.

The following options are available:

- Never Skip - channel IP addresses will never have SkipAsSource set.
- Always Skip - channel IP addresses will always have SkipAsSource set.
- Skip when Active - channel IP addresses will have SkipAsSource set when the server is active but not when passive.
- Skip when Active and Public Subnet - channel IP addresses will have SkipAsSource set if the server is active and the channel IP address is in the same subnet as a public IP address. When the server is passive the SkipAsSource setting is removed from the channel IP addresses.

Considerations for Passive Node Management Using Third Party Technology

The Engine cluster's passive nodes can be managed (e.g. updated) using third party tools like SCCM, Windows Server Update Services (WSUS) or Ivanti Patch. Each tool requires specific configuration, as described next. Management names and IPs must be defined for **all nodes** in the cluster.

SCCM 2012 R2

Read the following knowledge base article to learn how to deploy updates to passive servers using SCCM 2012 R2: **[Managing And Patching Neverfail Engine Clusters Using Cozen Passive Node Management Feature With System Center Configuration Manager \(SCCM\) 2012 R2.](#)**

Windows Server Update Services (WSUS)

- Make sure that SkipAsSource is disabled if the management IP address is in the same subnet as the public IP address.
- Configure the Group Policy's intranet update service to use the IP address of the WSUS server.

Ivanti Patch

Add the public name and all management names. Ivanti Patch will scan all names and ignore the management name of the active server.

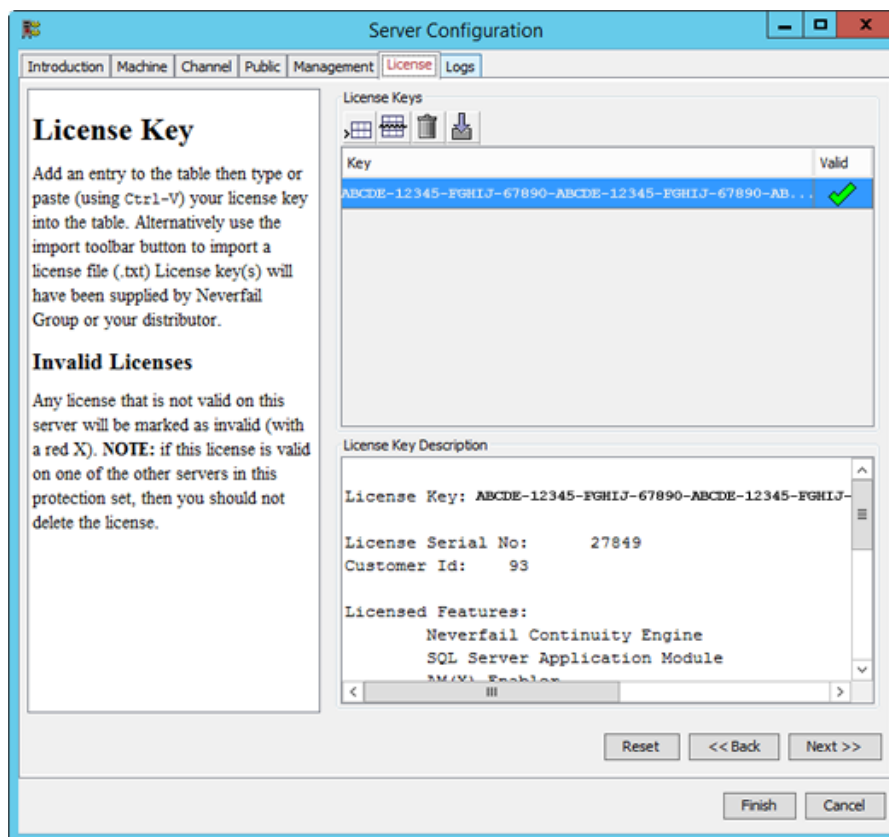
Note: Learn more about Engine's Passive Node Management use cases by reading the following article: **[When to Use Neverfail Patch Management Options.](#)**

Add Remove a Neverfail Engine License Key

Neverfail recommends using the Engine Management Service user interface for licensing Neverfail Engine (see the Installation Guide).

If requested by Neverfail Support, you can also use the Configure Server Wizard as follows:

1. To manage Neverfail Engine License Keys, select the **License** tab of the *Configure Server Wizard*.
2. To add an entry to the *License Keystable*, 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.

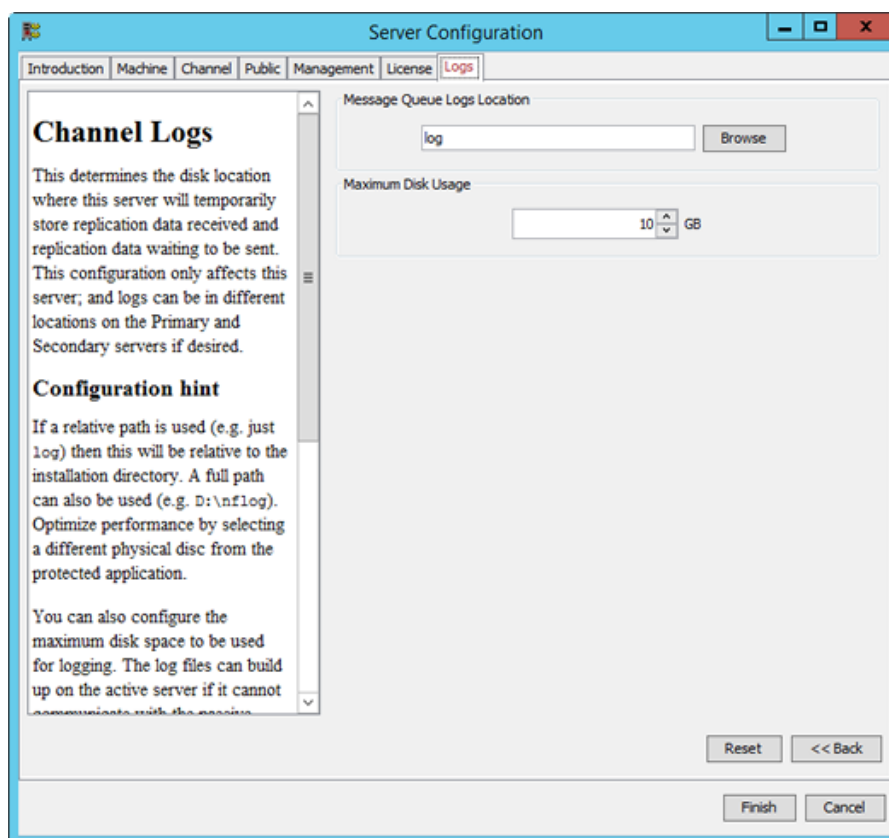


3. After entering your license keys click **Next** or **Finish**.

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 and Secondary servers.

1. 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.



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.

1. 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 Engine may run out of physical disk space, and when this happens, a shutdown and restart may be required before replication can resume.

2. Neverfail recommends a *Maximum Disk Usage* setting that leaves a little overflow space to enable Neverfail Engine 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.

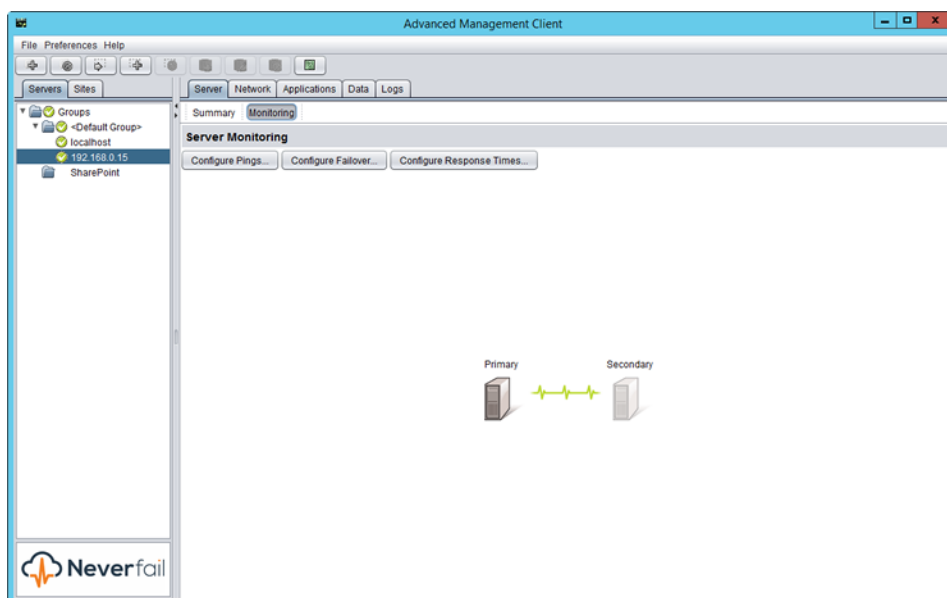
Server Protection

Protection against operating system or hardware failure affecting the active server is facilitated by multiple instances of Neverfail Engine 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**
- **Configure Neverfail Engine Settings**
- **Forcing a Switchover**
- **Failover versus Switchover**
- **Split-brain Avoidance**

Monitoring the Status of Servers

The Neverfail Advanced Management Client **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 Engine 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 Engine also sends a ping to ensure that the servers remain visible to one another.



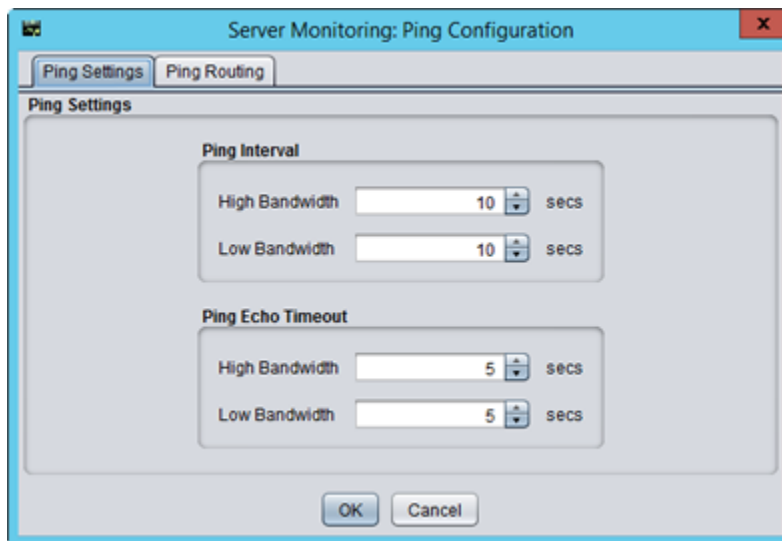
Configure Neverfail Engine Settings

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

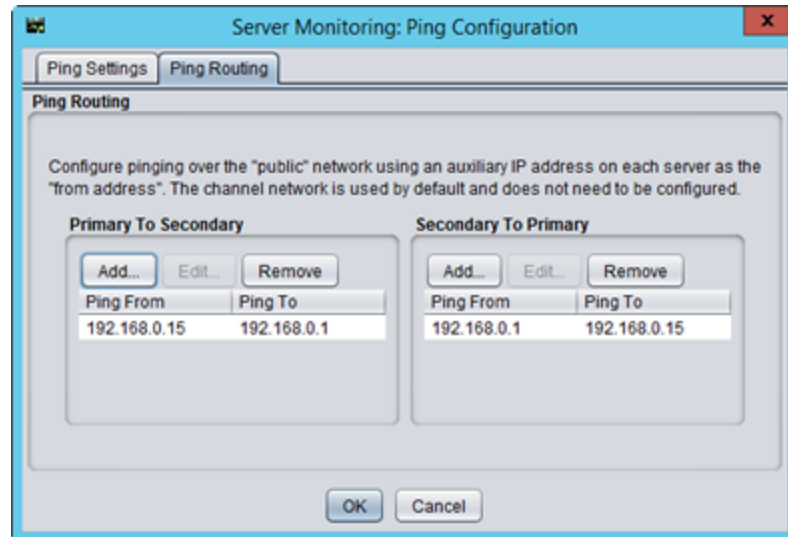
Configure Pings

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 Engine to send pings across the Neverfail Channel and the Public Network in addition to the heartbeat ("I am alive" messages) to confirm that the server is still operational and providing service.

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



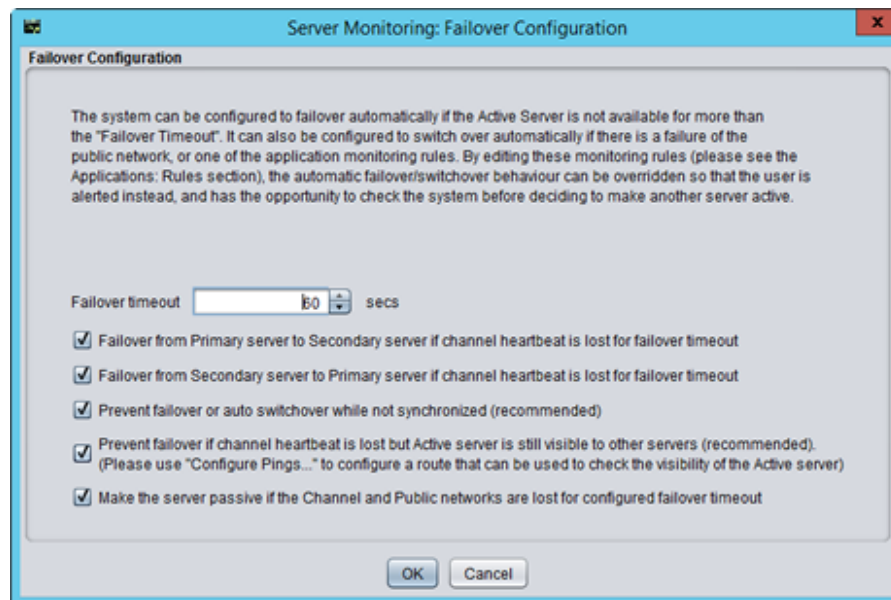
2. Select the **Ping Routing** tab and enter the auxiliary IP addresses of the appropriate servers.



Configure Failover

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

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



2. Type a new numeric value (seconds) in the *Failover timeout* text box or use the arrow buttons to set a new value.
3. 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.

Configure Response Times

Neverfail Engine also allows you to configure channel connection timeouts.

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

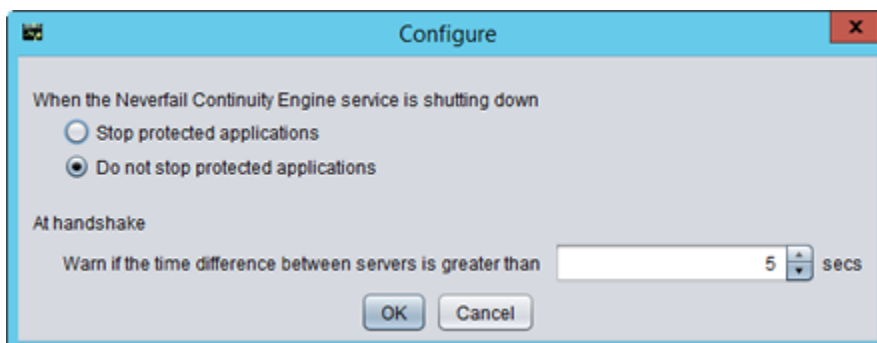


2. Type new numeric values (second) into the text boxes or use the arrow buttons to select new values.
3. Click **OK** to accept the changes or **Cancel** to dismiss the dialog without making any changes.

Common Administrative Tasks in Neverfail Engine

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

1. Click to open the *Configure* dialog.



2. Select the radio button corresponding to whether you want to stop or leave running the protected applications when Neverfail Engine 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.
3. 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.
4. Click **OK** to accept the changes or **Cancel** to dismiss the dialog without making any changes.

Forcing a Switchover

After Neverfail Engine is configured to protect all required applications and data, it allows the Secondary 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 re-booting 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.

Note: 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 Advanced Management Client.

Failover versus Switchover

Important: Do not confuse a failover with a switchover.

A switchover is a controlled switch (initiated from the Engine Management Service, Neverfail Advanced Management Client, or automatically by Neverfail Engine 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 Engine 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.

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 Engine 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*s exceeded.

When the configured *Failover timeout* value has elapsed, Neverfail Engine will evaluate, in order, the following pre-configured rules before taking action:

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

- Failover from Primary server to Secondary server if channel heartbeat is lost for fail-over timeout
- Failover from Secondary server to Primary server if channel heartbeat is lost for fail-over 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 Public network cards of each server to allow the passive server to send a ping via the Public network. Management IP addresses are additional IP addresses assigned to the network card connected to the Public network. They are used to allow the passive server to communicate, because unlike the Public IP address, they are not filtered. For information about how to configure Management IP addresses, see **Management IP Addressing**.

4. Click **OK**.

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

Typical Failover and Active Server Isolation Scenarios

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

The following scenario assumes the active server has failed and is no longer available. Upon detection of missed heartbeats, Neverfail Engine on the passive server performs the following steps:

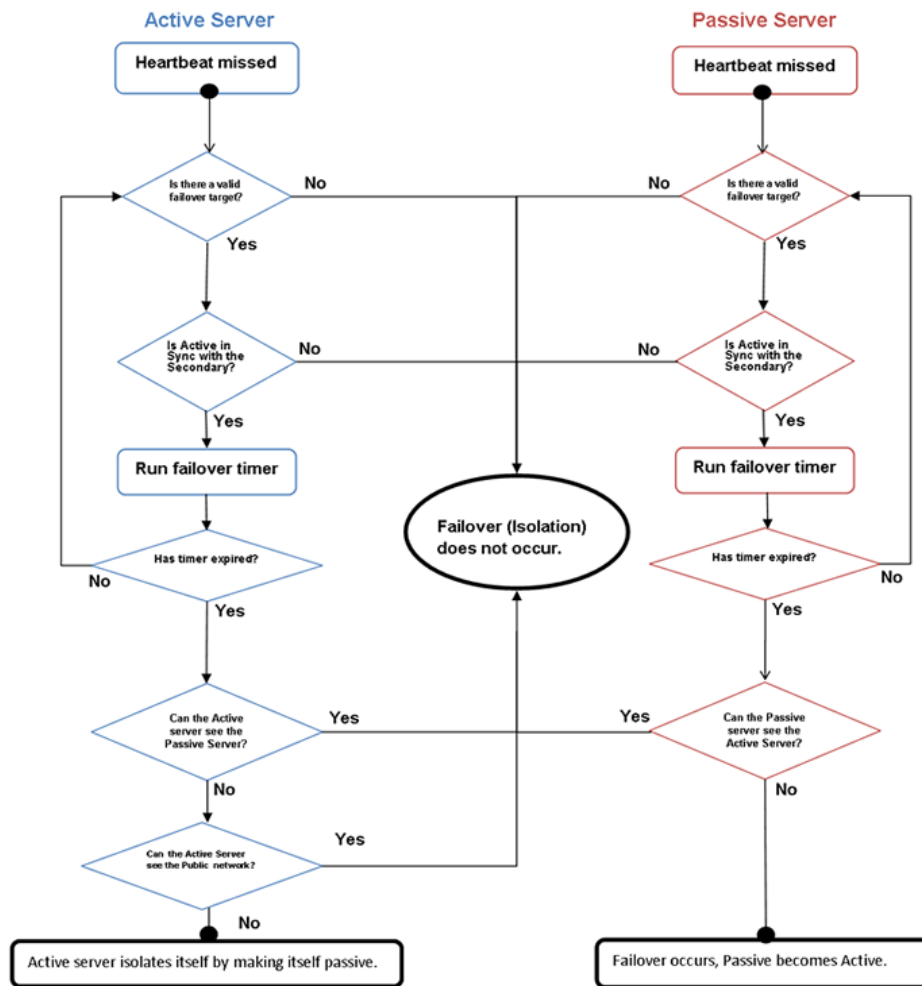
1. As soon as the passive server detects that the Neverfail Channel is experiencing missed heartbeats, it will determine if itself is a valid failover target to the currently active server.
2. As soon as the passive server detects that the Neverfail Channel is experiencing missed heartbeats, it will 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.

Note: The following scenario assumes that Neverfail Engine is deployed in a LAN with all rules selected in the **Server: Monitoring > Configure Failover > Failover Configuration** dialog.

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



Upon detection of missed heartbeats Neverfail Engine 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.

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 Engine in a configuration with the Primary server as active and the Secondary server as passive.

Note: When failover conditions, such as a power failure, cause failures in both active and passive servers, a condition may result that causes all servers to restart in Passive mode. In this situation, manual intervention is required. See **Two Passive Servers** 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 Engine 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 Engine 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 Engine.
 - **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 Advanced 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:
 1. If the Neverfail Advanced Management Client is unable to connect remotely, try running it locally. If you remain unable to connect locally then verify that the Neverfail service is running via the Service Control Manager. If it is not, review the event logs to determine a cause.
 2. 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 Engine is running, you can run the *Configure Server Wizard*, but you will not be able to make any changes. You must stop the Neverfail Engine service before attempting to make changes via the *Configure Server Wizard*.

3. Determine whether the protected application is accessible from clients. If it is, then start Neverfail Engine 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 Engine 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 Engine 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 Engine on the failed Primary server and then Start Replication and allow the system to synchronize. After a failover, replication does not start automatically giving you the opportunity to recover any lost information from the failed active server before you manually start replication from the new active server. 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 Public network. Split-brain Avoidance works by pinging from the passive server to the active server across the 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 Public network for the local and remote servers.

1. To enable Split-brain Avoidance, open the *Server Monitoring* page in the Neverfail Advanced Management Client.
2. Click **Configure Split-brain Avoidance** and go to *Split-brain Avoidance* tab. Edit ping routes between Primary and Secondary instances using existing auxiliary or management IP addresses. Split-brain avoidance works by pinging from the passive server to the active server across the client network. If the active server responds, the passive does not failover, even if the channel connection is lost.
3. Click **Configure Failover**.
4. 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.

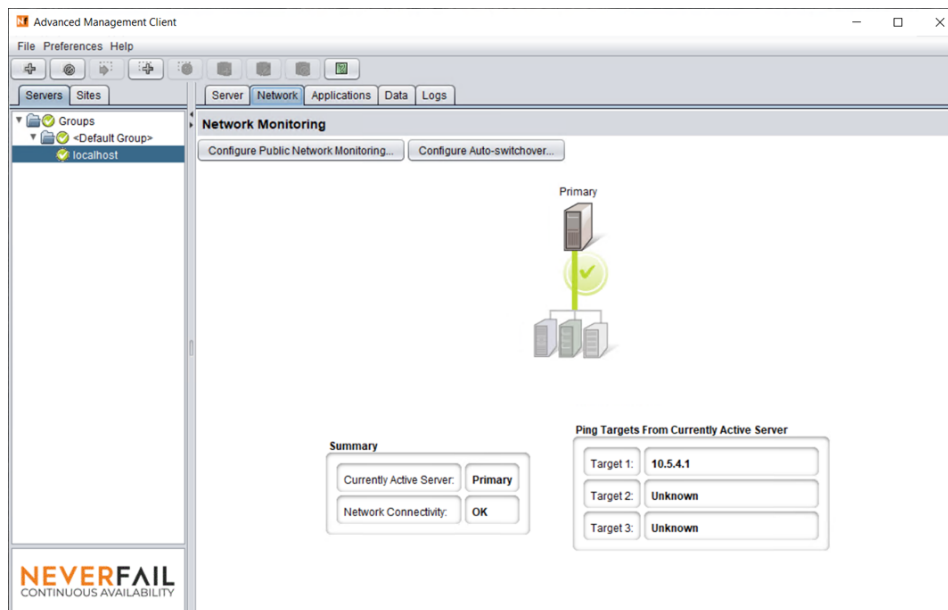
Note: You must configure Management IP addresses on the Public network cards of each server to allow the passive server to send a ping. Management IP addresses are additional IP addresses assigned to the network card connected to the Public network.

Additionally, the Passive Server can be configured to avoid false failover when it gets isolated from the active server and the public network, by configuring it with a Management IP address so it can ping the configured public network targets: this additional setting will avoid split-brain in situations where passive server fails-over after losing connection to active server and public network, followed by network connections recovery (original active server still remains active, hence split-brain will happen after the network reconnection occurs). The Management IP address can be added using the Configure Server Wizard.

Network Protection

Neverfail Engine 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 Engine can initiate a switchover, allowing the passive server to assume an identical network identity as the active server.

The Neverfail Advanced Management Client **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.



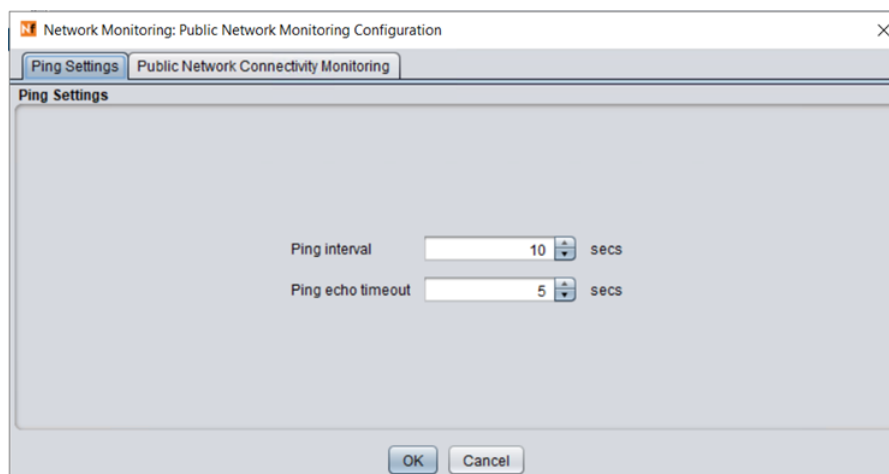
- **Configure Public Network Monitoring**
- **Enabling Automatic Switchover in a WAN**
- **Setting Max Server Time Difference**

Configure Public Network Monitoring

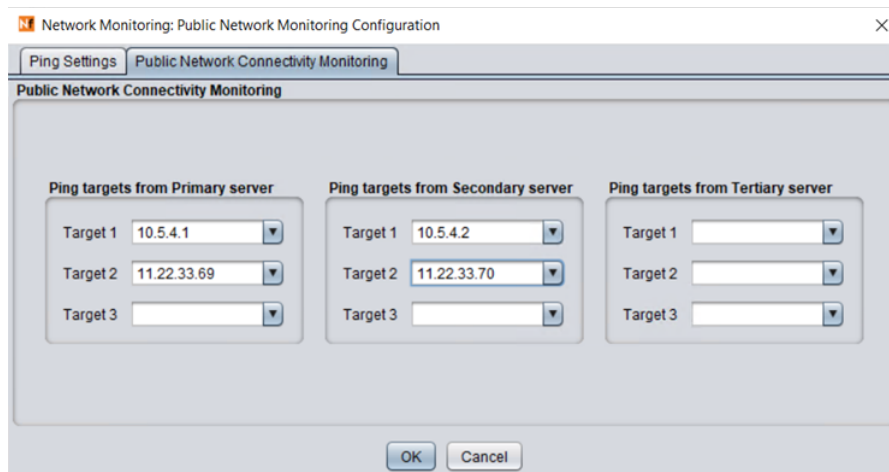
The Public network monitoring feature, previously discussed, is enabled by default during the installation of Neverfail Engine. This feature integrates the polling of the particular waypoints around the network through the active server's Public connection to ensure connectivity with the Public network is operational. By default, the IP addresses of the default gateway, the primary DNS server, and the Global Catalog server are all selected. When one or more of the automatically discovered waypoints are co-located on a physical machine (leading to duplication of IP addresses), the ability to specify additional waypoints manually becomes an advantage.

To configure Public Network Monitoring:

1. To specify a manual target for the Public network checking, click **Configure Public Network Monitoring**.



2. Select the **Public Network Connectivity Monitoring** tab to add to or modify the existing target IP addresses for each server to ping.



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

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 Engine 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:

1. In the Neverfail Advanced 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**.

WAN Auto-Switchover Configuration



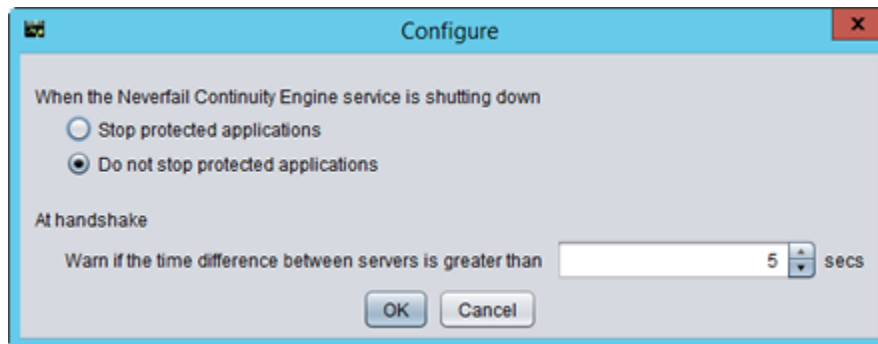
Setting Max Server Time Difference

Neverfail Engine 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.

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**.

Server: Summary Configure dialog



Application Protection

- **Applications Environment**
- **Applications Summary**
- **Applications Services**
- **Applications Tasks**

Applications Environment

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

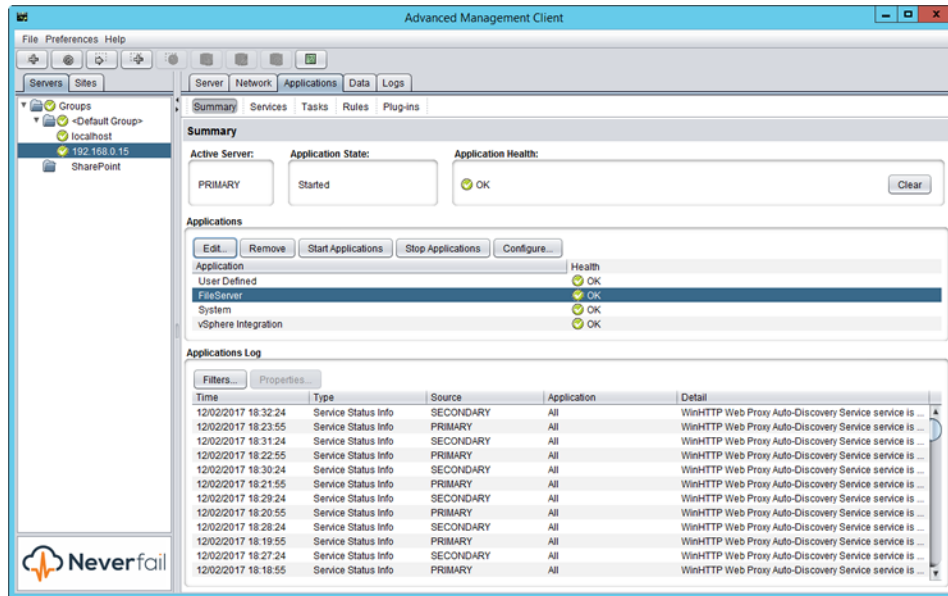
The AMFx provides additional functions while maintaining the traditional stability of Neverfail software. Use the AMFx to install and remove plug-ins on the fly while Neverfail Engine continues to provide protection to currently installed applications.

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 Engine 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 Neverfail Advanced Management Client **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.



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

View Application Status

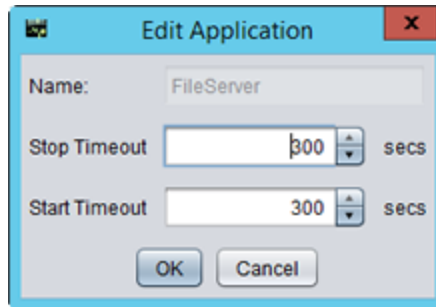
After an application starts and is running you can view its status in the *Applications* pane of the **Applications: Summary** page.

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.

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.



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 Engine.

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.

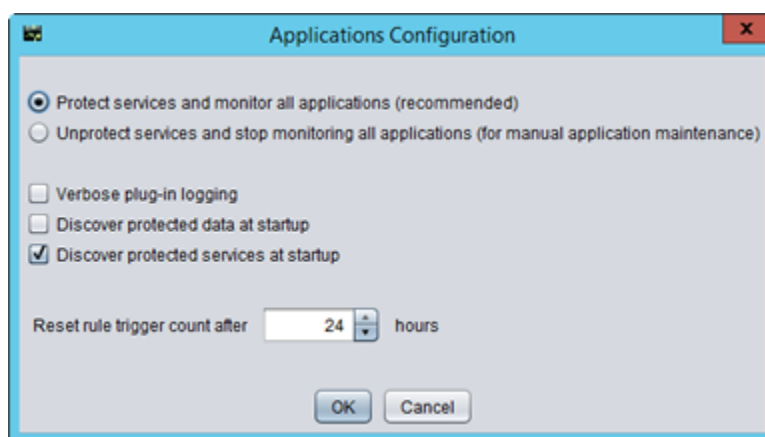
3. Click **Yes** to remove the selected application, or click **No** to dismiss the message without deleting the application.

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 Engine or taking the whole server offline. During installation, Neverfail Engine creates default settings for application configurations. The Neverfail Advanced Management Client **Applications: Summary** page allows you to change the settings.

To configure applications:

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



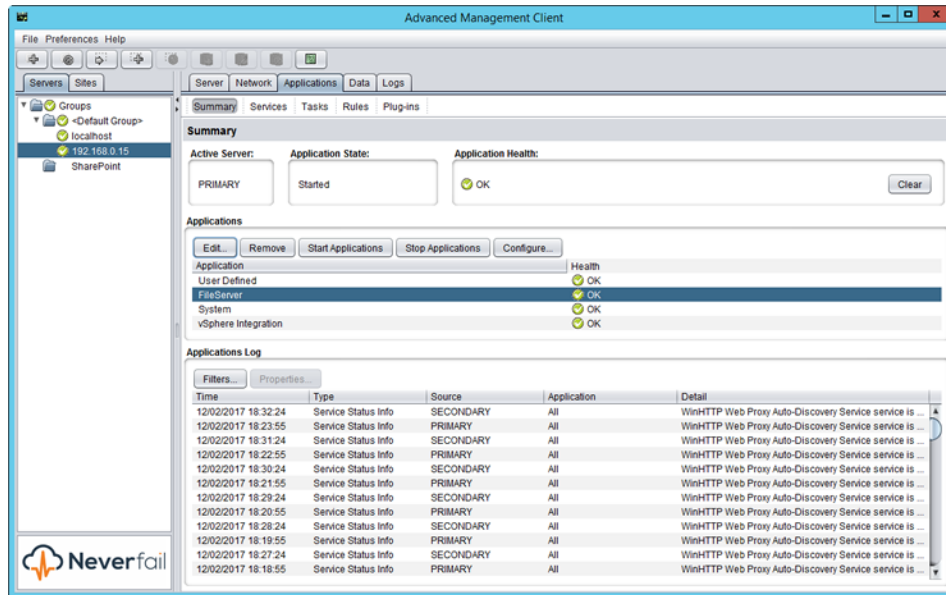
2. Select **Protect services and monitor all applications (recommended)** or **Unprotect services and stop monitoring all applications (for manual application maintenance)**.

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.



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.

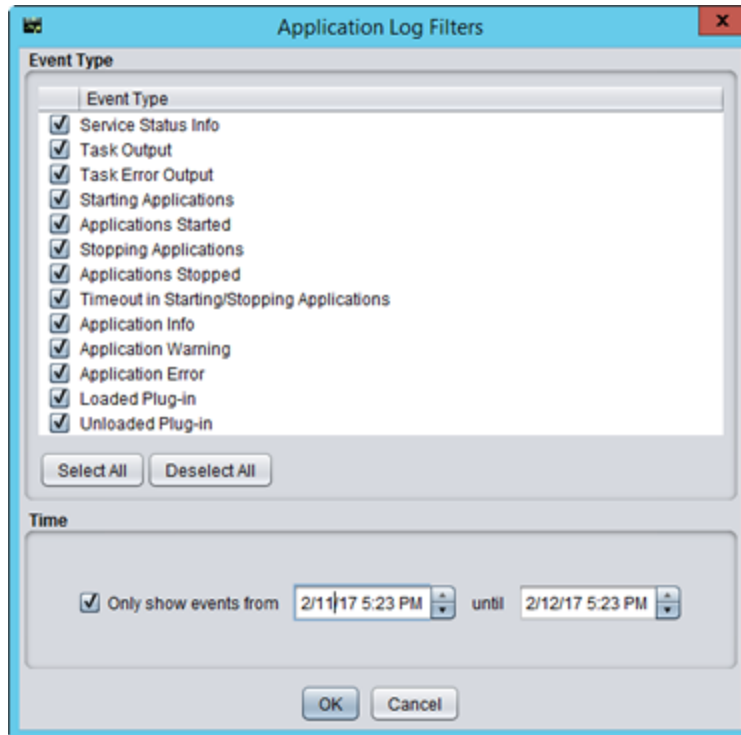
You also can filter *Applications Log* entries to reduce the number of events displayed, and use the *Applications Log* to troubleshoot application errors. For example, if an application fails, you can right-click on the associated event in the *Application Logs* and select **Properties** to open the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

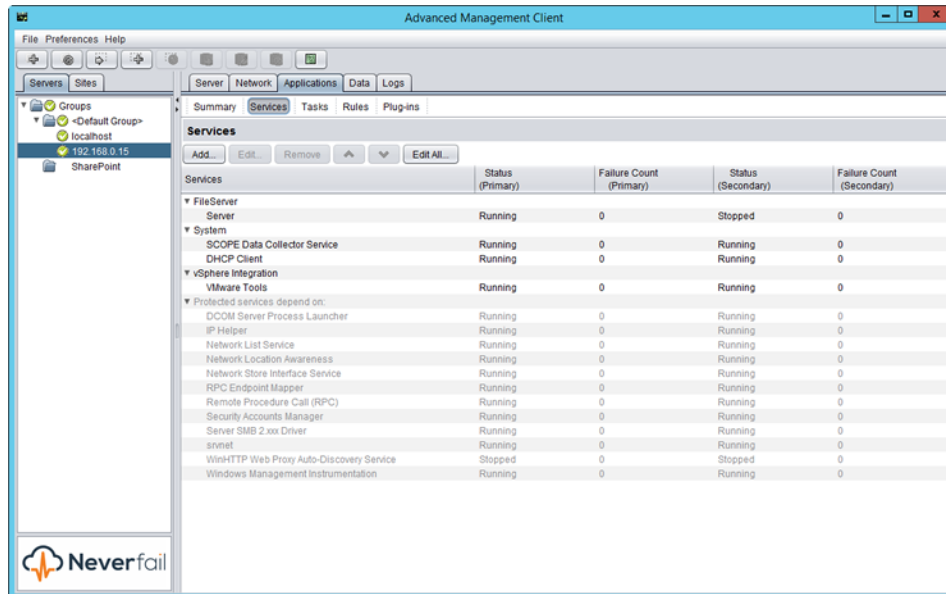
The *Application Log Filters* dialog appears.



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

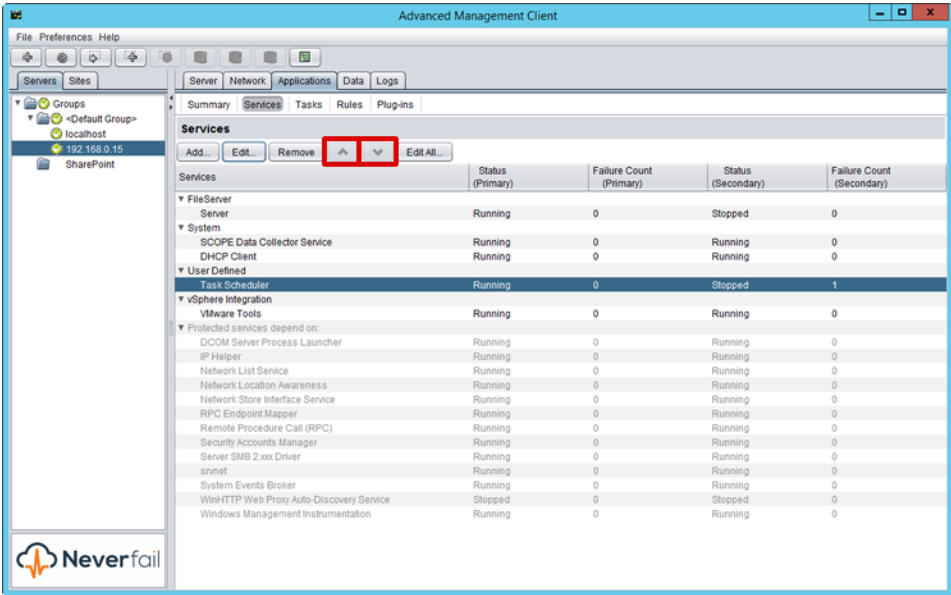
The Neverfail Advanced Management Client **Applications: Services** page shows services specified by plug-ins or by the user, and any services related by dependency.



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:

- 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

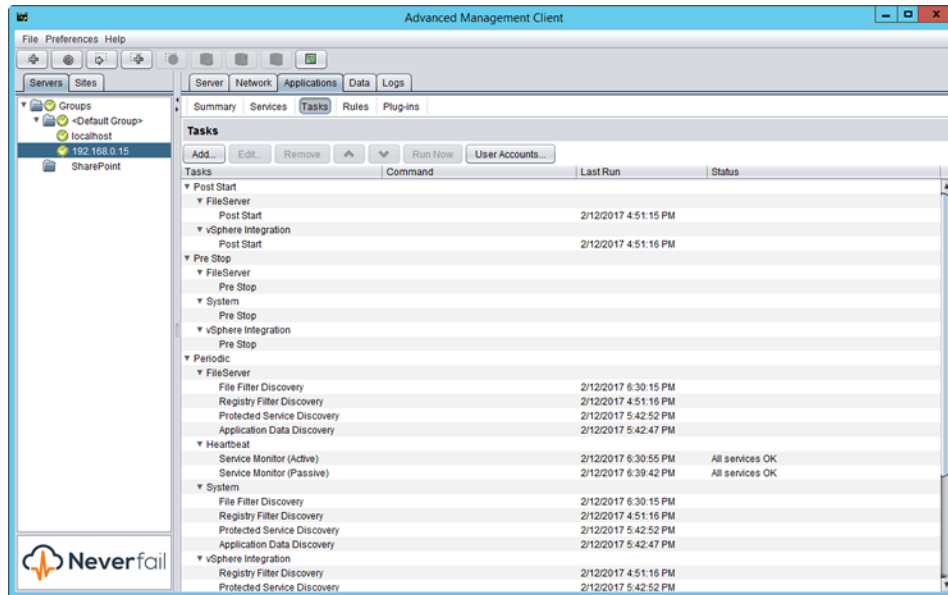
Tasks are a generalization and extension of the start, stop, and monitor scripts in earlier versions of this product.

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 and Secondary 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 (not available in this version).
- **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 Engine. 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 Advanced Management Client **Applications: Tasks** page provides a list of tasks and associated status information, as well as features to quickly manage tasks.



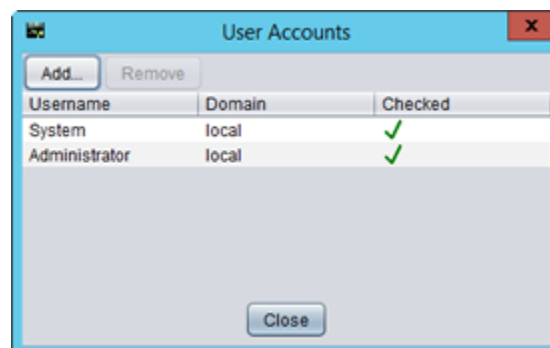
Change the Order of Tasks

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.

View, Add, and Remove User Accounts

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

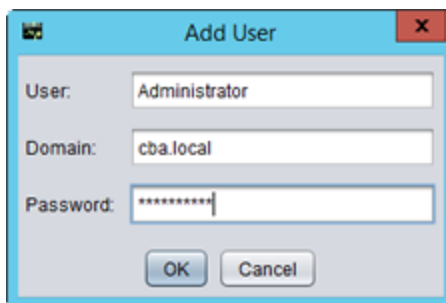
Click **User Accounts** (near the top of the **Applications: Tasks** page). The **User Accounts** dialog appears.



- To add a user account:

1. Click **Add**.

The *Add User* dialog appears.



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.

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

Data Protection

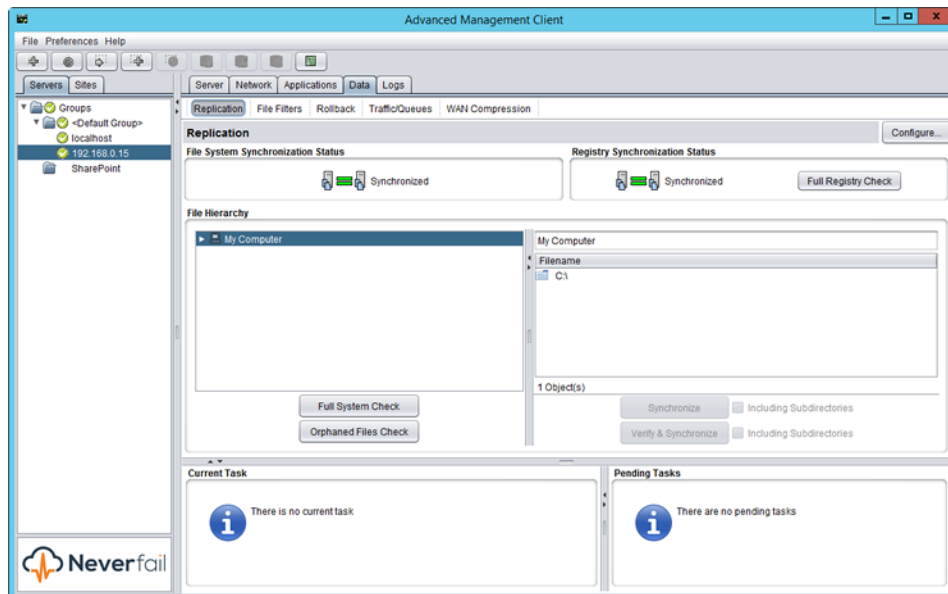
- **Data Replication**

Data Replication

Neverfail Engine 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 Engine program folder holds the send and receive queues on the active and 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.

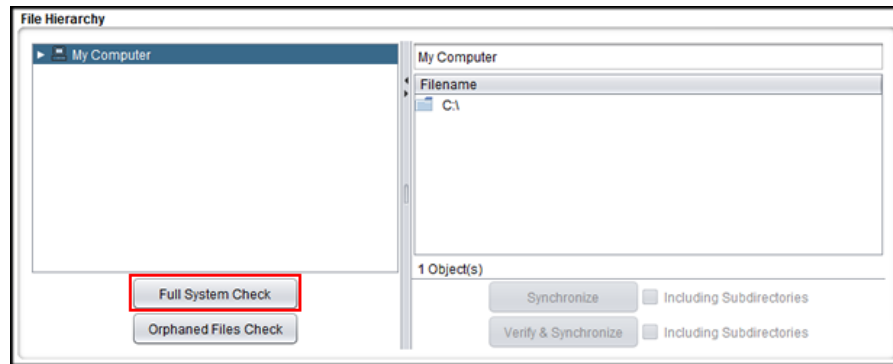


Initiate a Full System Check

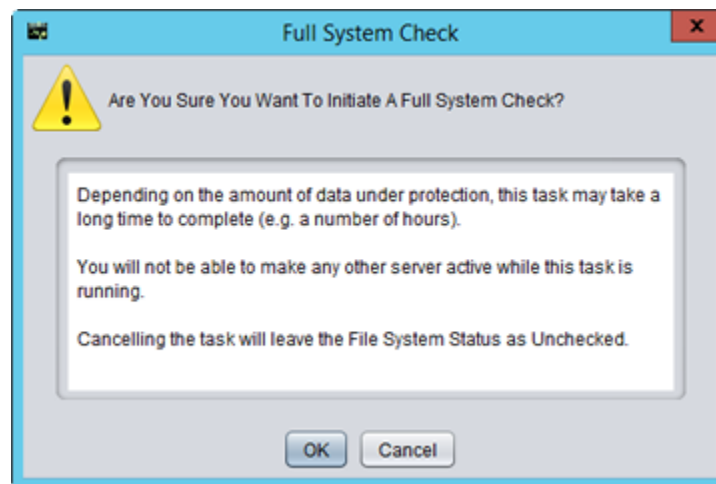
Certain system events, such as preceding a switchover or following a failover or split-brain syndrome, may require running a full system check to ensure that the entire protected file set is synchronized and verified. A full system check performs a block-level check identical to that performed during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

To initiate a full system check:

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



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

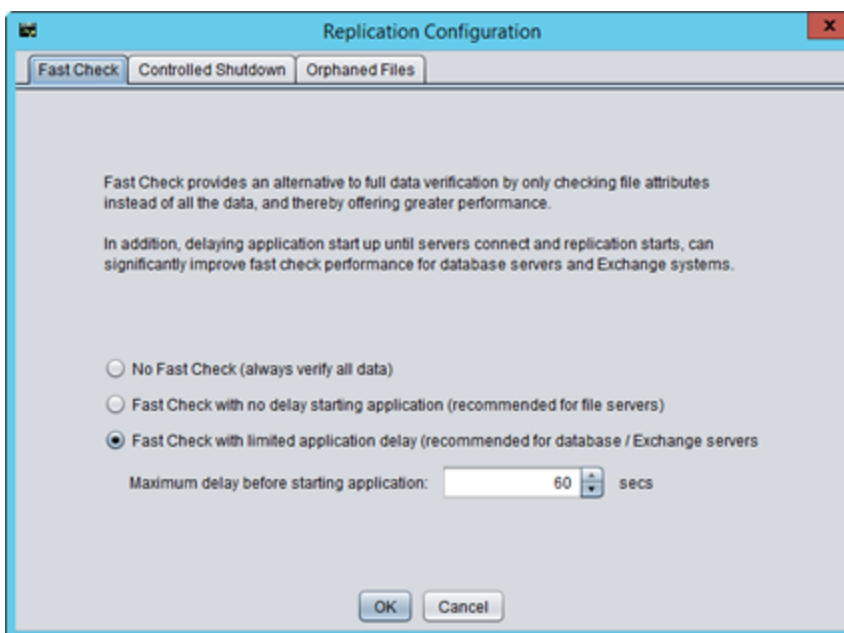
The Fast Check process is used by Neverfail Engine to rapidly verify files between servers prior to starting applications. Fast Check compares file time stamps and attributes rather than the check sums of the data thereby accelerating the startup and synchronization process. If the time stamp or attribute check fails, then the normal verification and synchronization process will initi-

ate. 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.

When combined with Controlled Shutdown, Fast Check provides the ability to perform scheduled unattended restarts of the servers. 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 Engine will delay the startup of the application while it attempts to establish replication between active and all passive nodes.
6. Click **OK**.



Note: When Fast Check is configured in addition to Controlled Shutdown, Neverfail Engine can be configured to perform an unattended restart. For more information about Controlled Shutdown, see **Controlled Shutdown**.

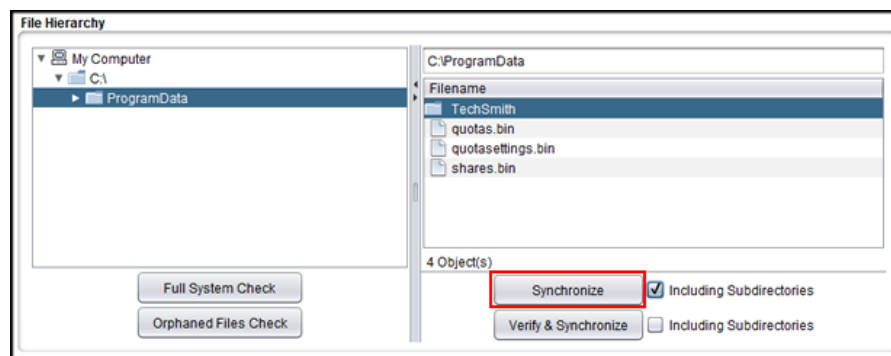
Manually Initiate File Synchronization

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.

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.



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**.

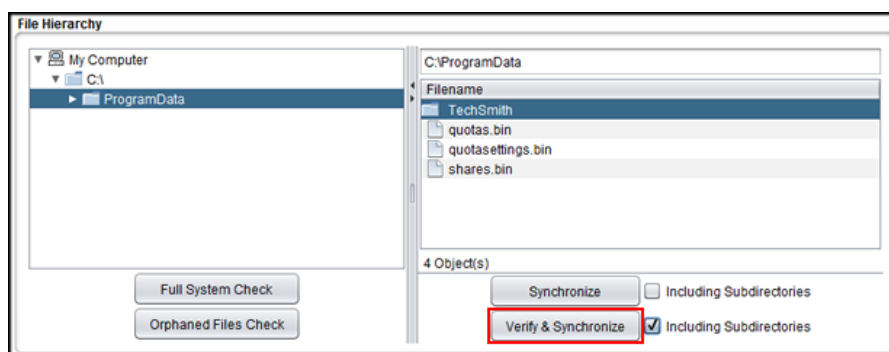
To manually verify and synchronize:

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.



Orphaned Files Management

Neverfail Engine 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.

Orphaned File Check can either delete or log files on the passive server that exist within the protected set; they were "orphaned" because Neverfail Engine 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

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

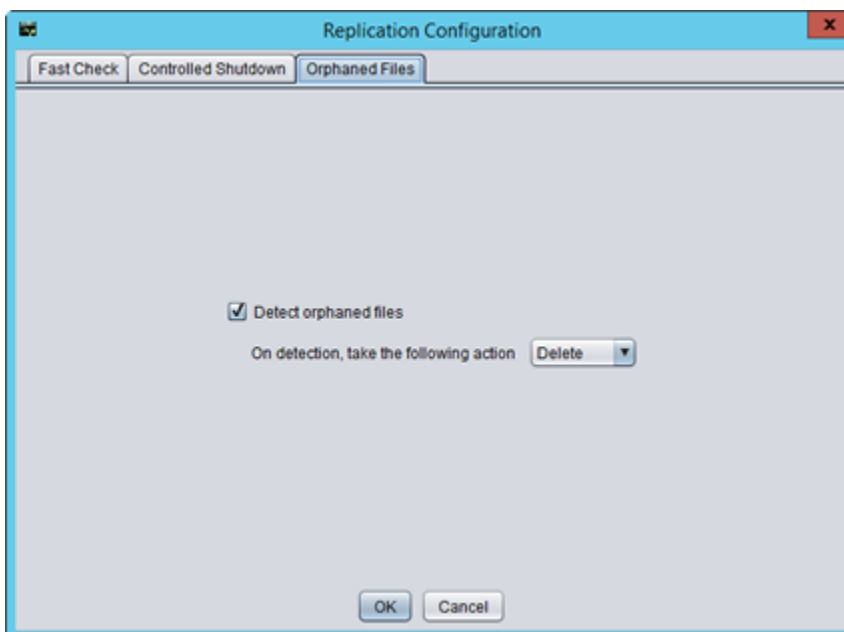
orphaned files. Should you want to log the files presence, see **Configure Orphaned Files Check**.

Configure Orphaned Files Check

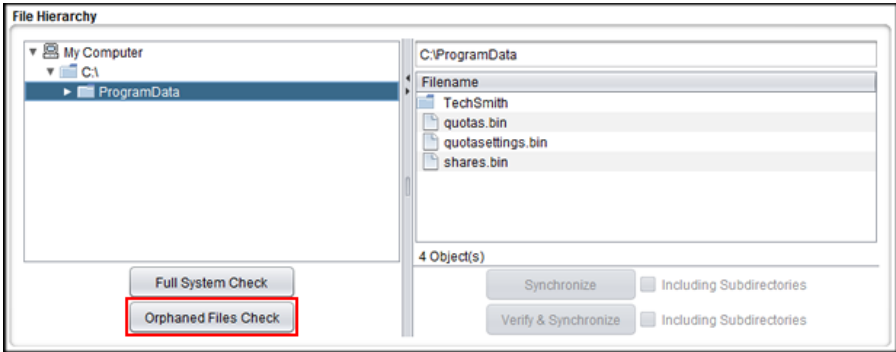
Prior to initiating an orphaned files check, you must configure the options for actions to take in the event orphaned files are found. By default, Orphaned Files Check is configured to delete orphaned files. Should you want to log the files presence, follow the steps below.

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.



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



Other Administrative Tasks

- **Post Installation Configuration**
- **Business Application Groups**
- **Configure Event Log Files**
- **Review Event Logs**
- **Recloning Secondary or Tertiary Server**

Post Installation Configuration

Upon completion of installation of Neverfail Engine, you should perform the following Post Installation tasks:

- **Configure the VmAdapter Plug-in**
- **Adding an Additional Network Interface Card**

Configure the VmAdapter Plug-in

After installation of Engine is complete, configure the VmAdapter Plug-in:

1. Launch the Engine Management Service UI, log in and select the protected server.
2. Navigate to **Server Details > Applications and Platforms**.
3. Locate the **vSphere Integration** plugin and click the **Edit** button.

The *Edit Plug-in* dialog is displayed.

4. For the Primary server, enter the Destination for VM migration of the Primary server by providing the following information:
 - Host (name or IP address as in vCenter)
 - Datastore
 - Resource Pool
5. For the Secondary server, enter the Destination for VM migration of the Secondary server by providing one of the following:
 - Host (name or IP address as in vCenter)
 - Datastore
 - Resource Pool
6. If integration with vSphere HA monitoring is desired, set the **Integrate with vSphere HA monitoring** to **True**.

This option requires vSphere HA Application monitoring for the cluster and VM.

7. Click **Save**.

Adding an Additional Network Interface Card

Neverfail Engine allows for installation using a single NIC on each Engine server in the Pair or Trio. When installed with a single NIC, Neverfail recommends that to prevent experiencing a single point-of-failure, an additional NIC be installed or configured on each server in a Pair or Trio with one NIC configured as the Public NIC and another configured for the Neverfail Channel.

Purpose: Add an additional network interface card (NIC) to allow moving the Channel IPs to a dedicated NIC.

Adding an additional NIC to a physical server will require that Engine be shutdown while the NIC is added and the server must be restarted. If the server is a virtual server, the shutdown is not necessary.

This procedure assumes that Engine is installed as a V2V Pair with the Primary server active and the Secondary server passive.

1. Shutdown Engine on all the nodes in the cluster and leave protected applications running.
2. On each node: Add a virtual NIC.
3. On each node: Open the *Configure Server* wizard, select the *Channel* tab, and double click the *Channel IP Routing* you are moving to the new NIC. Select the new NIC in the drop down list and click the **Edit** button.
4. On each node: Start Engine.
5. Allow the server to synchronize.

Business Application Groups

Neverfail Engine 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 the Business Application Plug-in installed, Neverfail Engine 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

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.
- Create the following folder `C:\Program Files\Ipswitch\Failover\R2\Scripts` on each server in the clusters participating in the BAG. The `StartSite` batch files scripts used by BAG will be placed in this folder.

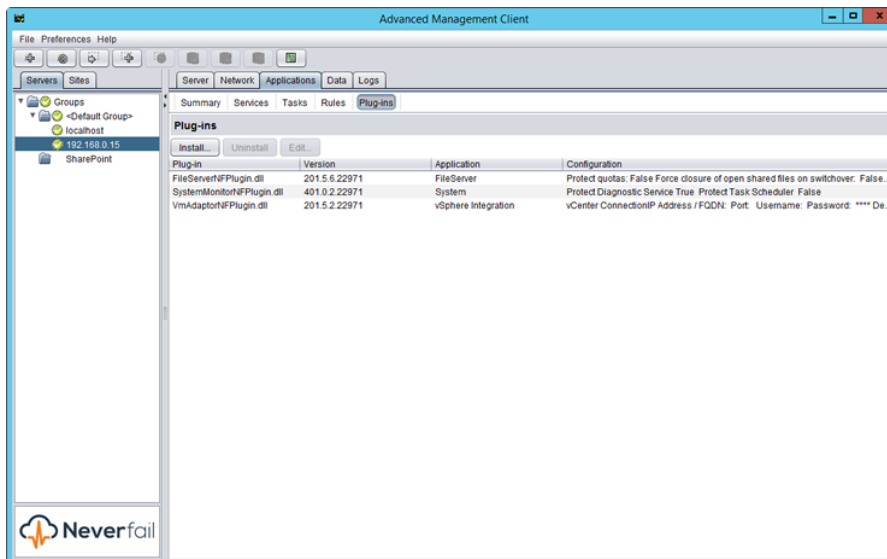
Note: Add persistent routes with a lower metric to allow them to be attempted first.

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

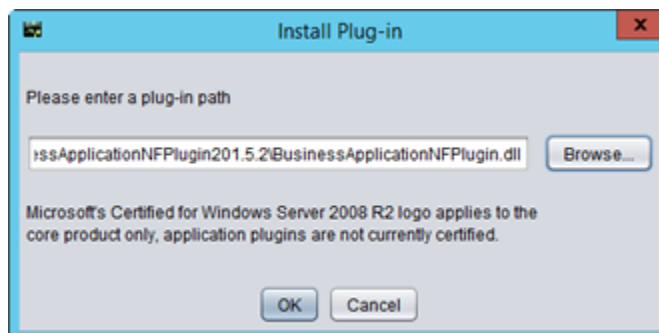
1. Download the `Z-SW-BusinessApplicationPlugin.201.5.[n].zip` file to a temporary location on the active server in the cluster.

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 Advanced Management Client and navigate to the **Applications: Plug-ins** page.



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



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

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

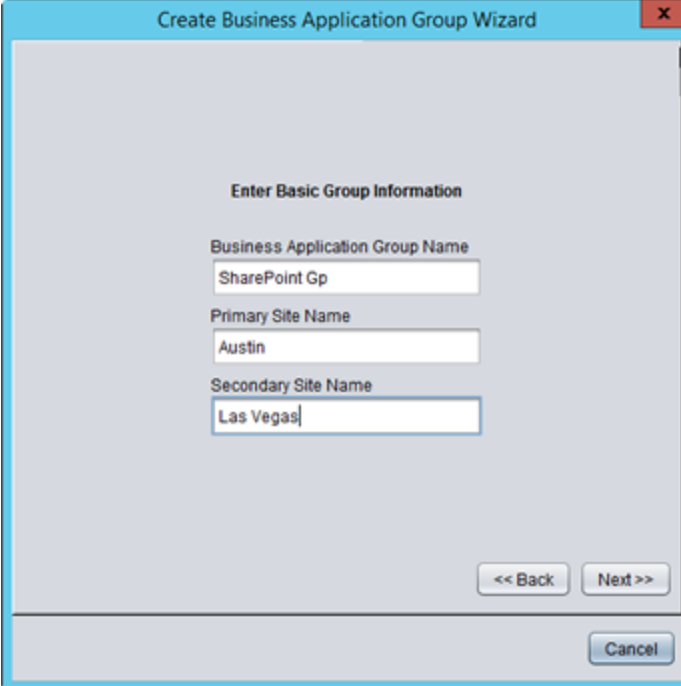
When the Neverfail Engine Business Application Plug-in is installed it is initially in an unconfigured state. The Unconfigured icon appears in the left pane of the Neverfail Advanced Management Client under Servers. All servers listed in the Unconfigured category are available as Business Application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

1. Launch the *Neverfail Advanced Management Client*.
2. Navigate to **File > Add Business Application Group**.

The *Business Application Group Wizard* is displayed.



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

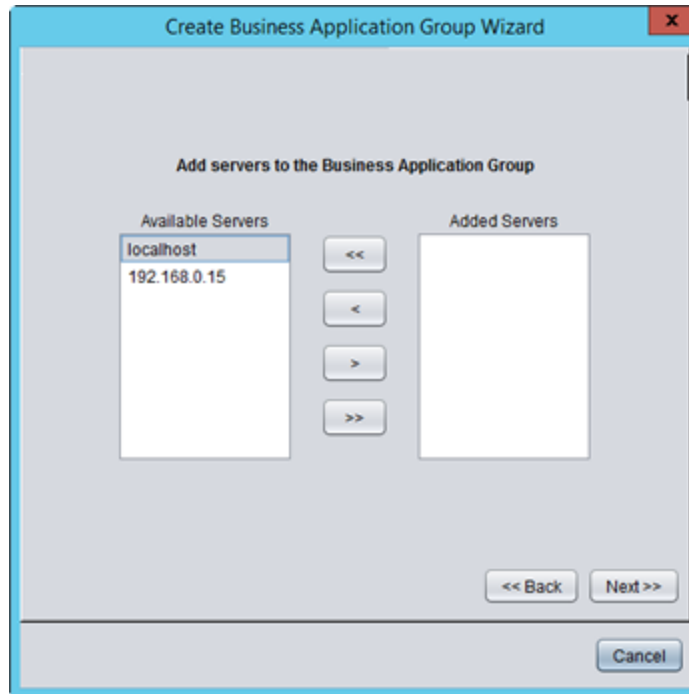


The screenshot shows a dialog box titled "Create Business Application Group Wizard" with a close button (X) in the top right corner. The main content area is titled "Enter Basic Group Information" and contains three text input fields:

- Business Application Group Name:** Contains the text "SharePoint Gp".
- Primary Site Name:** Contains the text "Austin".
- Secondary Site Name:** Contains the text "Las Vegas".

At the bottom right of the dialog, there are three buttons: "<< Back", "Next >>", and "Cancel".

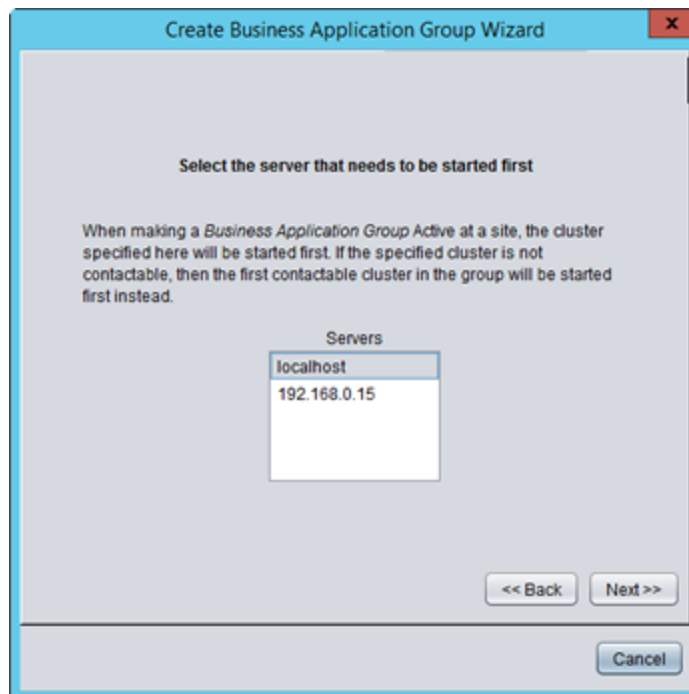
- Enter a name for the Business Application Group into the text field. The name of the Business Application Group cannot exceed 15 characters.
- Add the name of the Primary Site.
- Add the name of the Secondary (DR) site and click **Next**. The *Add Servers to Business Application Group* page is displayed. A list of available servers is displayed in the left pane of the dialog.



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.

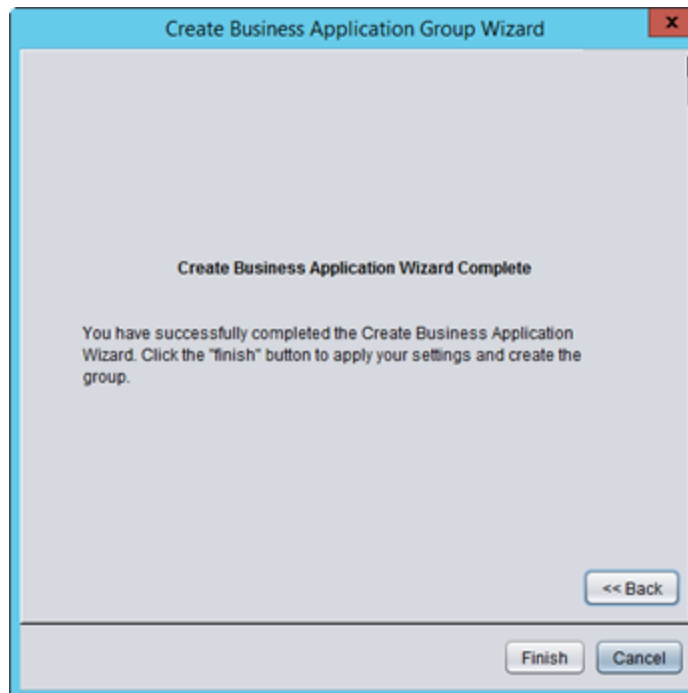
8. Select the server you want to be the first to switch within the Business Application Group. Click **Next**.



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

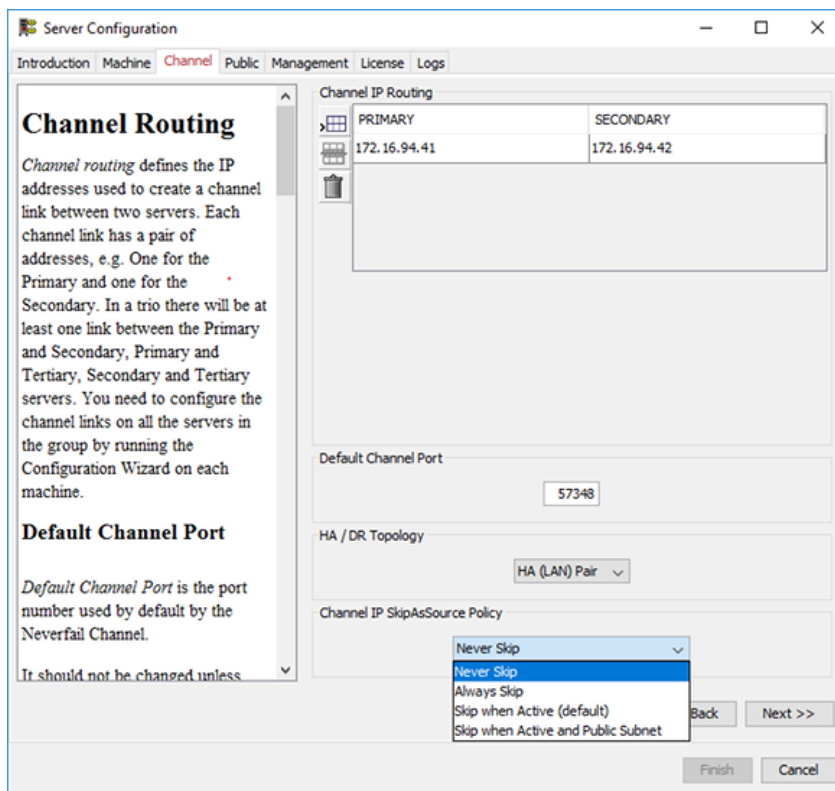
The *Create Business Application Wizard Complete* page is displayed.

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 Engine's Site Switchover capabilities discussed in **Site Switchover**. Click **Finish**.



Configuring Neverfail Engine for Business Application Group

1. If Public and Channel IP addresses share the same subnet then no changes to the **Channel IP SkipAsSource Policy** are needed.
2. If Public and Channel IP addresses belong to different subnets then configure the **Channel IP SkipAsSource Policy as Never Skip**. To do this:
 - Shutdown Engine on the active server.
 - Open *Configure Server Wizard*, go to **Channel** tab and configure the SkipAsSource policy as follows:



- Click **Finish**.
- Restart Engine on the active server.

Note: The SkipAsSource policy can be changed also without stopping Engine, by executing the following **nfclient** commands:

```
`setpe PublicIdentity AlwaysSkipChannelIPs false`
`setpe PublicIdentity ChannelIPSkipAsSourcePolicy NEVER`
```

3. For each server, add both Channel and Public IP addresses as **trusted clients** on the corresponding identity server on the other site (i.e. add Primary HA site IPs on Primary DR site and viceversa; add Secondary HA site IPs on Secondary DR site and viceversa) using the following command:

```
nfcmd localhost addTrustedClient <source\_IP\_address\> <user\_account\> administrator
```

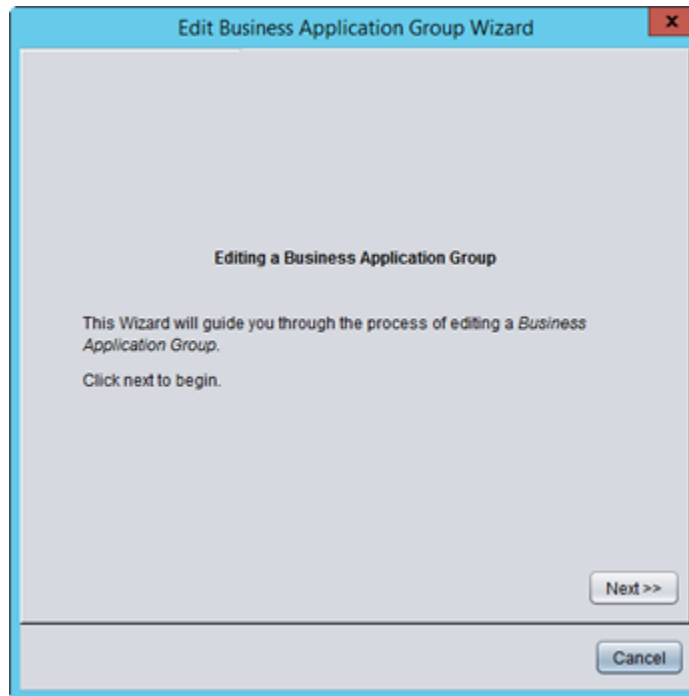
4. Verify (and make changes if needed) that Business Application Group tasks defined are configured to run with a user account added as trusted client in the above step : (e.g. if the *Administrator* account is used - the trusted client added should be *Administrator* in this case).

Editing a Business Application Group

Engine Management Service allows you to edit the configuration of an existing Business Application Group.

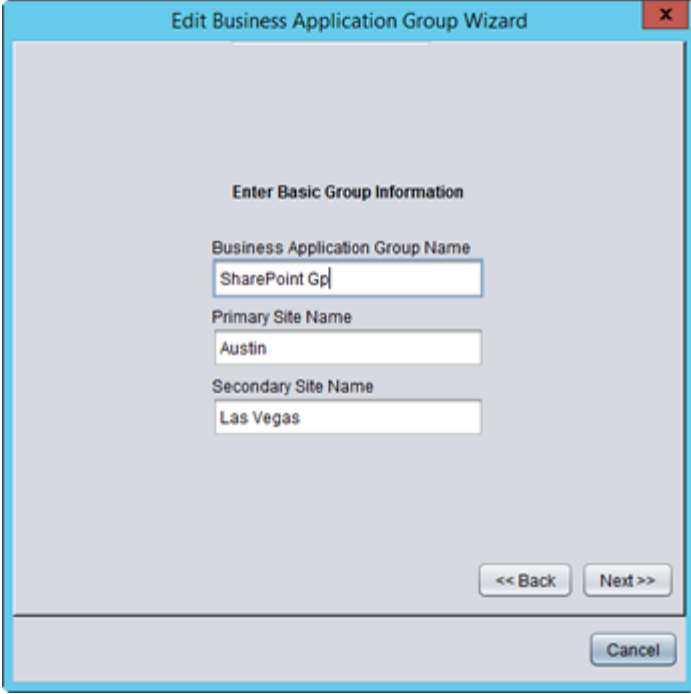
1. Navigate to **File > Edit Business Application Group** or click on the **Edit Business Application Group** button.

The *Edit Business Application Group Wizard* is displayed.



2. Click **Next**.

The *Enter Basic Group Information* page is displayed.



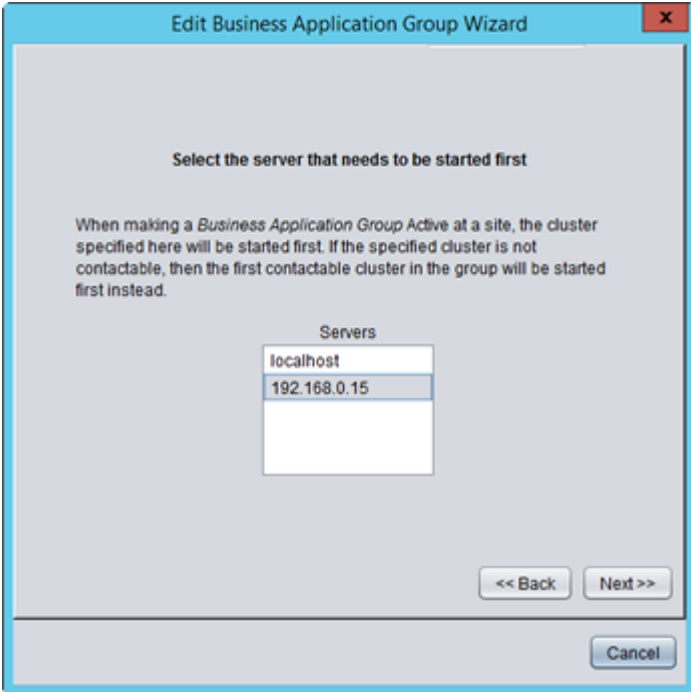
The screenshot shows a dialog box titled "Edit Business Application Group Wizard" with a close button (X) in the top right corner. The main area is titled "Enter Basic Group Information" and contains three text input fields:

- Business Application Group Name:** Contains the text "SharePoint Gp".
- Primary Site Name:** Contains the text "Austin".
- Secondary Site Name:** Contains the text "Las Vegas".

At the bottom right, there are three buttons: "<< Back", "Next >>", and "Cancel".

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.



The screenshot shows a dialog box titled "Edit Business Application Group Wizard" with a close button (X) in the top right corner. The main area is titled "Select the server that needs to be started first". Below the title is a paragraph of text:

When making a *Business Application Group Active* at a site, the cluster specified here will be started first. If the specified cluster is not contactable, then the first contactable cluster in the group will be started first instead.

Below the text is a list box titled "Servers" containing two entries:

- localhost
- 192.168.0.15

The "192.168.0.15" entry is selected. At the bottom right, there are three buttons: "<< Back", "Next >>", and "Cancel".

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**.

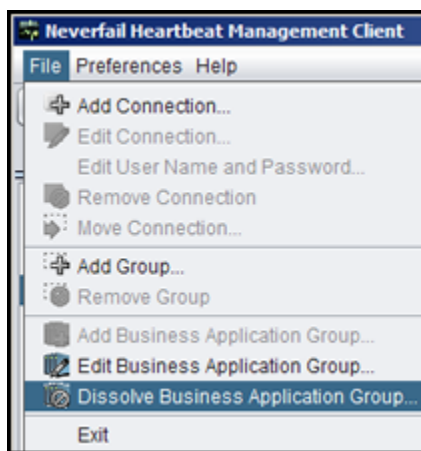
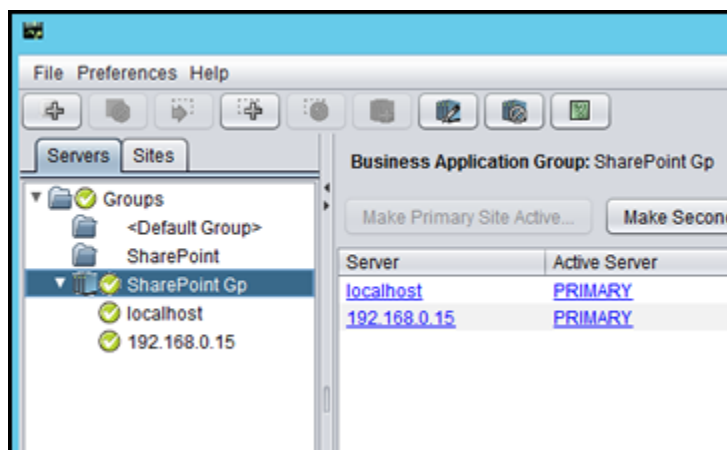
Dissolve a Business Application Group

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

1. Using the Neverfail Advanced 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**.



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



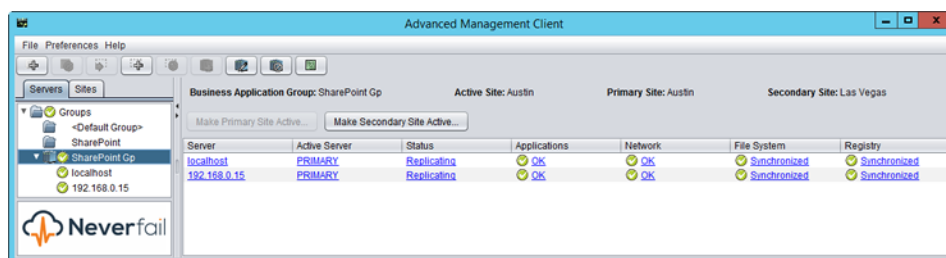
3. Click **Yes** to dissolve the Business Application Group.

Business Application Switchover

Neverfail Engine 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.




Performing a Business Application Switchover

1. Launch the Neverfail Advanced 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



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

Site Switchover

When Neverfail Engine is deployed for Disaster Recovery in a pair, Neverfail Engine 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 Engine 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.

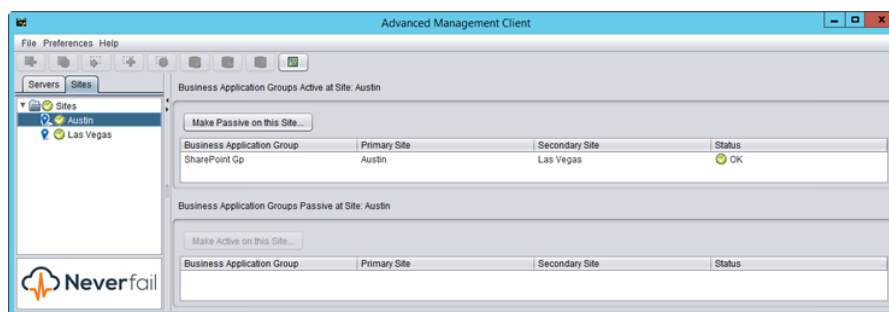
If the server that fails is the server configured to switch first, the Neverfail Advanced Management Client will be unable to connect to the host name and after a retry, will attempt to connect via the Alternate IP address. If the Alternate IP address has not been configured, then the connection will 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

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



3. Select the Site to change. Click:

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

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:

1. Launch the Neverfail Advanced Management Client.
2. Login to Neverfail Engine 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

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.

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

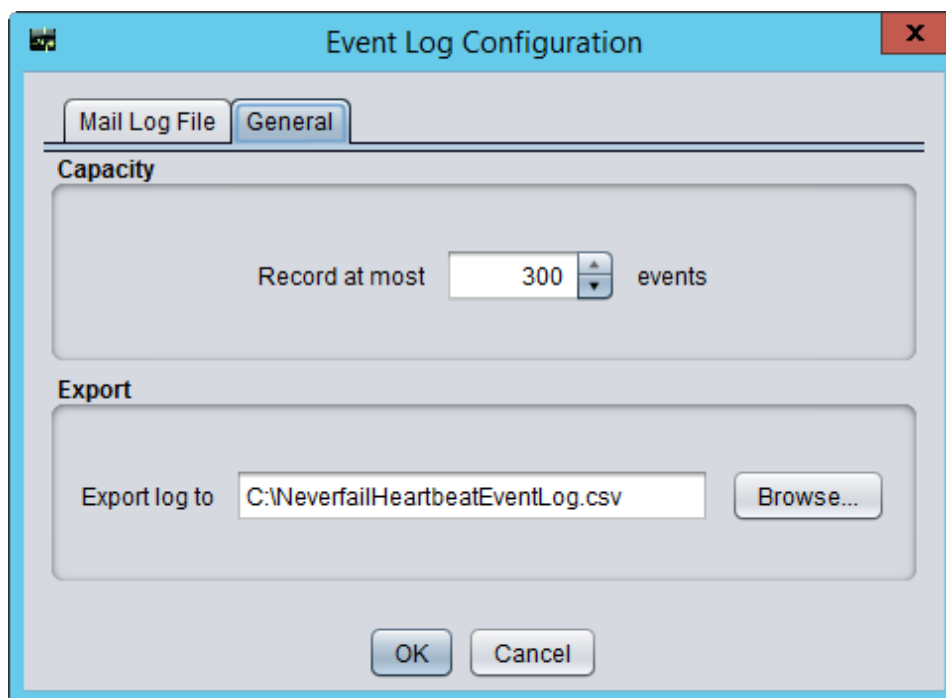
1. After dissolving the Business Application Group, select the server to have the Business Application Plug-in uninstalled.
2. Navigate to the *Applications: Plug-ins* page of the Neverfail Advanced Management Client.
3. Select the server on which to uninstall the Business Application Plug-in.
4. Select 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.

Configure Event Log Files

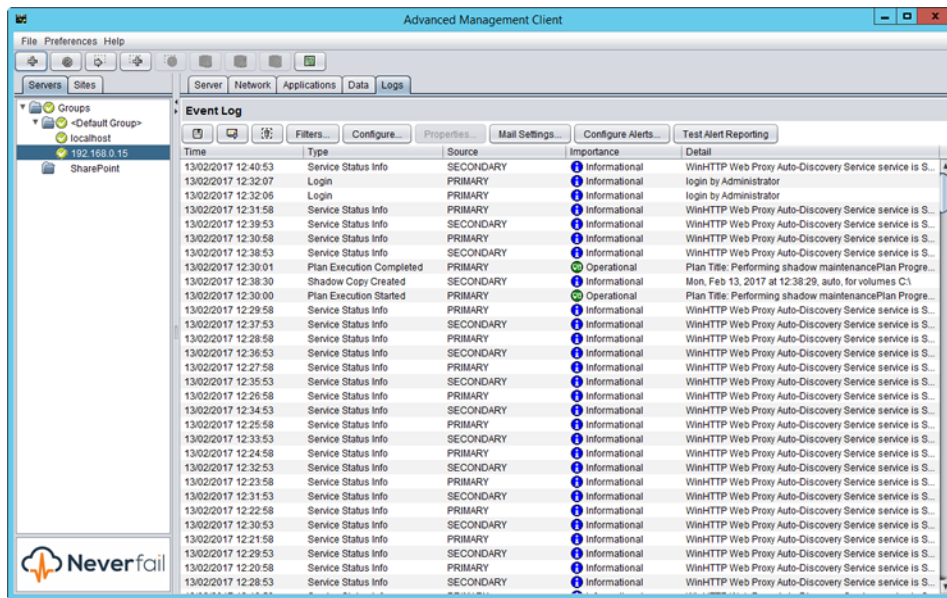
To configure default settings for log files, click **Configure** to invoke the *Event Log Configuration* dialog. Select the **General** tab to configure the log file. This dialog allows you to define where the exported comma separated variable file is stored and the name of the file by entering the path 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.



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.

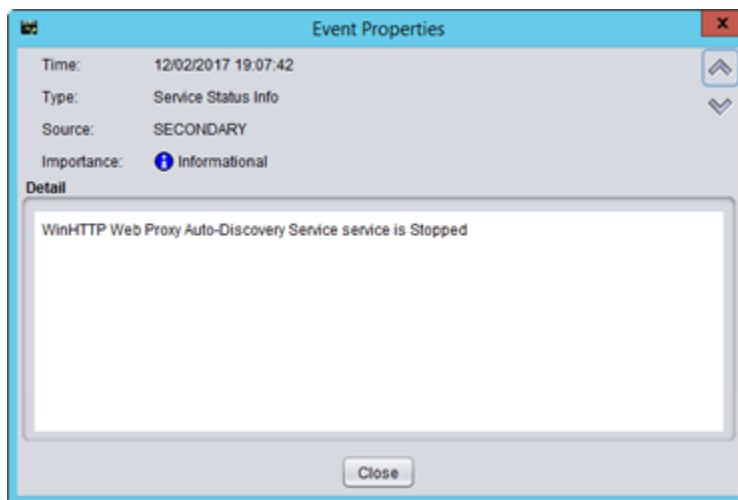
Review Event Logs

The events that Neverfail Engine 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.







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.

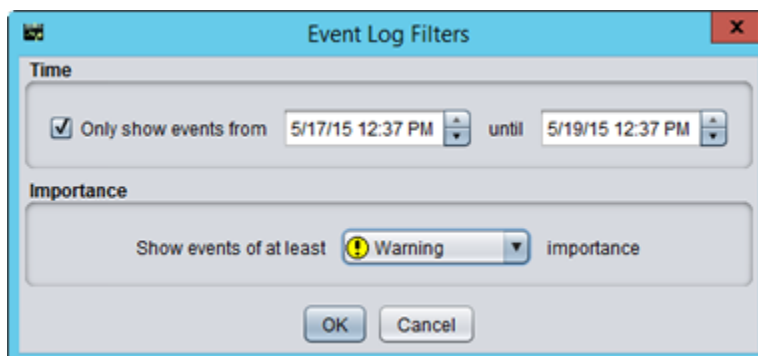


The *Event Properties* dialog gives the full detail and trace of the log that caused the event along with the source of the error aiding in troubleshooting. Further logs can be reviewed without having to close this window by using the **Up** and **Down** arrows of the dialog box to scroll through the 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 Engine by default is configured to log:

Icon	Definition
	These are critical errors within the underlying operation of Neverfail Engine and can be considered critical to the operation of the system.
	Warnings are generated where the system finds discrepancies within the Neverfail Engine operational environment that are not deemed critical to the operation of the system.
	System logs are generated following normal Neverfail Engine operations. Review these to verify the success of Neverfail Engine 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 Engine Server service itself such as logging on etc.

The list of logs that Neverfail Engine 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.



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**.

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 <i>Event Log Configuration</i> dialog (click Configure).
	Mail event log to recipients immediately - Click to email the list to recipients immediately.

Recloning Secondary or Tertiary Server

Recloning the Secondary or Tertiary standby instances flow helps you recreate your passive instances from your active, Production instance of your cluster. This flow can be employed, for example, when you need to update or upgrade applications running on your protected server cluster (read more in the **Upgrade Applications** article).

Note: While recloning a passive instance, you cannot alter the channel connection configuration. Only the storage host of the instance can be configured when recloning.

It can also provide a quick way to reconstruct your standby instances, when you don't need to alter the channel configuration. For example, you can rebuild an unavailable passive instance with the exact same configuration as before, or you could employ this flow to move a passive instance to a different storage location.

The screenshot shows a wizard window titled "Reclone Secondary or Tertiary | lj-w10". The wizard has four steps: "Select clone type", "Select nodes", "Select location", and "Ready to complete". The "Select clone type" step is active and shows three options: "Automated cloning" (selected), "Create a temporary powered-off DR cloned VM locally, so that the VMDK files can be transferred manually", and "Assisted cloning". Below these options is a "DR VM clone type" section with explanatory text. At the bottom right, there are "Back" and "Next" buttons.

Reclone Secondary or Tertiary | lj-w10

Select clone type | Select nodes | Select location | Ready to complete

Automated cloning

- Create and power-on the DR cloned VM automatically after cloning
- Create a temporary powered-off DR cloned VM locally, so that the VMDK files can be transferred manually
- Assisted cloning

DR VM clone type

You can select to reclone the stand-by VMs immediately, either automated or assisted. Alternatively, you can opt for planning a [Scheduled auto recloning](#).

If you have a reliable, high-bandwidth connection to the remote site, you can choose to create a stand-by server VM directly on its host. This is recommended only if you have previously cloned VMs to the remote site with success.

Alternatively, you can create the VM in a temporary location on a local host. The VM will not be powered-on. You can then transfer the VMDK files to the remote site, e.g. using detachable storage or FTP. Scheduled recloning is not available with this option.

Back Next

Note: You can find out more about the use cases in which the Engine's recloning use is recommended here: [When to Use Neverfail Patch Management Options](#)

When triggering a server reclone, certain prerequisites must be met before the procedure starts:

- the Primary node is running (active) and serving applications.
- for automated recloning: the VMware vCenter Server connection must be set up correctly in the Engine Management Software.
- for automated recloning: VMware vCenter Server Converter must be configured if the Primary node is not a VMware virtual machine.

When the above prerequisites are met, the cluster is in the Ready State. The Engine cluster may be complete or incomplete: any of the passive servers, Secondary or Tertiary, may be present or not.

Recloning Passive Nodes with configured Static Routes - supported scenarios:

- IPv4 static routes created using the route command.

The route command is used to view and modify the network routing tables of an IP network. For example:

```
route add 192.168.33.63 mask 255.255.255.255 192.168.33.254 IF 12 -p
```

The above command adds a persistent static route for the 192.168.33.63 destination IP address, associated with the NIC interface defined by index 12, using the 192.168.33.254 address as next gateway.

- All the single NIC deployments.
- All virtual-to-virtual (V2V) deployments, where the passive nodes recloning is done via VMware vCenter cloning method.
- All virtual-to-virtual-to-virtual (V2V2V) deployments, where the passive nodes recloning is done via VMware vCenter cloning method.
- Automated, Assisted and Scheduled Automated recloning options, considering the above conditions are met.

To start the recloning flow, open the **Reclone Secondary or Tertiary** dialog from the **Actions** menu.

1. The Select clone type section allows you to choose between **Automated cloning** and **Assisted cloning**:

- The automated cloning option is available only if the standby instance to be reclone was previously cloned in an automated way. While for HA instances, the process is straight forward, there are two automated cloning methods for DR instances:

- Powered-on clone, which will create the cloned DR VM and power it on.
 - Powered-off clone, which will create the cloned DR VM without powering it on, allowing the user to transfer the VM manually to a remote host.
 - The assisted cloning option allows the user to perform all the cloning and transfer operation manually.
2. The Select nodes section allows you to choose what to be cloned:
- **Full cluster reclone** will reclone all standby instances.
 - **Partial cluster reclone** is available only in trio cluster configurations and allows you to select the standby instance to be reclone.
3. The Select location section allows you to choose the location of the reclone instance:
- **Current location**, while keeping the original VM.
 - **Current location**, deleting the original VM once reclone.
 - **New location**, in which case you need to specify the **Host** and **Datastore**.
4. The Ready to complete section will summarize the reclone options. Clicking the **Finish** button will initiate the reclone process and its progress will be visible in the **Operations in progress** section.

Note: If the new location has not been configured for the reclone job or the EMS is not able to retrieve the original reclone target location from the Secondary or Tertiary nodes, the VMware vCenter Server Default Host will be used for the target reclone location. The Default Host needs to be configured in the VMware vCenter Server connection settings.

Schedule recloning

The schedule recloning feature allows you to trigger the automatic reclone flow at a predefined time and date. This can prove very useful when performing recurring maintenance tasks on the cluster.

Note: To use the Schedule recloning feature, an automated recloning flow must have been previously applied to the standby instances to be scheduled for recloning. Scheduling only works for automated recloning. Assisted recloning cannot be scheduled.

Auto recloning | ij-w10 ✕

Select schedule Ready to complete

A schedule cannot be created for a server that was not fully cloned via vCenter Server.

A schedule cannot be created for a server that was created via vCenter Server, but was a powered-off VM.

A scheduled automated recloning will be executed successfully only when the vCenter Server configured is the same as that of the original clone.

Clear schedule

Disable schedule

Enable schedule

Once every month on the

Twice every month on the

Every second week on

Begin recloning:

Starting at

Some time between the hours of and

Delete original Secondary VM

Delete original Tertiary VM

Next

To configure a scheduled automated reclone, open the **Auto recloning** dialog by clicking the Auto reclone **Configure** button in the Summary Status panel of the Server Details page. The dialog provides the following options:

- **Clear schedule:** deletes the current recloning schedule, if existing.
- **Disable schedule:** stops the existing scheduled reclone, while keeping the schedule configuration for a later enable.
- **Enable schedule:** enables the scheduled reclone operation using the specified settings:
 - Once every month on the specified day of month.
 - Twice every month on the specified days of the month.

Note: that the execution days always have a two weeks period between them (unless the first execution is set to the 14th of the month, then the second execution will be triggered in the last day of the month, regardless of how many days the month has).

- Every second week on the specified day.
- Starting time of the recloning operation. It can either be exactly specified or the starting time can be set to an optimal time inside a specified time interval.
- Original VM deletion, which can be enabled individually for Secondary and Tertiary instances. This will remove the clone source VM once the reclone operation is successfully completed.

Note: As the Scheduled recloning procedure is depending on the time of Primary node, it is mandatory that the time is correctly configured on both EMS and Engine.

The Ready to complete section of the dialog will resume the recloning schedule. Clicking **Finish** will save and apply the defined schedule.

Troubleshooting

- **Two Active Servers**
- **Two Passive Servers**
- **Invalid Neverfail Engine License**
- **Synchronization Failures**
- **Channel Drops**
- **MaxDiskUsage Errors**
- **Application Slowdown**

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.

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) and S/A (Secondary and active).
2. An IP address conflict may be detected in a Cluster running Neverfail Engine on the Public IP address.
3. A name conflict may be detected in a Cluster running Neverfail Engine. In a typical WAN environment, the Primary and Secondary servers 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 Engine.

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 Engine software

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 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 Engine. 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 Engine 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.
Do not change the Identity of the server (Primary or Secondary).
 5. Click **Finish** to accept the changes. Reboot this server.
 6. Start Neverfail Engine (if required) and verify that the task bar icon now reflects the changes by showing **PI-** (Primary and passive) or **SI-** (Secondary 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 Engine (if required) and verify that the task bar icon now reflects the changes by showing **PIA** (Primary and active) or **SIA**(Secondary and active).
11. Log into the Neverfail Advanced Management Client.
12. Verify that the servers have connected and replication has started.

Two Passive Servers

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

The first indication that Neverfail Engine 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.

- Two passive servers generally results from some kind of sudden failure on the active server - for example, unexpected termination of the Neverfail Engine 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 Engine without reconfiguring one server as active, Neverfail Engine responds with the warning: No active server amongst [PRIMARY, SECONDARY]

To resolve two passive servers, perform the following steps.

1. Determine which server to make active.

2. If Neverfail Engine 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 all servers. This ensures that all protected application services are stopped on the passive servers and started on the active server.
6. Start Neverfail Engine on both servers.

Invalid Neverfail Engine License

The Neverfail Engine 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.

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

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 Engine License has expired - If a licensing problem arises during an implementation, Neverfail 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 Engine from starting when the date is exceeded.
 - Windows Management Instrumentation (WMI) hung or not running. Neverfail Engine uses WMI to validate the license on the Primary server and if WMI is hung or not running validation cannot complete.
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 Engine License key, simply generate request a new license key for the Primary server.
 2. If the invalid license error is not due to expiration of a temporary or time-limited Neverfail Engine 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 Engine.

Synchronization Failures

When Neverfail Engine 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

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

A service that is running on the passive server may open a protected file for exclusive access. If Neverfail Engine 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 Engine 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 a Neverfail Engine for IIS deployment (if the `\inetpub` folder is shown as *Out-of-sync*)
- IISAdmin service in a Neverfail Engine for IIS deployment (if `C:\WINDOWS\system32\inetsrv\MetaBase.xml` is shown as *Out-of-sync*). IISAdminservice starts on the passive after a reboot of the server and must be stopped manually.

Until the file is closed on the passive server, Neverfail Engine 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 a Neverfail Engine 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 a Neverfail Engine 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

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*.

The most common Neverfail Channel configuration errors are:

- Channel IP addresses configured in different subnets (in LAN configurations)
 - In a WAN configuration, no static routes between the channel NICs
1. Verify that channel IP addresses are properly configured.
 2. In a WAN configuration, verify that static routes between channel NICs are properly configured.
 3. 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:

When Neverfail Engine 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 Engine requires the same drive letter configuration on the Primary and Secondary servers, and only supports protection of NTFS partitions.

Different partition structures on the Primary and Secondary 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 server
- The Primary server has drive G:, which is a valid NTFS partition; but the equivalent drive on the Secondary server is a CD / DVD drive or a FAT / FAT32 partition, which cannot be protected by Neverfail Engine.

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**.

1. 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 Passive Server has Less Available Space than the Active Server

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.

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

1. Free up some additional disk space on the 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 passive server.
2. When complete, you must manually start replication.

Unprotected File System Features

Another possible reason why Neverfail Engine 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 Engine 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*.

Neverfail Engine does not synchronize if the replication set contains files with unprotected file system features. Unprotected file system features are described by category the following KB: [Neverfail for File Server: Unprotected Features of the Windows NTFS File System](#).

1. Two methods of dealing with these Unprotected Features are described in the following KB: [Neverfail for File Server: Unprotected Features of the Windows NTFS 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 Engine 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

Neverfail Engine logs the following error message:

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

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

1. 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

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

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

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

Channel Drops

Performance Issues

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.

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 Engine R2 service itself experiences problems and may hang or terminate unexpectedly.
1. Determine the likely source of the hang or reboot by examining the Windows event logs.
 2. 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

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

The passive server does not meet the recommended hardware requirements for Neverfail Engine 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 passive server at the rate that the data is sent to the is passive server.

1. To avoid reinstalling your Neverfail Engine 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 or Secondary) of the affected server before you perform the upgrade.

Hardware or Driver Issues on Channel NICs

The Neverfail Channel drops or disconnects and reconnects intermittently.

- 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
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 - Neverfail Engine 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 Engine, 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.

The Neverfail Channel cannot connect or connects and disconnects continuously.

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.

1. 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

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

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 Engine 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 Public NIC on the Secondary server; or the name and IP address of the Public NIC may be transferred to a channel NIC. 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 server.

1. It is part of the normal Neverfail Engine installation process to manually assign the correct IP addresses to each NIC on the Secondary server. If there is no channel connection between the servers, verify that the IP addresses on the Secondary server's channel NICs are correctly configured. Verify the settings for the 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 server prior to installing Neverfail Engine, 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 server to a file called `ipconfig.txt`, which is present on the server after the Plug and Play phase of the Neverfail Engine 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.

Subnet or Routing Issues In a LAN

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

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

1. If Neverfail Engine is deployed in a LAN environment, the 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 or Routing Issues In a WAN

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

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 Engine is deployed in a WAN, it is generally not possible for the 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.

1. Refer to KB: [How to Create a Static Route for the Neverfail Channel Connection in a WAN Environment](#) 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 Engine uses queues to buffer the flow of replication data from the active server to the passive server. This configuration provides resilience in the event of user activity spikes, 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 Engine 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 Engine considers the send queue as 'unsafe' because the data in this queue is awaiting replication across the channel to the 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 Engine, 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 Engine considers the receive queue safe because the data in this queue has already been transmitted across the channel from the active to the 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 both servers) are stored on-disk, by default in the <Neverfail Engine 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 Engine started, open the Neverfail Advanced Management Client and select **Data: Traffic/Queues**. Click the **Configure** button. Configure the value for the *Max Disk Usage* and click **OK**. It is necessary to shut down and restart Neverfail Engine (specify that the stopping of protected applications is not necessary) for the change to take effect.
- With Neverfail Engine 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 Engine is a symmetrical system, and can operate with either server in the active role. For this reason, the queue size is always set to the same value for both servers.

MaxDiskUsage Errors

If Neverfail Engine 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 Engine reports [L9] Exceeded the maximum disk usage (NFChannelExceeded-MaxDiskUsageException), 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

Replication stops and restarts or stops completely (if the event occurs while a Full System Check is in progress) and the Neverfail Engine Event Log displays the error [L9]Exceeded the maximum disk usage, originating from the ACTIVE server.

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.

1. Assuming there are no other channel connection issues (see KB: [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 Engine to utilize any free disk space to store the queues.

L9 Exceeded the Maximum Disk Usage on a PASSIVE Server

Replication stops and restarts or stops completely (if the event occurs while a Full System Check is in progress) and the Neverfail Engine Event Log displays the error [L9]Exceeded the maximum disk usage, originating from a PASSIVE server.

- 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 server is more powerful than the Primary server, hardware-related issues might only occur while the Secondary server is in the active role.

If you have multiple physical disks on each server, it may be worth locating the Neverfail Engine 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 Engine.
2. Open the **Server Configuration Wizard** and select the *Logs* tab.
3. Set the intended path for *Message Queue Logs Location* and click **Finish**.
4. Start Neverfail Engine 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 Engine 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 Engine 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 Engine). 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

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

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.

1. 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 Engine and for servers running any other application. Neverfail Engine 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 Engine meet recommended hardware requirements and are powerful enough to cope with the load imposed by Neverfail Engine, 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 Engine is running.

Poor Application Performance

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

This may be due to the active server's resource usage in one or more areas being close to the maximum possible before Neverfail Engine was installed.

1. 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 per-

forming any recommended hardware upgrades on the active server before Neverfail Engine is installed.

Servers Could Accommodate the Initial Load but the Load has Increased

Application response times have slowed in response to increased user activity.

It is also possible that the servers may be able to operate normally when Neverfail Engine 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.

1. 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

Applications operate normally when the Primary server is active but slow when the Secondary server is active (or vice versa).

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 is active (or vice-versa). Problems may arise even when the more powerful server is active, such as when resource intensive tasks are running.

1. 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

Resource-intense scheduled tasks impact performance at certain times.

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 Engine 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.

1. 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.