



Administrator's Guide

For Neverfail Continuity Engine v8.5

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 © 2018 Neverfail, LLC. All rights reserved.

About This Book

The Administrator's Guide provides information about configuring and performing the day-to-day management of Neverfail Continuity 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, and 2016 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.

Overview of Content

This guide is designed to give guidance on the configuration and administration of Neverfail Engine, and is organized into the following sections:

- **Preface** — About This Book (this chapter) provides an overview of this guide and the conventions used throughout.
- **Chapter 1** — Neverfail Engine Concepts presents an overview of Neverfail Engine architecture and the five levels of protection provided by Neverfail Engine.
- **Chapter 2** — Status and Control describes how to connect to Neverfail Engine using the Engine Management Service or the Neverfail Advanced Management Client to review the status of and manage a Cluster.
- **Chapter 3** — Configuring Neverfail Engine discusses how to configure Neverfail Engine using the Configure Server Wizard.
- **Chapter 4** — Server Protection discusses how the Neverfail Engine solution protects users from server system failure or server hardware crash.
- **Chapter 5** — Network Protection describes how Neverfail Engine protects against network failure by ensuring that the network identity of the production server, IP address, etc. are provided to users.
- **Chapter 6** — Application Protection discusses how Neverfail Engine maintains the protected application environment ensuring that applications and services stay alive on the network.

- **Chapter 7** — Data Protection discusses how Neverfail Engine intercepts all data written by users and protected applications and maintains a copy of this data for use in case of failure.
- **Appendix A** — Other Administrative Tasks discusses additional tasks for the administrator to configure system logging and alerting functions. Neverfail vAdministrator's Guide
- **Appendix B** — Troubleshooting discusses common issues that may appear and techniques to troubleshoot the issue and includes two active servers or two passive servers, application slowdown, channel drops, and MaxDiskUsage errors.
- **Appendix C** — Neverfail SCOPE Data Collector discusses how to use Neverfail SCOPE to measure bandwidth, and interrogate your server environment to prepare for installation.

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 Continuity 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 Continuity 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 Continuity Engine Documentation Library

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

Document	Purpose
Installation Guide	Provides detailed setup information.
Administrator Guide	Provides detailed configuration and conceptual information.
Online Help	Provides help for every window in the Engine Management Service user interface
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.
Logical OR, as in value1 value2	Exclusive command parameters where only one of the options can be specified.

Table of Contents

Chapter 1. Neverfail Continuity Engine Concepts	1
1. Architecture	1
2. Protection	3
3. Neverfail Continuity Engine Networking Configuration	4
4. Neverfail Continuity Engine Communications	6
5. Neverfail Continuity Engine Switchover and Failover Processes	8
6. Recovery from a Failover	11
Part I. Configuration	12
Chapter 2. Status and Control	13
1. Using the Engine Management Service User Interface	13
1.1. Configure Connection to VMware vCenter Server	13
1.2. Configure VMware vCenter Converter	16
1.3. Protected Servers	17
1.4. Management	18
1.4.1. Deploy	18
1.4.1.1. Configure Windows Firewall for Deployment	18
1.4.1.2. Deploy to a Primary Server	18
1.4.1.3. Upgrade the Selected Server	21
1.4.1.4. Uninstall from the Selected Server	24
1.4.1.5. Add a Stand-by Server for High Availability	25
1.4.1.6. Add a Stand-by Server for Disaster Recovery	29
1.4.1.7. Create Secondary and Tertiary stand-by VMs for HA and DR	33
1.4.1.8. Upgrade Applications	40
1.4.1.9. Recloning Secondary or Tertiary Server	41
1.4.2. Manage	48
1.4.2.1. Discover Protected Servers	48
1.4.2.2. Add a Protected Server	49
1.4.2.3. Remove the Selected Server	50
1.4.2.4. Download the Advanced Management Client	51
1.4.2.5. Open or View Support Ticket	52
1.4.2.6. Send Product Feedback	52
1.4.3. Integrate	54
1.4.3.1. Log in to VMware vSphere Client	54
1.4.3.2. Create VMware SRM Plan Step for Selected Server	54
1.4.4. License	57
1.4.4.1. Configure an Internet Proxy Server for Licensing	57
1.4.4.2. License the Selected Server	58
1.5. Summary	60

1.6. Status	61
1.6.1. Summary Status	62
1.6.2. Plan Execution	63
1.6.3. Applications and Platforms	63
1.7. Events	64
1.8. Services	65
1.8.1. Add a Service	65
1.8.2. Edit a Service	67
1.8.3. Configure Service Recovery Options for Protected Services	68
1.8.4. Remove a Service	69
1.9. Data	69
1.9.1. Add File Filters	71
1.9.2. Edit Filters	72
1.9.3. Remove Filters	73
1.10. Shadows	73
1.10.1. Best Practices for Using Volume Shadow Copy Service & DRM	73
1.10.2. Configure Shadow Creation Options	75
1.10.3. Configure the Shadow Copy Schedule	76
1.10.4. Configure Shadow Keep Options	78
1.10.5. Manually Create Shadow Copies	79
1.10.6. Delete a Shadow Copy	80
1.10.7. Roll Back Protected Data to a Previous Shadow Copy	80
1.11. Tasks	82
1.11.1. Run Now	83
1.11.2. Add Task	83
1.11.3. Edit Task	84
1.11.4. Remove Task	85
1.12. Rules	86
1.12.1. Check a Rule Condition	87
1.12.2. Edit a Rule	87
1.13. Settings	88
1.13.1. Configure Plug-ins	89
1.13.2. Alert Settings	90
1.13.3. Using WScript to Issue Alert Notifications	91
1.13.4. Alert Triggers	92
1.13.5. Email Settings	93
1.13.6. Wan Compression	94
1.13.7. Replication Queue Settings	95
1.14. Actions	97
2. Managing Neverfail Continuity Engine Clusters	100
3. Review the Status of Neverfail Continuity Engine Clusters and Groups	101
4. Exit Neverfail Advanced Management Client	102

5. Shutdown Windows with Neverfail Continuity Engine Installed	102
6. Controlled Shutdown	102
Chapter 3. Configuring Neverfail Continuity Engine	104
1. Configure Server Wizard	104
2. Configure Machine Identity	105
3. Configure Server Role	106
4. Change the Client Connection Port	107
5. Configure Channel IP Routing	107
6. Configure the Default Channel Port	108
7. Configure Low Bandwidth Optimization	109
8. Configure Public IP Addressing	110
9. Management IP Addressing	111
10. Considerations for Passive Node Management Using Third Party Technology	113
11. Add/Remove a Neverfail Continuity Engine License Key	113
12. Configure the Message Queue Logs	114
13. Configure Maximum Disk Usage	115
Part II. Management	117
Chapter 4. Server Protection	118
1. Monitoring the Status of Servers	118
2. Configure Neverfail Continuity Engine Settings	119
2.1. Configure Pings	119
2.2. Configure Failover	120
2.3. Configure Response Times	121
2.4. Common Administrative Tasks in Neverfail Continuity Engine	122
3. Forcing a Switchover	123
4. Failover versus Switchover	123
4.1. Configuring Failover and Active Server Isolation	124
4.2. Recover From a Failover	128
5. Split-brain Avoidance	130
Chapter 5. Network Protection	132
1. Configure Public Network Monitoring	133
2. Enabling Automatic Switchover in a WAN	134
3. Setting Max Server Time Difference	135
Chapter 6. Application Protection	136
1. Applications Environment	136
2. Applications: Summary	136
3. Applications: Services	140
4. Applications: Tasks	142
Chapter 7. Data Protection	145
1. Data: Replication	145
1.1. Initiate a Full System Check	146

1.2. Fast Check	147
1.3. Manually Initiate File Synchronization	148
1.4. Manually Initiate Verify and Synchronize	149
1.5. Orphaned Files Management	150
Part III. Reference	153
Chapter 8. Other Administrative Tasks	154
1. Post Installation Configuration	154
1.1. Configure the VmAdapter Plug-in	154
1.2. Adding an Additional Network Interface Card	155
2. Business Application Groups	155
2.1. Installing the Business Application Plug-in	156
2.2. Creating a Business Application Group	158
2.2.1. Configuring Neverfail Engine for Business Application Group	161
2.3. Editing a Business Application Group	163
2.4. Dissolve a Business Application Group	165
2.5. Business Application Switchover	166
2.5.1. Performing a Business Application Switchover	167
2.6. Site Switchover	167
2.6.1. Performing a Site Switchover	168
2.6.2. Perform a Site Switchover when the First Server to Switch is Unavailable	169
2.7. Uninstall the Business Application Plug-in	169
3. Configure Event Log Files	170
4. Review Event Logs	170
5. Recloning Secondary or Tertiary Server	172
5.1. Automated Recloning	173
5.2. Manual Recloning	178
5.3. Scheduled Recloning	179
6. Deploying Neverfail Engine 8.5 Cluster in Amazon Web Services Cloud Environment	181
6.1. Installing Neverfail Engine 8.5 DR Pair in different Amazon Web Services VPCs	182
7. Deploying a Passive Node in an Amazon Web Services Cloud Environment	188
Chapter 9. Troubleshooting	197
1. Two Active Servers	197
2. Two Passive Servers	199
3. Invalid Neverfail Continuity Engine License	200
4. Synchronization Failures	201
4.1. Services Running on the Passive Server	201
4.2. Neverfail Channel Incorrectly Configured	202
4.3. Incorrect or Mismatched Disk Configuration	203
4.4. The Passive Server has Less Available Space than the Active Server	203

4.5. Unprotected File System Features	204
4.6. Registry Status is Out-of-Sync	205
4.6.1. Resource Issues	205
4.6.2. Registry Security Issues	205
5. Channel Drops	206
5.1. Performance Issues	206
5.2. Passive Server Does Not Meet Minimum Hardware Requirements	206
5.3. Hardware or Driver Issues on Channel NICs	207
5.4. Firewall Connection	208
5.5. Incorrect Neverfail Channel Configuration	208
5.6. Subnet/Routing Issues — In a LAN	209
5.7. Subnet/Routing Issues — In a WAN	210
6. MaxDiskUsage Errors	210
6.1. [L9]Exceeded the Maximum Disk Usage on the ACTIVE Server	212
6.2. [L9]Exceeded the Maximum Disk Usage on a PASSIVE Server	212
6.3. [L20]Out of disk space (NFChannelOutOfDiskSpaceException)	214
7. Application Slowdown	214
7.1. Poor Application Performance	215
7.2. Servers Could Accommodate the Initial Load but the Load has Increased	215
7.3. One Server is Able to Cope, but the Other Cannot	216
7.4. Scheduled Resource Intensive Tasks	216
Chapter 10. Neverfail SCOPE Data Collector Service Overview	218
1. Using Neverfail SCOPE Data Collector Service	218
1.1. Configuring Neverfail SCOPE Data Collector Service	218
1.1.1. Configure the General tab	219
1.1.2. Configure the Data Files tab	220
1.1.3. Configure the Connectivity tab	221
1.1.4. Configure the Support tab	222
1.1.5. Automatic Configuration	223
1.1.6. Manual Configuration	223
1.1.7. Neverfail SCOPE Data Collector Service Parameters	224
1.1.8. Configure Bandwidth Measurement	227
1.2. Neverfail SCOPE Data Collector Service Network Ports	228
1.3. Daylight Savings Time	228
2. Neverfail SCOPE Analysis Reports	228
2.1. Neverfail SCOPE Reports	228
2.2. Neverfail SCOPE Graphs	230
2.3. Neverfail SCOPE Performance Counters	230

Chapter 1. Neverfail Continuity Engine Concepts

Neverfail Continuity 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 Continuity 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 Continuity Engine Cluster. Refer to the **Glossary** for more information about Neverfail Engine Clusters.

Related information

- [Architecture](#)
- [Protection](#)
- [Neverfail Continuity Engine Networking Configuration](#)
- [Neverfail Continuity Engine Communications](#)
- [Neverfail Continuity Engine Switchover and Failover Processes](#)
- [Recovery from a Failover](#)

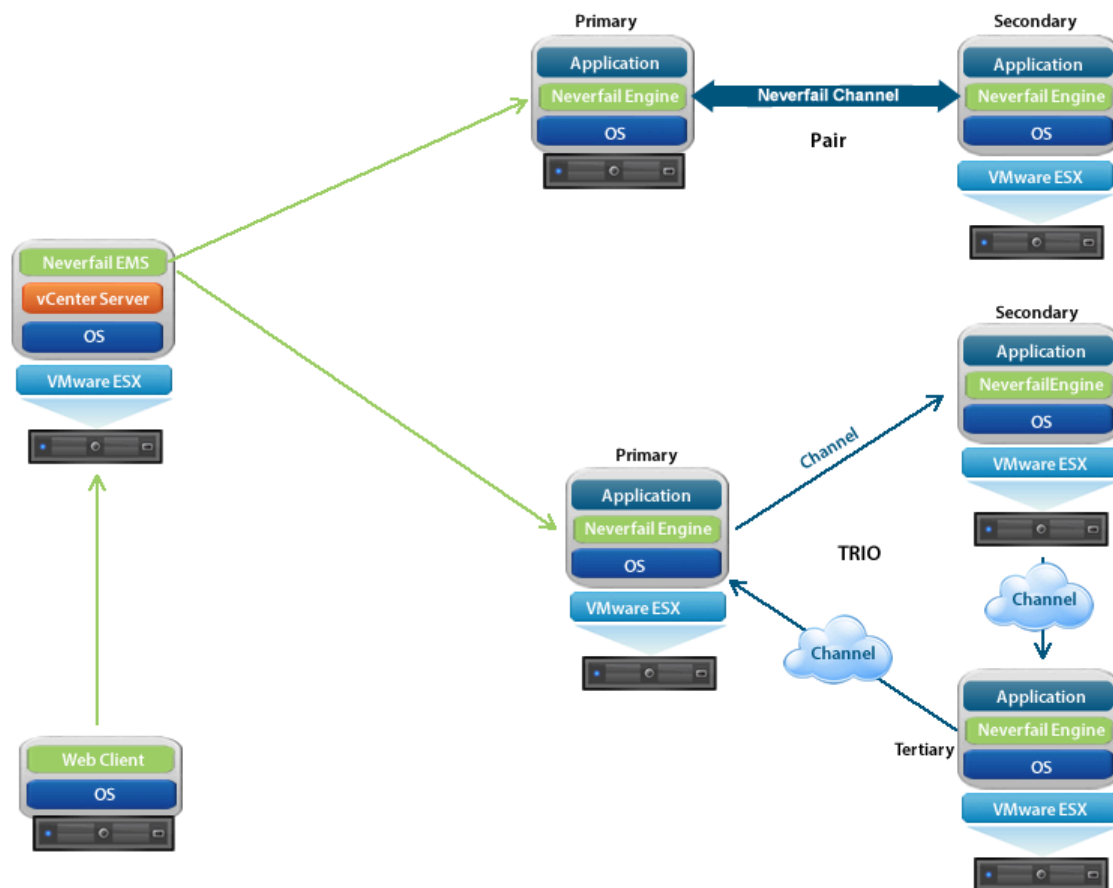
1. Architecture

Neverfail Continuity 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 Continuity Engine brings a best in class solution for protecting critical business systems.

Neverfail Continuity 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.

Figure 1-1. Deployment Architecture



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

Neverfail 9 Administrator's Guide 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 Continuity 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)

2. Protection

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

- **Server Protection** — Neverfail Engine continues to provide availability to end-user clients in the event of a hardware failure or operating system crash. When deployed, Neverfail 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** — Neverfail 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, Neverfail 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** — Neverfail 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, Neverfail Engine will first try to restart the application on the active server. If a restart of the application fails, then Neverfail 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** — Neverfail 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.

Neverfail Engine Plug-ins provide these monitoring and preemptive repair capabilities. Neverfail 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, Neverfail 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** — Neverfail Engine ensures the data files that applications or users require in the application environment are made available should a failure occur. Once installed, Neverfail 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.

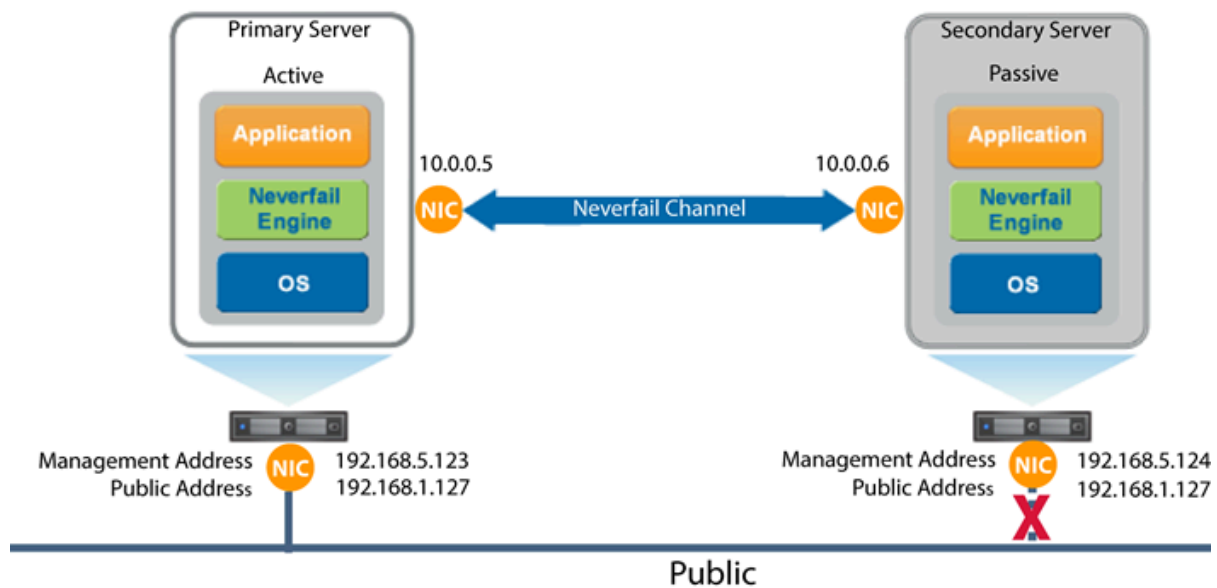
The passive server's receive queue is then read in numerical order and a duplicate set of file operations are applied to the disk of the passive server.

Neverfail 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.

3. Neverfail Continuity 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.

Figure 1-2. Neverfail Continuity Engine Network Configuration

**Note**

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

Neverfail Continuity 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 Continuity 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 100Mbps per second or more across standard Cat-5 cabling.

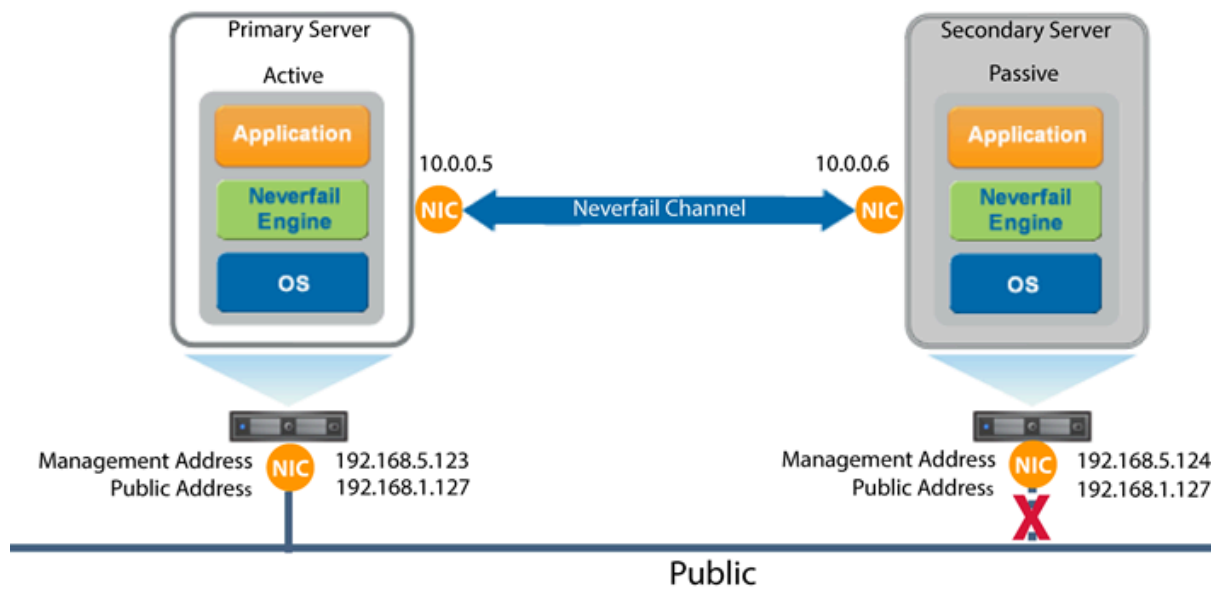
Note

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

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

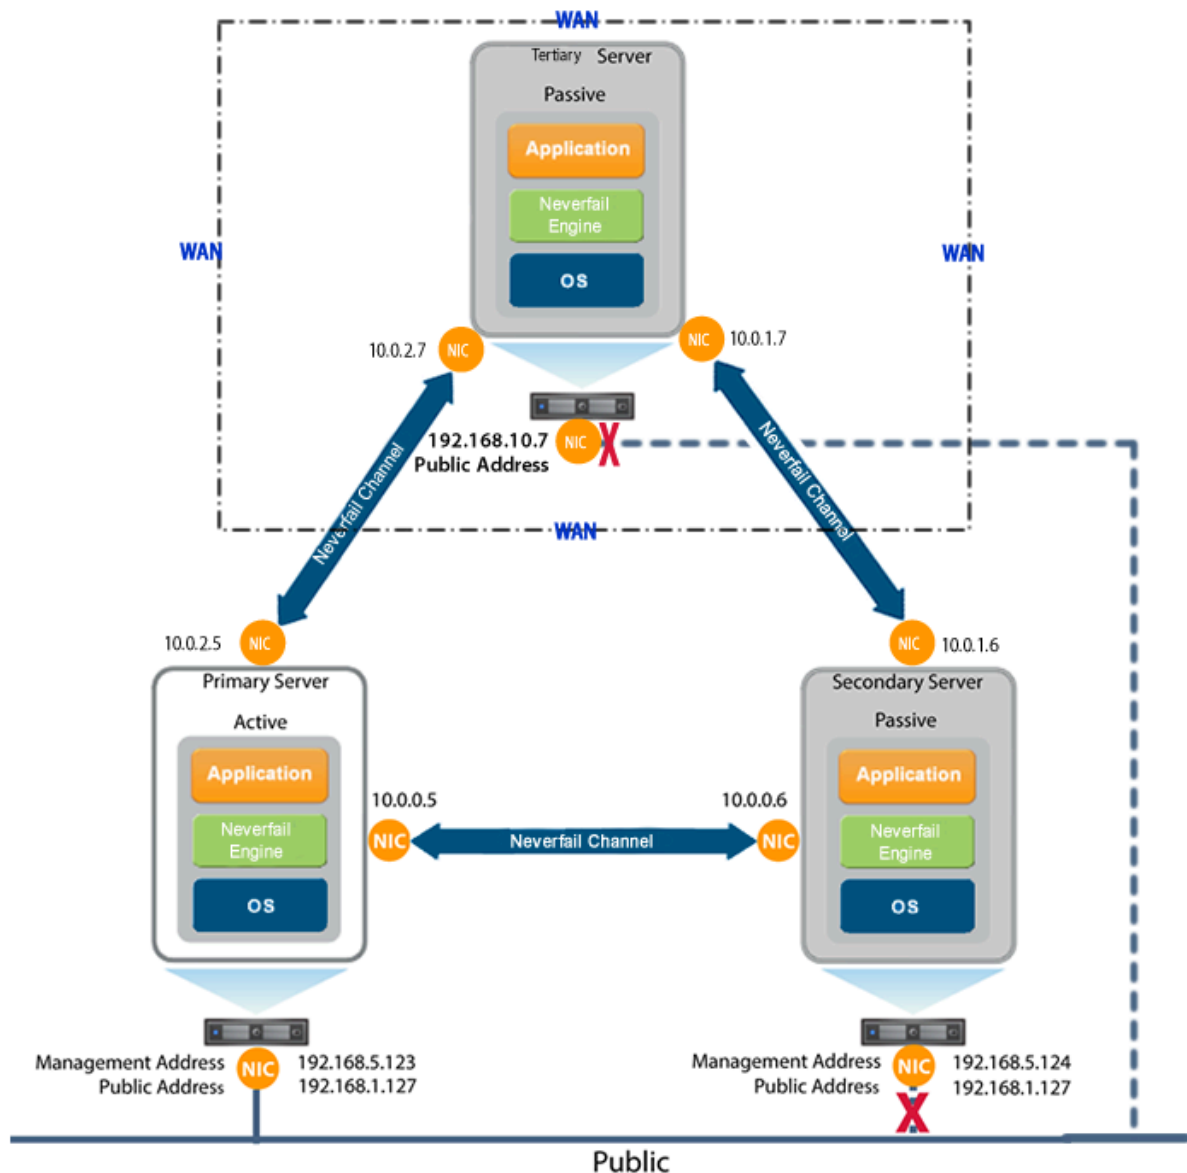
4. Neverfail Continuity Engine Communications

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.

Figure 1-3. Neverfail Continuity Engine Pair Communications**Note**

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

Figure 1-4. Trio Configuration



5. Neverfail Continuity Engine Switchover and Failover Processes

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

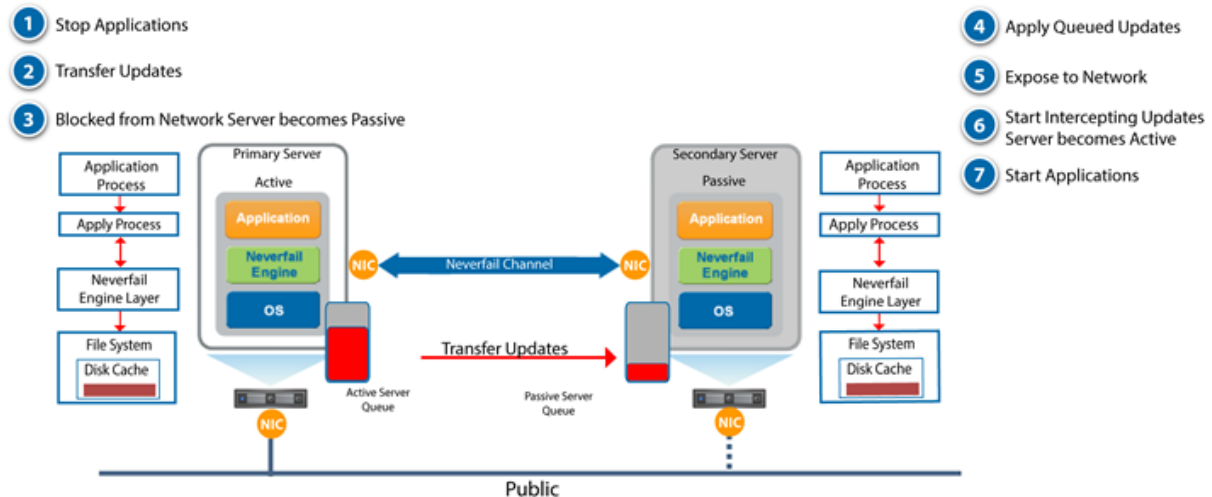
Note

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.

Figure 1-5. Neverfail Continuity Engine Switchover Process



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

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

The managed switchover is complete.

The Automatic Switchover Process

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

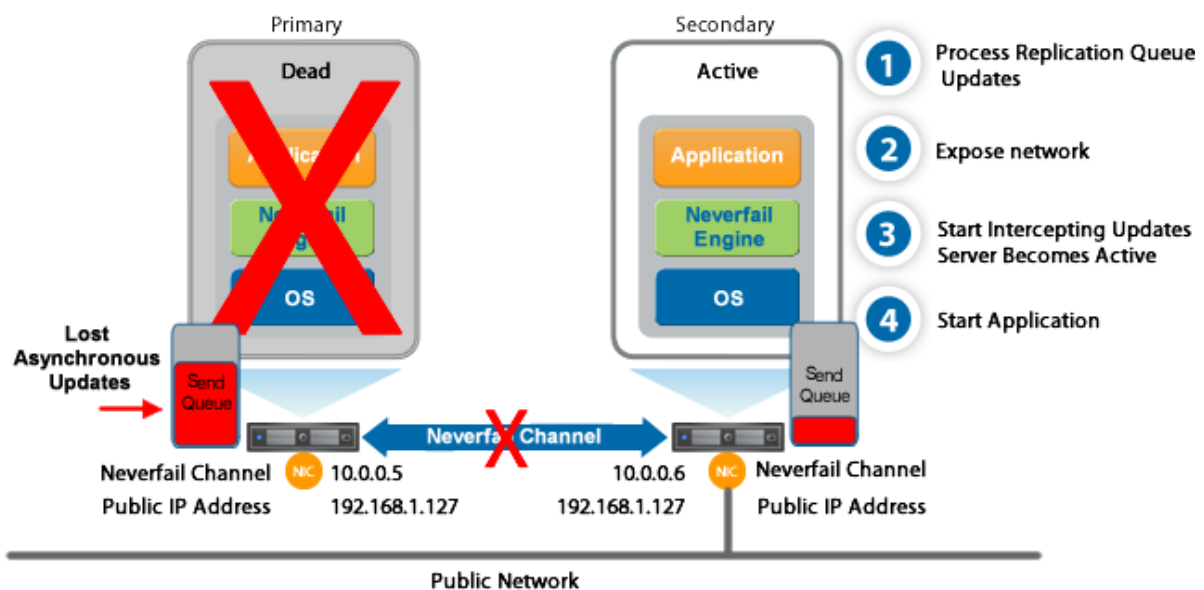
An auto-switchover is different from a managed switchover in that although the server roles are changed, Neverfail Engine is stopped on the previously active server to allow the administrator to verify the integrity of the data on the newly passive server and to investigate the cause of the auto-switchover.

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.

Figure 1-6. Neverfail Continuity Engine Failover Process



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

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

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

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

It enables the public identity of the server. The active and passive servers both use the same system name and same 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.

6. 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.

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 become passive and begin replication. If it connects to its peers and no other server is active, it will become active, and begin replication. 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 connect and start to resynchronize the data on the Primary 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.

Part I. Configuration

Chapter 2. Status and Control

Related information

- [Using the Engine Management Service User Interface](#)
- [Managing Neverfail Continuity Engine Clusters](#)
- [Review the Status of Neverfail Continuity Engine Clusters and Groups](#)
- [Exit Neverfail Advanced Management Client](#)
- [Shutdown Windows with Neverfail Continuity Engine Installed](#)
- [Controlled Shutdown](#)

1. Using the Engine Management Service User Interface

The Engine Management Service is the primary tool used for deployment and normal daily control of Neverfail Continuity Engine. Most routine operations can be performed from the Engine Management Service User Interface thereby providing a lightweight, easily accessible, method of conducting Neverfail Continuity Engine operations.

1.1. Configure Connection to VMware vCenter Server

Before you begin

The Configure Connection to VMware vCenter Server feature provides the ability to select and deploy Neverfail Engine on a powered-on VM, with VMtools running, from the vCenter inventory. Also, a VMware vCenter Server connection is required to automatically create a stand-by Secondary and/or Tertiary VM server from the cluster and place them on a specific Host/Datastore.

About this task

Configuring a connection to VMware vCenter Server:

Procedure

1. Click the **vCenter** button to display the *Configure Connection to VMware vCenter Server* page. The Configure vCenter section allows you to configure the connection to VMware vCenter Server or to remove any configured connection using the **Clear** button.
2. Enter the URL for the VMware vCenter Server, the username, and the password for a user account with the minimum privileges required by EMS to operate (see KB 2901), and then click **Next**.

Figure 2-1. Configure vCenter

Configure Connection to VMware vCenter Server

1) **Configure vCenter**
 2) Select default host
 3) Select default datastore
 4) Ready to complete

Enter the URL for the VMware vCenter Server

Enter the name of an account on the VMware vCenter Server

Enter the password for the account

☒ I want to select a default host and datastore for automated recloning

Account Privileges

Privilege	Allowed
Extension.Update	true
Task.Create	true
Extension.Register	true

VMware vCenter Server

VMware vCenter Server is used to create virtual Secondary Servers from VMware virtual Primary Servers. It is also used to create virtual Tertiary Servers from VMware virtual Secondary Servers.

The URL should have the format `https://vCenterServerFQDN/sdk`

For more information on required vSphere privileges, see article KB 2901

3. Select the **default host** used for *Recloning Secondary or Tertiary Server*. The *default host* configuration is used when recloning a node to a new location that has *not been defined* for the reclone operation and the EMS is *unable to retrieve the original reclone target location* from the Secondary or Tertiary servers.

Figure 2-2. Select Default Host

Configure Connection to VMware vCenter Server

1) Configure vCenter
 2) **Select default host**
 3) Select default datastore
 4) Ready to complete

Select a datacenter and host for the virtual machine


- 172.21.0.11
- 172.21.0.14
- 172.21.0.16
- 172.21.0.17
- 172.21.0.19
- 172.21.0.20**
- 172.21.0.21
- 172.21.0.24
- 172.21.0.25
- 172.21.0.26
- 172.21.0.27
- 172.21.0.51
- 172.21.0.52

Select Host

This host will be used for recloning operations where the host is not otherwise specified.

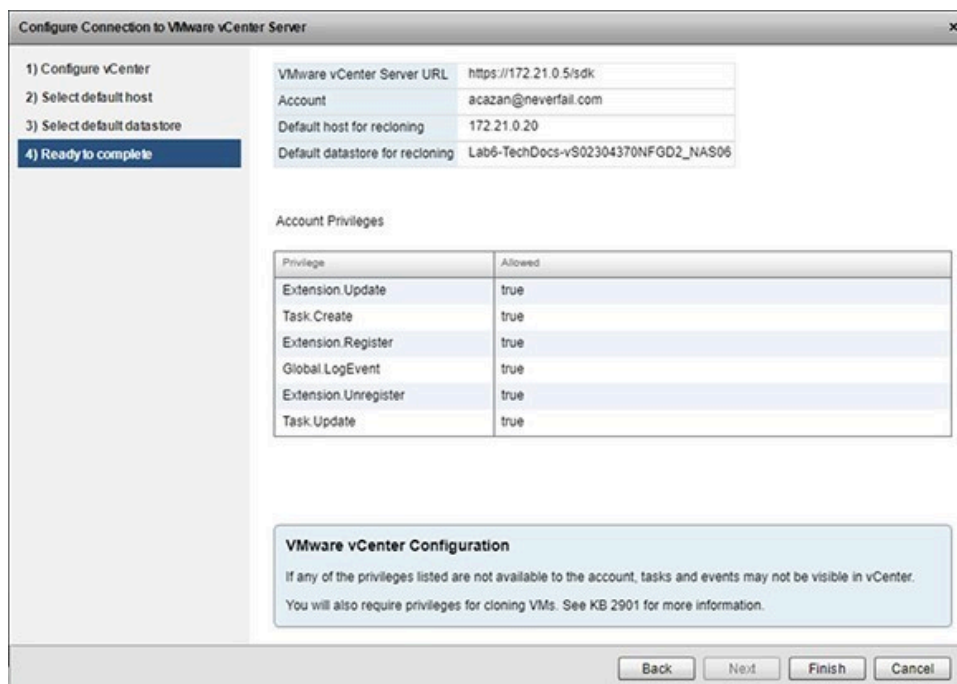
4. Select the **default storage** used for *Recloning Secondary or Tertiary Server* when the storage is *not specified* and the EMS *can not retrieve the storage location* from the Secondary or Tertiary servers.

Figure 2-3. Select Default Datastore



5. Review the information in the *Ready to Complete* dialog and then click **Finish**.

Figure 2-4. Ready to Complete



1.2. Configure VMware vCenter Converter

Before you begin

Use the *Configure VMware vCenter Converter* feature to convert physical Primary or VMs with a different hypervisor than ESXi to virtual Secondary and/or Tertiary servers during the automated cloning process used by Neverfail Continuity Engine Management Service to create the Secondary and/or Tertiary servers.

VMware vCenter Converter 5.5 or later must be installed manually.

About this task

To configure the VMware vCenter Converter:

Procedure

1. Click the **Converter** button to display the *Configure Connection to VMware vCenter Converter* page.

The *Configure Converter* section allows you to configure the connection to the VMware vCenter Converter or to remove any configured connection using the **Clear** button.

Figure 2-5. Configure VMware vCenter Converter

Configure Connection to VMware vCenter Converter

1) Configure Converter
2) Ready to complete

Enter the URL for the VMWare vCenter Converter

Enter the name of an administrator account on the VMware vCenter Converter server

Enter the password for the account

VMware vCenter Converter

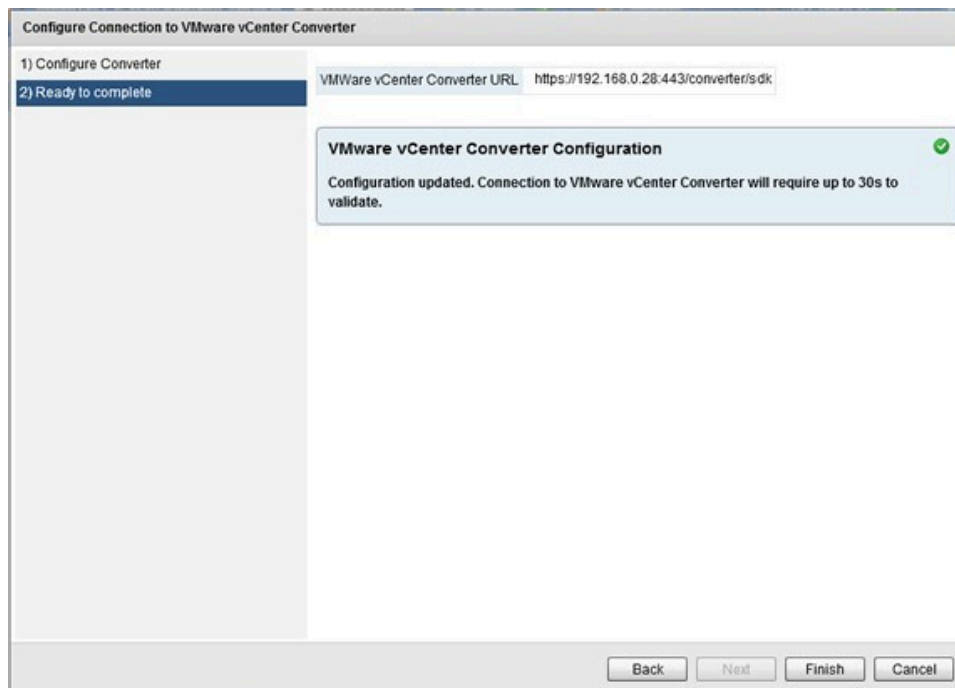
VMWare vCenter Converter is used to create virtual Secondary Servers from physical Primary Servers or VMs with a different hypervisor type.

VMware vCenter Converter installation must meet these requirements:

- 1) Be a supported version: versions 5.5, 6.0, 6.1.1 and 6.2 are supported
- 2) Installed in advanced (client/server) mode with remote access enabled
- 3) Have network visibility to Neverfail CE Management Service, vCenter Server and the target Primary server(s)
- 4) Where co-located with vCenter, the default port for converter is changed from 443

[Obtain VMware vCenter Converter](#)

2. Enter the URL to where VMware vCenter Converter resides.
3. Enter the Username and Password for an account with Administrator permissions on the VMware vCenter Converter server. Click **Next**.

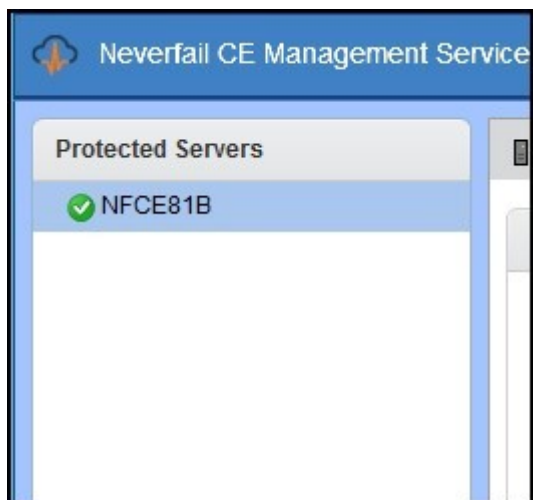
Figure 2-6. Ready to Complete

4. Click **Finish** to accept the configuration parameters.

1.3. Protected Servers

The *Protected Servers* pane provides a view of all servers that are currently protected by Neverfail Continuity Engine and managed by Neverfail Continuity Engine Management Service.

To view the status of a protected server, simply select the intended protected server.

Figure 2-7. Protected Servers

1.4. Management

The *Management* drop-down provides access to all of the key functions to deploy Neverfail Continuity Engine and get Neverfail Engine up and running. It provides the ability to Deploy, Manage, Integrate, and License Neverfail Engine.

1.4.1. Deploy

The Deploy group is focused on deployment actions and provides the functions to deploy Neverfail Continuity Engine as a Primary, Secondary, or Tertiary server.

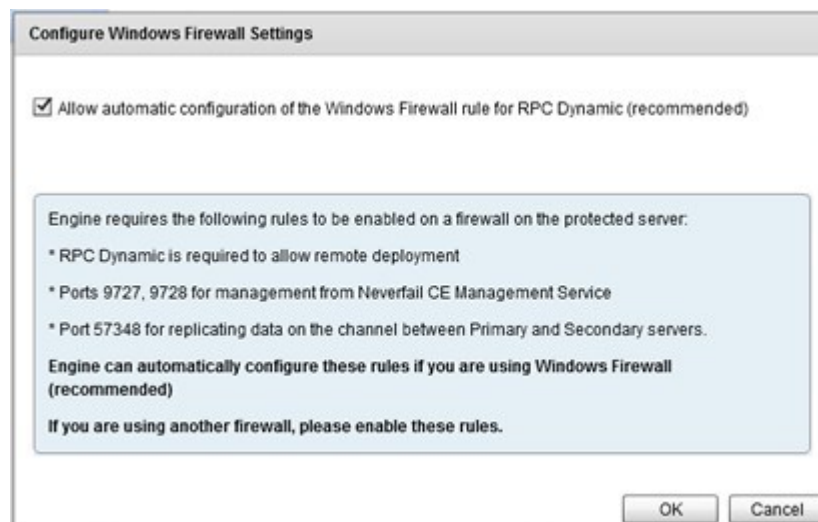
1.4.1.1. Configure Windows Firewall for Deployment

Neverfail Continuity Engine Management Service, by default, automatically configures Windows Firewall rules for RPC Dynamic (recommended). In the event that a non-Windows firewall is being used, you must manually configure firewall rules to allow for deployment and operations.

Configure the following firewall rules:

- RPC Dynamic is required to allow remote deployment.
- Ports 9727, 9728 for management from Neverfail Continuity Engine Management Service.
- Port 57348 for replicating data via the Neverfail Channel between the Primary and Secondary servers.

Figure 2-8. Configure Windows Firewall Settings



1.4.1.2. Deploy to a Primary Server

Before you begin

When this option is selected, Neverfail Engine is installed onto the Primary server.

Prior to attempting installation of Neverfail Engine on the Primary server, ensure that the server meets all of the pre-requisites stated in the **Pre-Install Requirements** section of the Neverfail Engine Installation Guide.

Important

Neverfail Engine requires that Microsoft™ .Net Framework 4 be installed prior to Neverfail Engine installation. If .Net Framework 4 is not installed, Neverfail Engine will prevent installation until .Net Framework 4 is installed.

About this task

To Deploy Neverfail Engine:

Procedure

1. Having verified all of the environmental prerequisites are met, click on **Management** and navigate to **Deploy > Deploy to a Primary Server**.

When deploying a Primary server, use an account with full administrator permissions to successfully deploy the Primary server.

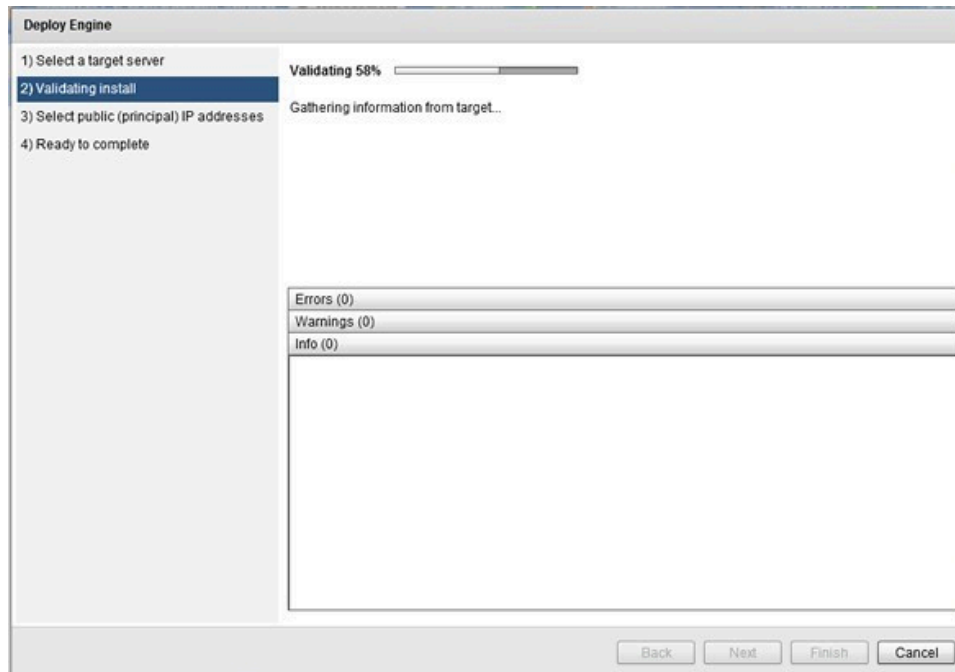
The *Deploy Engine* page is displayed.

Figure 2-9. Deploy Neverfail Engine step

2. Enter the DNS name or IP address of the server that will be the Primary server, or select a virtual server from the inventory. Enter credentials for a user account with full administrator permissions on the target server and click **Next**.


The *Validating Install* step is displayed. Neverfail Engine automatically configures Windows firewalls to allow installation to continue and communications via the Neverfail Channel and the Neverfail Continuity Engine Management Service.

Figure 2-10. Validating Install step



3. Once the Validating Install dialog completes and displays that the server is a valid target, click **Next**. The *Select public (principal) IP addresses* step is displayed.


Figure 2-11. Select public (principal) IP addresses step



4. Verify that the proper IP address for the Public IP address is configured/selected and that the check box is selected. Click **Next**.

The *Ready to complete* step is displayed.

Figure 2-12. Ready to complete step



5. Review the information and click **Finish**.

The installation of the Primary server proceeds.

6. Once installation of the Primary server is complete, in the *Protected Servers* pane, select the Primary server to display the *Server Summary* page .

1.4.1.3. Upgrade the Selected Server

Before you begin

Neverfail Continuity Engine Management Service provides a simple process incorporating a wizard to upgrade from previous versions of the product.

Procedure

1. From the **Management** drop-down, navigate to **Deploy > Upgrade the selected server**.

The *Upgrade Engine* page is displayed.

Figure 2-13. Upgrade Engine

Upgrade Engine

1) Provide credentials
2) Validating upgrade
3) Ready to complete

Server cluster to upgrade: lj-psc67U1.eng.cj

Enter the name of a local administrator account for the server

Enter the password for the account

☐ I confirm that no users are logged on to the Primary, Secondary (or Tertiary) Servers

☒ Upgrade all server nodes in cluster (recommended)
☐ Upgrade only a specific server in the cluster

IP address specific to server node (management IP address)

Upgrade

Account and User Access Control (UAC)

If you have UAC enabled on the target server, you must use the built-in local Administrator account.

If UAC is not enabled, you may use any local account with membership in the local Administrators group on the target server.

All server nodes (Primary, Secondary and Tertiary if relevant) will be upgraded, unless a single node upgrade is selected.

Single node upgrades should only be used where upgrade of the whole cluster has failed, for example because of connection loss during upgrade. Single node upgrades require a unique management IP address assigned to the node.

Please see KB 2886 before using single node upgrade

Back Next Finish Cancel

2. Enter the name of the local Administrator account and password. After confirming that no users are logged into the Primary, Secondary (or Tertiary) servers, select the check box.

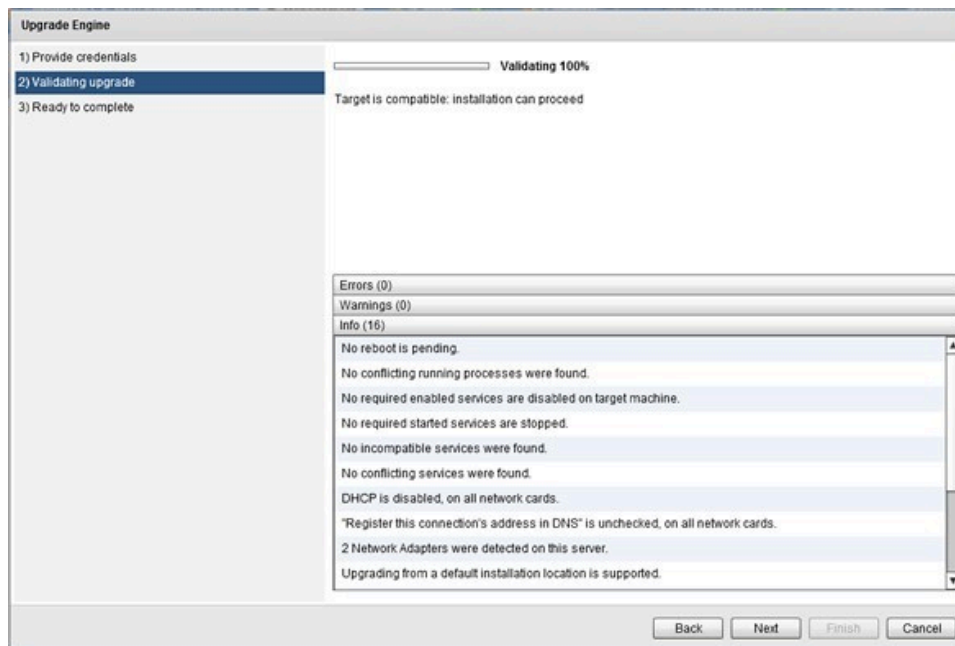
It is mandatory to have the same the local Administrator account and password on all of the Engine nodes in the cluster, for a successful full cluster upgrade.

3. Select to either upgrade all server nodes or only a specific server in the cluster. Click **Next**.

Single node upgrades should only be used in the event the upgrade of the whole cluster has failed. If you select to upgrade only a specific server in the cluster, you must configure a Management IP address on the target server prior to attempting the upgrade. A new instance will then be added in the Protected Servers list represented by the management IP.

The *Validating upgrade* step is displayed.

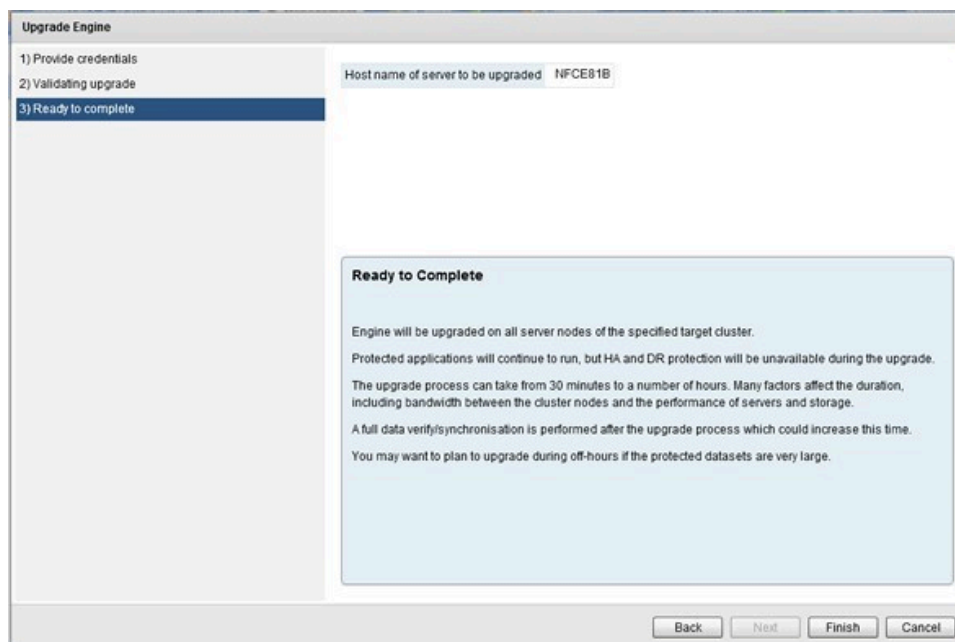
Figure 2-14. Validating upgrade step



4. Once validation is complete, click **Next**.

The *Ready to complete* step is displayed.

Figure 2-15. Ready to complete step



5. Review the information and click **Finish** to initiate the upgrade of the selected cluster or single server.

1.4.1.4. Uninstall from the Selected Server

About this task

The Neverfail Continuity Engine Management Service allows you to uninstall Neverfail Engine from a selected cluster.

To uninstall from the selected server:

Procedure

1. Select the intended server and from the **Management** drop-down, navigate to **Deploy > Uninstall from the Selected Server**.

The *Uninstall Engine* step is displayed.

Figure 2-16. Uninstall Engine

2. Select one of the available (and applicable) uninstall options for Secondary (and Tertiary - if present).
 - Delete VM (Recommended, requires vCenter) - this option will delete the VM.
 - Reconfigure host name and IP address - specify the new host name for the formerly passive server.

This option is only available if you attempt to uninstall a v8.1 or later cluster from Neverfail Continuity Engine Management Service v8.1 (or later).

3. Choose one of the available options:

- **Disable NICs** - this option will uninstall Engine and disable all the existing NICs on the formerly passive server. The server will be shutdown and removed from the domain if it was previously a domain member.
- **Change Public IP address** - this option will uninstall Engine then configure the newly specified IP address on the formerly passive server. The server will be left running.

In both cases, the passive server(s) will be removed from the domain.

4. After verifying that no users are logged onto the Primary, Secondary, or Tertiary (if installed) servers, select the confirmation check box and provide the local Administrator account valid on all servers. Click **OK**.

The Uninstall Validation process will start. If no issues are found, Neverfail Engine is uninstalled from the Primary, Secondary and Tertiary (if installed) servers.

1.4.1.5. Add a Stand-by Server for High Availability

About this task

The *Add a stand-by server for high availability* feature is used to create a Secondary server when deployed for high availability. Deploying for high availability means that failover will occur automatically when the active server fails. This feature can also be used to add a stand-by server for high availability to an existing disaster recovery pair. In this case, the new server will become the Secondary server and the existing Secondary/DR server will be re-labeled as the Tertiary.

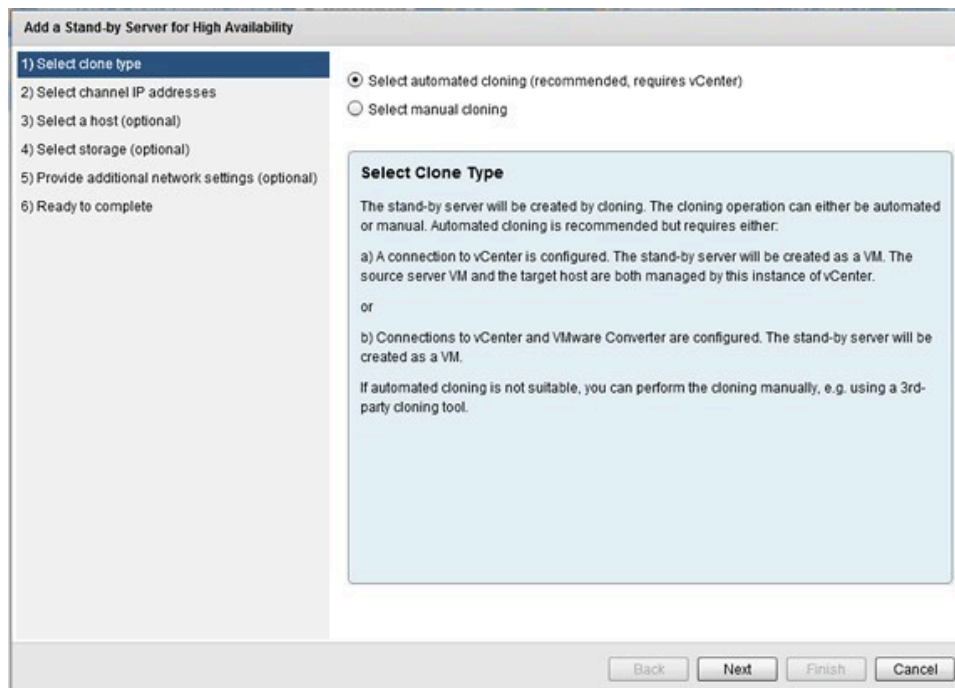
To add a stand-by VM for high availability:

Procedure

1. On the Neverfail Continuity Engine Management Service user interface, click the **Management** drop-down and navigate to **Deploy > Add a stand-by Server for high availability**.

The *Add a Stand-by Server for High Availability* page is displayed.

2. Select clone type – select to use either automated cloning (recommended) or manual (using a third-party cloning tool) to clone a specific server. Click **Next**.

Figure 2-17. Select Clone Type step

The *Select channel IP addresses* step is displayed.

3. Select the NIC which is to host the Channel IP addresses. Enter the Channel IP addresses for the Primary and Secondary servers. Manually enter the subnet mask or leave blank to set to the default subnet mask. If you are adding high-availability to an existing DR pair, enter the IP addresses and associated information for the Secondary-Tertiary and Tertiary-Primary (when deployed) Channel. Click **Next**.

If the IP addresses chosen are not already present on the server's NICs, they will be added automatically.

Figure 2-18. Select Channel IP Address step

The screenshot shows a wizard window titled "Add a Stand-by Server for High Availability". On the left, a list of steps is shown: 1) Select clone type, 2) Select channel IP addresses (highlighted), 3) Select a host (optional), 4) Select storage (optional), 5) Provide additional network settings (optional), and 6) Ready to complete. The main area is divided into two tabs: "Primary server to Secondary server" (selected) and "HA Stand-by server to DR Stand-by server". Under the selected tab, there are fields for "Select a network adapter for the channel" (set to "Channel"), "Enter an IPv4 address for the Primary server" (192.168.5.5), "Subnet mask for the Primary server (blank for default)" (255.255.255.0), "Enter an IPv4 address for the Secondary server" (192.168.5.6), and "Subnet mask for the Secondary server (blank for default)" (255.255.255.0). Below these fields is a section titled "Channel IP Addresses" with the text: "The addresses will be automatically added to each server to allow Engine to communicate and replicate data." At the bottom right are buttons for "Back", "Next", "Finish", and "Cancel".

The *Select a host (optional)* step is displayed.

4. Select the Datacenter and Host where the Secondary server will be created and click **Next**.

If the Primary server is a virtual machine, then the Secondary server should be on a separate host to protect against host failure.

Figure 2-19. Select Host step

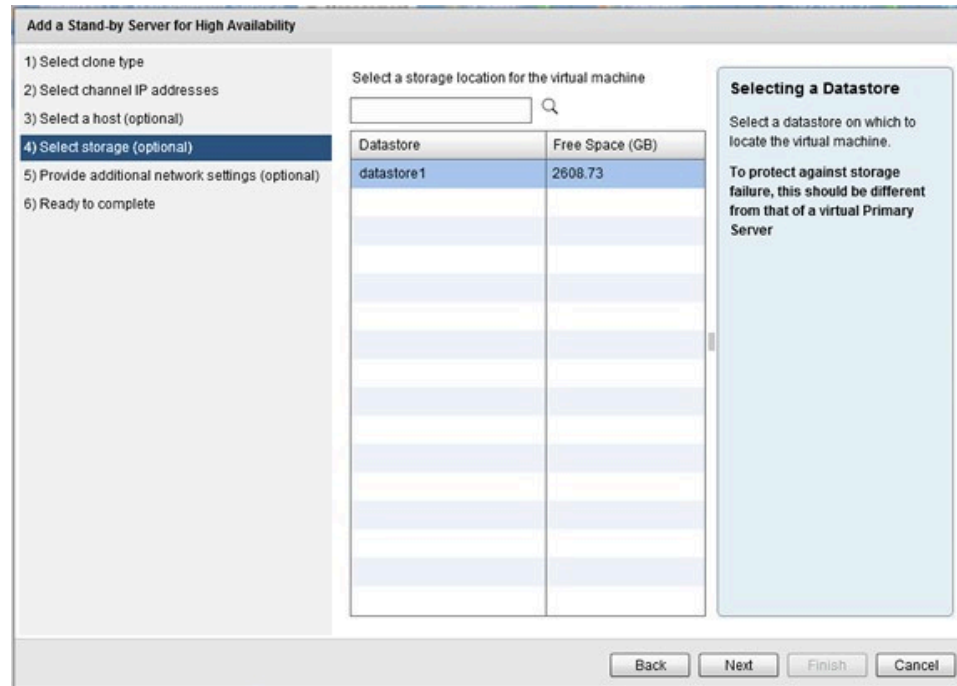
The screenshot shows the same wizard window, now at Step 3: "Select a host (optional)". The left sidebar highlights this step. The main area is titled "Select a datacenter and host for the virtual machine". It features a search bar and a list box showing "Datacenter1" expanded, with a sub-entry "192.168.0.141" selected. To the right of the list box is a section titled "Selecting a Host" with the following text: "A new VM will be created to provide High Availability. To protect against server failure, this should be a separate host from a Primary VM. To provide high-availability, it should have a reliable, high-bandwidth connection with the Primary." At the bottom right are buttons for "Back", "Next", "Finish", and "Cancel".

The *Select Host* step is displayed.

5. Select a storage location for the virtual machine. Click **Next**.

The option to provide additional network settings is not available if Engine is deployed on a Windows based server.

Figure 2-20. Select Storage step



The *Select Storage* step is displayed.

6. Click **Finish** to initiate installation of the Secondary server.

Once installation of the Secondary server is complete, automatic reconfiguration of the Secondary server will take place requiring only a few minutes to complete.

Figure 2-21. Ready to Complete step

Add a Stand-by Server for High Availability

1) Select clone type
2) Select channel IP addresses
3) Select a host (optional)
4) Select storage (optional)
5) Provide additional network settings (optional)
6) Ready to complete

Primary VM name	NFCE81B
Primary channel IP address	192.168.5.5
Subnet mask	255.255.255.0
Secondary channel IP address	192.168.5.6
Subnet mask	255.255.255.0
Cloning mechanism	Automatic
Datacenter for Secondary server	Datacenter1
Host for HA Stand-by server	192.168.0.141
Datastore for HA Stand-by server	datastore1

Ready to Complete

The VM will be cloned to the specified location.

Cloning may take some time, depending on volume of data and available bandwidth.

Once the cloning has completed, the servers will begin replicating automatically.

Back Next Finish Cancel

The *Ready to complete* step is displayed.

- Once complete, perform *Post Installation Configuration* tasks listed in this guide.

1.4.1.6. Add a Stand-by Server for Disaster Recovery

About this task

The *Add a stand-by server for disaster recovery* feature is used to create a Secondary server when deployed for disaster recovery. A Secondary server created for disaster recovery will typically be located at a different site from that of the Primary server. By default, automatic failover is disabled between the active and passive servers. This feature can also be used to add a stand-by server for disaster recovery to an existing high availability pair.

To add a stand-by server for disaster recovery:

Procedure

- On the Neverfail Continuity Engine Management Service user interface, click the **Management** drop-down and navigate to **Deploy > Add a stand-by server for Disaster Recovery**.

The *Add a stand-by server for disaster recovery* page is displayed.

- Select either of the following and click **Next**:
 - The public (principal) IP address will be identical to the Primary server.
 - The public (principal) IP address will be different than the Primary server. Here you can set up the new public IP address (IP, subnet mask, gateway and preferred/alternate DNS servers). You must also add credentials to be used for updating DNS.

Figure 2-22. Select Public IP Address step

Add a Stand-by Server For Disaster Recovery

1) Select public IP address
 2) Select channel IP addresses
 3) Select clone type
 4) Select host (optional)
 5) Select storage (optional)
 6) Configure helper VM (optional)
 7) Ready to complete

☐ The public (principal) IP address will be identical to the Primary server
☒ The public (principal) IP address will be different than on the Primary server

Network adapter: **Ethernet1**

IP address	Subnet mask	
192.168.2.99	255.255.255.0	Add...
		Remove

Enter the gateway: 192.168.2.1

Enter the preferred DNS server: 192.168.2.2

Enter the alternate DNS server (optional):

Enter the user name for updating DNS servers: administrator

Enter the password: *****

Public IP Addresses
 If the Primary and DR site use different subnets, the DR server requires a separate public IP address.
 In this case, an account capable of updating the DNS servers must be specified.
 On switchover or failover, DNS servers will then be updated with the IP address of the active server.

Back Next Finish Cancel

The *Select Channel IP Addresses* step is displayed.

- Enter the Neverfail Channel IP addresses for the Primary and Secondary servers. Manually enter the subnet mask or leave blank to set to the default subnet mask. If you are adding Disaster Recovery to an existing pair, then enter the IP Addresses and associated information for the Primary-Tertiary and Secondary-Tertiary channels. Click **Next**.

Figure 2-23. Select Channel IP Addresses step

Add a Stand-by Server For Disaster Recovery

1) Select public IP address
2) Select channel IP addresses
 3) Select clone type
 4) Select host (optional)
 5) Select storage (optional)
 6) Configure helper VM (optional)
 7) Ready to complete

Primary server to Secondary server | Secondary server to Tertiary server | Tertiary server to Primary server

Select a network adapter for the channel: **Ethernet1**

Enter an IPv4 address for the existing Secondary: 10.0.1.5

Secondary Subnet Mask (blank for default): 255.255.255.0

Enter an IPv4 address for the Tertiary (new server): 10.0.2.6

Tertiary Subnet Mask (blank for default): 255.255.255.0

Channel IP Addresses
 The addresses will be automatically added to each server to allow Engine to communicate and replicate data.
 A persistent static route should be configured for the channel connection where routing is required

Back Next Finish Cancel

The *Select Clone Type* step is displayed.

4. Select whether to clone the Primary server to create a Secondary server and power-on the Secondary server or to clone the Primary server to create the .vmdk files to be ported manually to the DR site. Additionally, you can select to perform a manual clone using a third-party cloning tool to clone a specific server. Click **Next**.

If you have selected to move the .vmdk files, this refers to where the files will be created, not the final destination.

Figure 2-24. *Select Clone Type* step

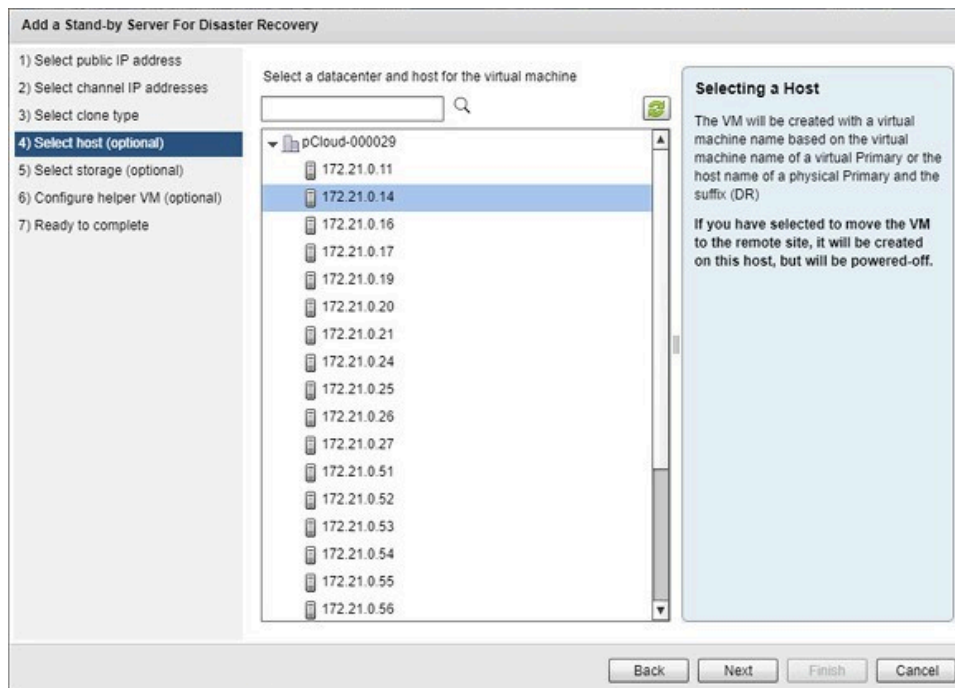
The screenshot shows a wizard window titled "Add a Stand-by Server For Disaster Recovery". On the left, a list of steps is shown: 1) Select public IP address, 2) Select channel IP addresses, 3) Select clone type (highlighted), 4) Select host (optional), 5) Select storage (optional), 6) Configure helper VM (optional), and 7) Ready to complete. The main area contains three radio button options: "Select automated cloning (recommended, requires vCenter)" (selected), "Create and power-on the stand-by VM automatically after cloning" (selected), and "Create a temporary powered-off clone locally, so that the VMDK files can be transferred manually" (unselected). Below these is an option for "Select manual cloning" (unselected). A large text box titled "Automated Cloning" explains that if the DR stand-by server is a VMware VM, automated cloning is the simplest option. It mentions that if you have a reliable, high-bandwidth connection to a remote site, you can create the DR VM directly on its host. Alternatively, you can create the DR VM in a temporary location on a local host and then transfer the VMDK files to the remote site. Below this is a section titled "Manual Cloning" which states that if the DR stand-by server is another type of VM or a physical machine, manual cloning can be used. At the bottom right, there are four buttons: "Back", "Next", "Finish", and "Cancel".

The *Select Host* step is displayed.

5. Select a Datacenter and Host for the virtual machine. Click **Next**.

If you have selected to move the .vmdk files, this refers to where the files will be created, not the final destination.

Figure 2-25. Select Host step

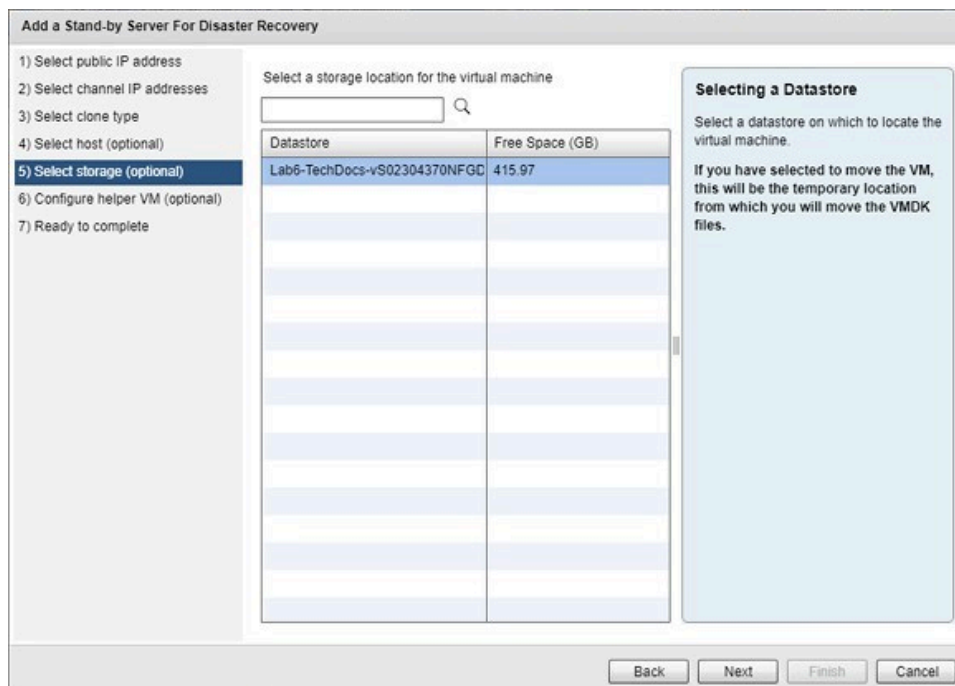


The *Select Storage* step is displayed.

6. Select the storage location for the virtual machine. Click **Next**.

The option to Configure helper VM (optional) is not available if Engine is deployed on a Windows based server.

Figure 2-26. Select Storage step



- Review the information on the *Ready to Complete* step and if accurate, click **Finish** to create the Secondary server.

Figure 2-27. Ready to Complete step

Add a Stand-by Server For Disaster Recovery

- 1) Select public IP address
- 2) Select channel IP addresses
- 3) Select clone type
- 4) Select host (optional)
- 5) Select storage (optional)
- 6) Configure helper VM (optional)
- 7) Ready to complete**

Primary server	NFCE81B
Cloning mechanism	Automatic
Secondary Datacenter	Datacenter1
Secondary Host	192.168.0.141
Secondary Datastore	datastore1
Public IP addresses	192.168.0.15
Gateway	192.168.0.1
Preferred DNS server	192.168.0.2
Alternate DNS server	
Primary channel IP address	10.0.0.5
Subnet mask	255.255.255.0
Secondary channel IP address	10.0.1.6
Subnet mask	255.255.255.0

Ready to complete

The DR VM will be cloned to the specified location.

Cloning may take some time, depending on volume of data and available bandwidth.

If you have selected to move the VM, once the cloning has completed, copy the VMDK files to the remote site, and power-on the VM.

Otherwise, once the cloning has completed, the servers will begin replicating automatically.

Back Next Finish Cancel

1.4.1.7. Create Secondary and Tertiary stand-by VMs for HA and DR


About this task

This feature works to extend capabilities of Neverfail Continuity Engine to incorporate both High Availability and Disaster Recovery by deploying both a Secondary server (for HA) and a Tertiary server (for DR).

To deploy Secondary and Tertiary VMs for High Availability and Disaster Recovery:

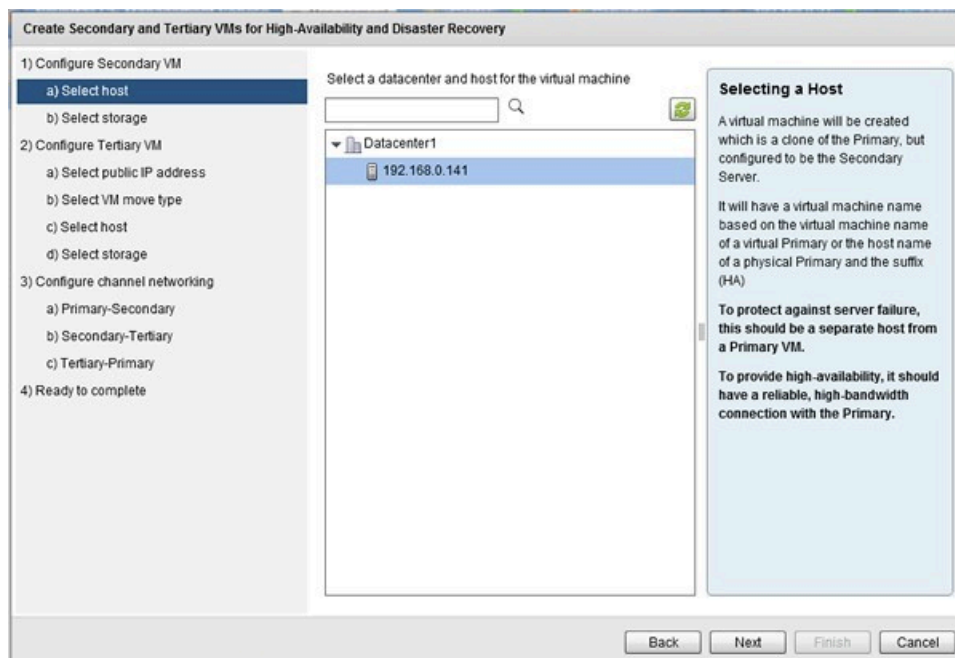
Procedure

- On the Neverfail Continuity Engine Management Service, navigate to the **Management > Deploy** drop-down and select **Create Secondary and Tertiary stand-by VMs for HA and DR**.

Figure 2-28. Configure Secondary VM step

The *Create Secondary and Tertiary VMs for High Availability and Disaster Recovery* page is displayed.

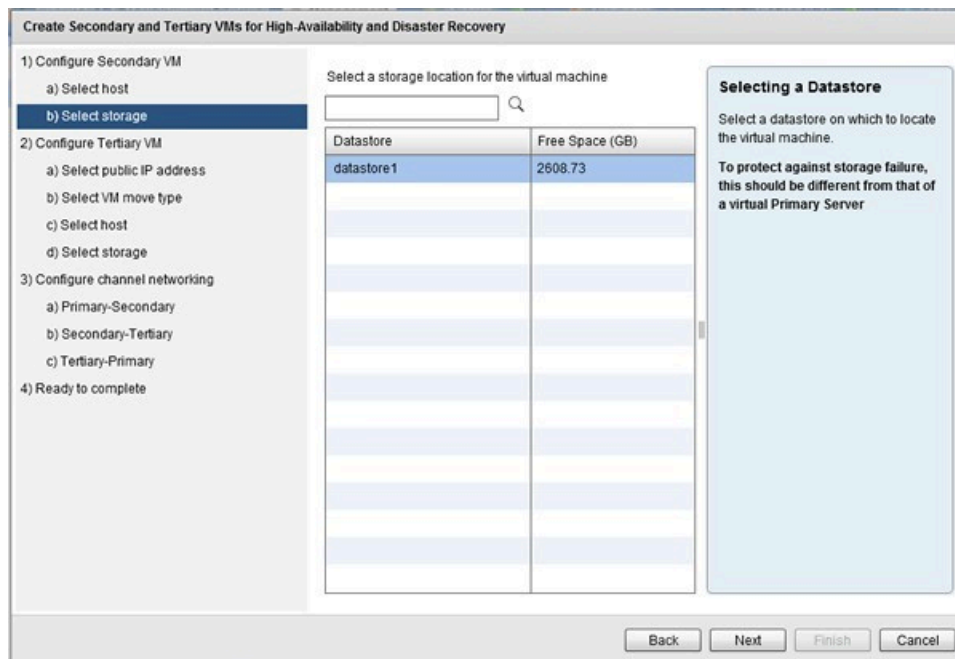
2. Review the information in the step and then click **Next**.

Figure 2-29. Select host step

The *Select host* step is displayed.

3. Click on the appropriate Datacenter to display all available hosts. Select the intended host for the Secondary server and then click **Next**.

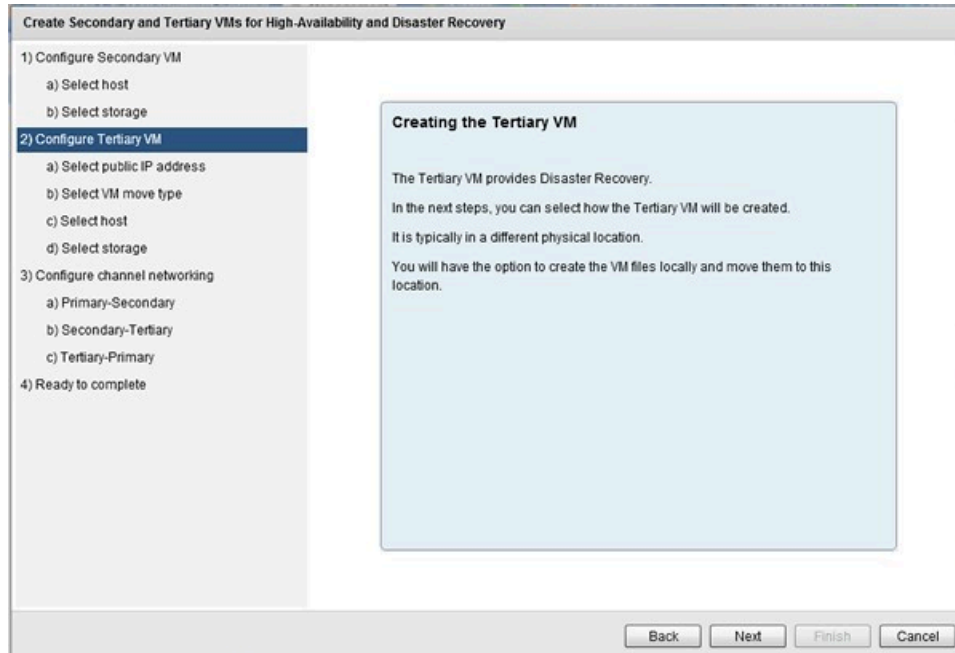
Figure 2-30. Select storage step



The *Select storage* step is displayed.

4. Select the intended datastore for the Secondary VM, and then click **Next**.

Figure 2-31. Configure Tertiary VM step



The *Configure Tertiary VM* step is displayed.

5. Review the contents of the step and then click **Next**.

Figure 2-32. Select public IP address step

Create Secondary and Tertiary VMs for High-Availability and Disaster Recovery

1) Configure Secondary VM

- a) Select host
- b) Select storage

2) Configure Tertiary VM

- a) Select public IP address**
- b) Select VM move type
- c) Select host
- d) Select storage

3) Configure channel networking

- a) Primary-Secondary
- b) Secondary-Tertiary
- c) Tertiary-Primary

4) Ready to complete

☐ The public (principal) IP address will be identical to the Primary Server

☒ The public (principal) IP address will be different than on the Primary Server

Network adapter: **Ethernet1**

IP address	Subnet mask
192.168.2.35	255.255.255.0

Buttons: Add... Remove

Enter the Gateway: 192.168.2.1

Enter the Preferred DNS Server: 192.168.2.2

Enter the Alternate DNS Server (optional):

Enter the user name for updating DNS Servers: administrator

Enter the password: *****

Public IP Addresses

If the Primary and DR site use different subnets, the Tertiary server requires a separate public IP address.

In this case, an account capable of updating the DNS servers must be specified.

On switchover or failover, DNS servers will then be updated with the IP address of the active server.

Buttons: Back Next Finish Cancel

The *Select public IP address* step is displayed.

6. If the public IP address will be different than the Primary server, select which NIC this should be assigned to and add a static IP address and the subnet mask. Additionally, add the Gateway IP, Preferred DNS server IP, and the user name and password of an account used for updating DNS servers. Click **Next**.

Figure 2-33. Select VM move type step

Create Secondary and Tertiary VMs for High-Availability and Disaster Recovery

1) Configure Secondary VM

- a) Select host
- b) Select storage

2) Configure Tertiary VM

- a) Select public IP address
- b) Select VM move type**
- c) Select host
- d) Select storage

3) Configure channel networking

- a) Primary-Secondary
- b) Secondary-Tertiary
- c) Tertiary-Primary

4) Ready to complete

☒ Create and power-on the Tertiary VM automatically after cloning

☐ Create a temporary powered-off clone locally, so that the VMDK files can be transferred manually

VM Move Type

If you have a reliable, high-bandwidth connection to remote site, you can choose to create the Tertiary 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 Tertiary VM in a temporary location on a local host. The VM will not be powered-on.

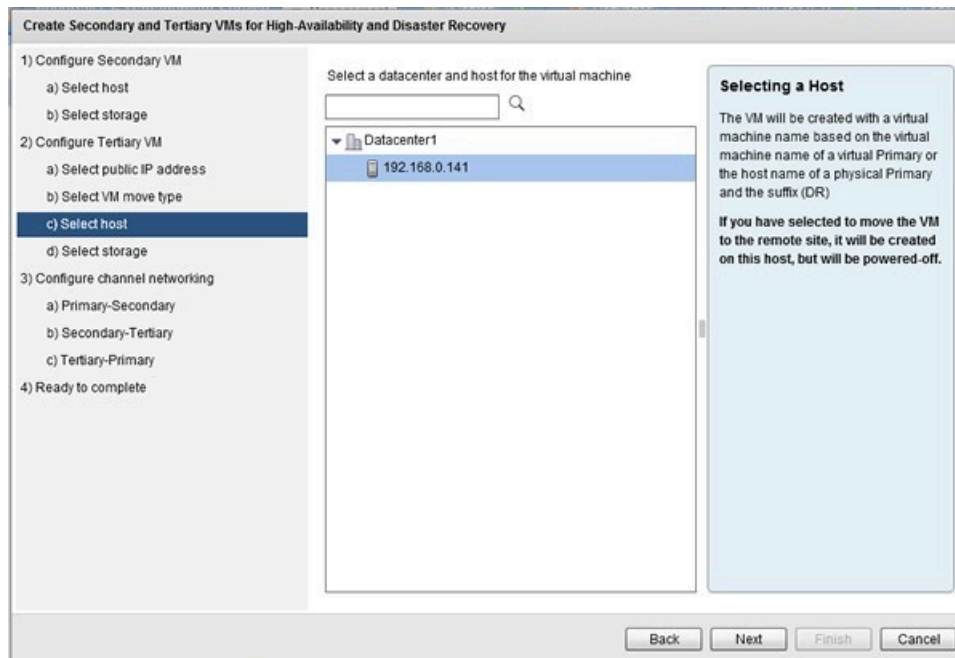
You can then transfer the VMDK files to the remote site, e.g. using detachable storage or FTP.

Buttons: Back Next Finish Cancel

The *Select VM move type* step is displayed.

7. Review the definitions of the options and then select whether the VM will be transferred manually or not. Click **Next**.

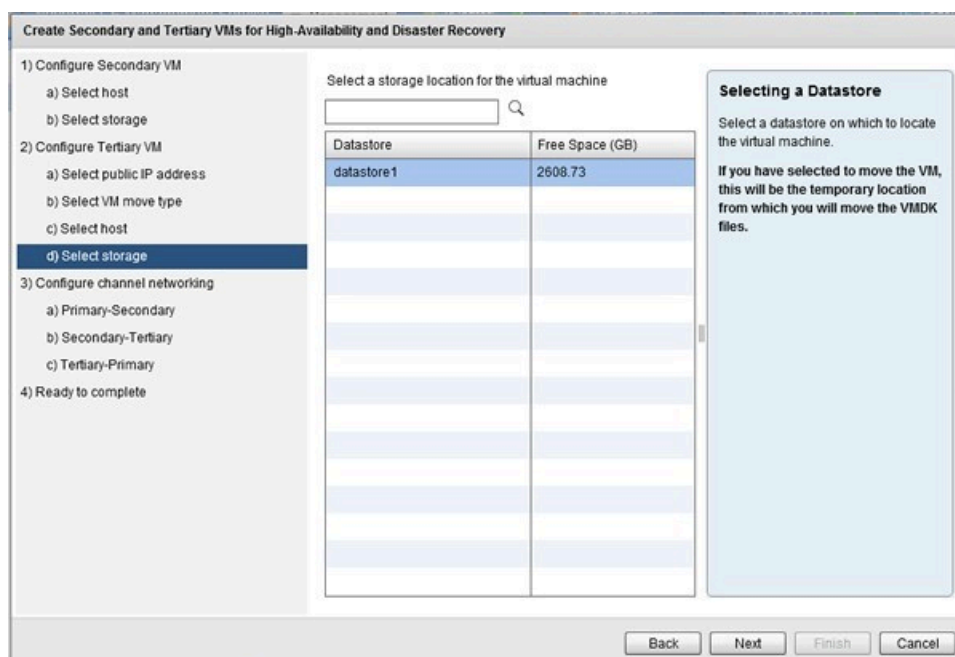
Figure 2-34. Select host step



The *Select host* step is displayed.

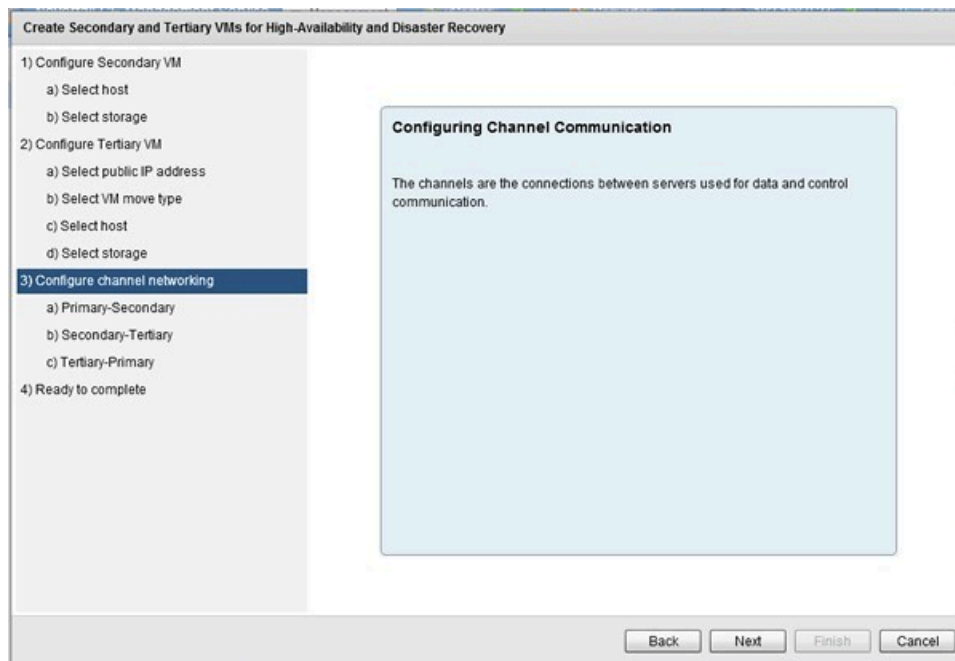
8. Click on the appropriate Datacenter to display all available hosts. Select the intended host for the Tertiary server and then click **Next**.

Figure 2-35. Select storage step




The *Select storage* step is displayed.

9. Select the intended datastore for the Tertiary VM, and then click **Next**.

Figure 2-36. Configure channel networking step

The *Configuring Channel Communications* step is displayed.

10. Review the contents of the step and then click **Next**.

Figure 2-37. Primary-Secondary step

The *Primary-Secondary* step is displayed.

11. Select the appropriate network adapter and then enter the channel IP addresses for Primary-Secondary communications. Click **Next**.

Figure 2-38. Secondary-Tertiary step

The screenshot shows the 'Create Secondary and Tertiary VMs for High-Availability and Disaster Recovery' wizard. The left pane shows the progress: 1) Configure Secondary VM (a) Select host, b) Select storage; 2) Configure Tertiary VM (a) Select public IP address, b) Select VM move type, c) Select host, d) Select storage; 3) Configure channel networking (a) Primary-Secondary, **b) Secondary-Tertiary**, c) Tertiary-Primary; 4) Ready to complete. The main pane is for the 'Secondary-Tertiary Channel'. It has a dropdown 'Select a network adapter for this channel' set to 'Public'. Below are input fields: 'Enter an IPv4 address for the Secondary' (10.0.1.6), 'Secondary Subnet Mask (blank for default)' (255.255.255.0), 'Enter an IPv4 address for the Tertiary' (10.0.1.7), and 'Subnet Mask (blank for default)' (255.255.255.0). A right-hand box titled 'Secondary-Tertiary Channel' contains text: 'The addresses will be automatically added to each server to data replication and cluster communication' and 'A persistent static route should be configured for the channel connection where routing is required'. At the bottom are 'Back', 'Next', 'Finish', and 'Cancel' buttons.

The *Secondary-Tertiary* step is displayed.

12. Select the appropriate network adapter and then enter the channel IP addresses for Secondary-Tertiary communications. Click **Next**.

Figure 2-39. Tertiary-Primary step

The screenshot shows the same wizard, but now the 'Tertiary-Primary' step is selected in the left pane. The main pane is for the 'Tertiary-Primary Channel'. The dropdown 'Select a network adapter for this channel' is set to 'Channel'. Input fields include: 'Enter an IPv4 address for the Tertiary' (10.0.2.7), 'Tertiary Subnet Mask (blank for default)' (255.255.255.0), 'Enter an IPv4 address for the Primary' (10.0.2.5), and 'Primary Subnet Mask (blank for default)' (255.255.255.0). The right-hand box titled 'Tertiary-Primary Channel' contains text: 'The addresses will be automatically added to each server to allow Engine to communicate and replicate data.' and 'A persistent static route should be configured for the channel connection where routing is required'. At the bottom are 'Back', 'Next', 'Finish', and 'Cancel' buttons.

The *Tertiary-Primary* step is displayed.

13. Select the appropriate network adapter and then enter the channel IP addresses for Tertiary-Primary communications. Click **Next**.

Figure 2-40. Ready to complete step

Create Secondary and Tertiary VMs for High-Availability and Disaster Recovery

1) Configure Secondary VM
a) Select host
b) Select storage
2) Configure Tertiary VM
a) Select public IP address
b) Select VM move type
c) Select host
d) Select storage
3) Configure channel networking
a) Primary-Secondary
b) Secondary-Tertiary
c) Tertiary-Primary
4) Ready to complete

Primary VM Name	NFCE81B	P-S Channel IP Address	10.0.0.5
Secondary Datacenter	Datacenter1	P-S Subnet Mask	255.255.255.0
Secondary Host	192.168.0.141	S-P Channel IP Address	10.0.0.6
Secondary Datastore	datastore1	S-P Subnet Mask	255.255.255.0
Tertiary Datacenter	Datacenter1	S-T Channel IP Address	10.0.1.6
Tertiary Host	192.168.0.141	S-T Subnet Mask	255.255.255.0
Tertiary Datastore	datastore1	T-S Channel IP Address	10.0.1.7
Tertiary Public IP Address	192.168.10.7	T-S Subnet Mask	255.255.255.0
Location for Tertiary VM	Use Tertiary host location	T-P Channel IP Address	10.0.2.7
Gateway	192.168.10.1	T-P Subnet Mask	255.255.255.0
Preferred DNS	192.168.10.2	P-T Channel IP Address	10.0.2.5
Alternate DNS		P-T Subnet Mask	255.255.255.0

Ready to complete
The Secondary and Tertiary VMs will be cloned to the specified locations.
Cloning may take some time, depending on volume of data and available bandwidth.
If you have selected to move the VM, once the cloning has completed, copy the VMDK files to the remote site, and power-on the Tertiary.
Otherwise, once the cloning has completed, the servers will begin replicating automatically.

Back Next Finish Cancel

The *Ready to complete* step is displayed.

- Review all of the summary information on the step. If any errors are found, use the **Back** button to navigate to the step with the error and correct it. If no errors are found, click **Finish** to deploy the Secondary and Tertiary servers.

1.4.1.8. Upgrade Applications

About this task

This feature prepares the upgrade of protected applications on a Neverfail Continuity Engine cluster.

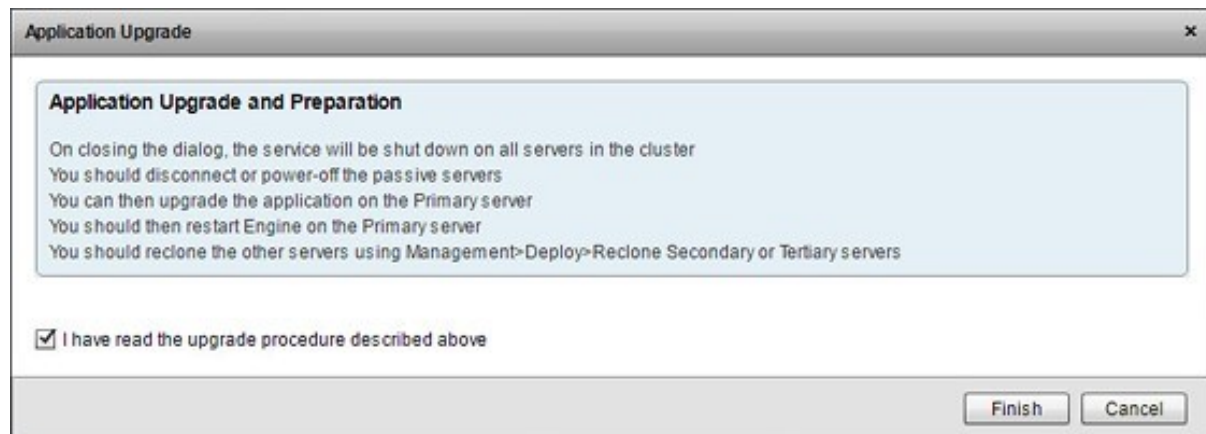
The upgrade of protected applications is done manually on the Primary server, while the passive servers are disconnected or powered down. Once the protected applications are upgraded on the Primary server, the cluster can be redeployed using the *Recloning Secondary or Tertiary Server*.

Note

The Upgrade applications feature prepares the cluster for the upgrade procedure, the user having to execute all the steps manually.

The upgrade procedure can be started by selecting the **Application upgrade** option in the **Management > Deploy** menu.

Figure 2-41. Upgrade Applications



The following steps are required in order to safely upgrade a protected application:

Procedure

1. Carefully read the procedure described in the *Application Upgrade* dialog and check the **I have read the upgrade procedure described above** check-box. Checking the safety check-box will enable the **Finish** button.
 2. Click **Finish** to close the dialog.
- The Neverfail Continuity Engine service will be shut down on all the servers in the cluster.
3. Disconnect or shut down the passive servers.

The current Secondary and Tertiary servers will not be used anymore but should be kept as backup, in case the application upgrade procedure does not go as planned.

4. Perform the upgrade on the desired application on the Primary server.
5. Restart the Neverfail Continuity Engine once the application has been upgraded successfully.
6. Reclone the Secondary and Tertiary servers.

Use the *Recloning Secondary or Tertiary Server* to safely redeploy the Secondary and Tertiary servers (if applicable). The old passive servers can be either discarded or stored for backup purpose.

1.4.1.9. Recloning Secondary or Tertiary Server

The *Reclone Secondary or Tertiary Server* feature allows the administrator to perform a server reclone, either in place or scheduled. The **Secondary node, Tertiary node or both Secondary and Tertiary nodes** can be redeployed using this feature, while the **Primary node needs to be active** to serve as source for the recloning operation. This feature is available in the **Management** menu, at the bottom of the **Deploy** section.

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 *Configure Connection to VMware vCenter Server* 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, Manual and Scheduled Automated recloning options, considering the above conditions are met.

The Reclone Secondary or Tertiary Server feature provides three cloning options in the Select clone type section:

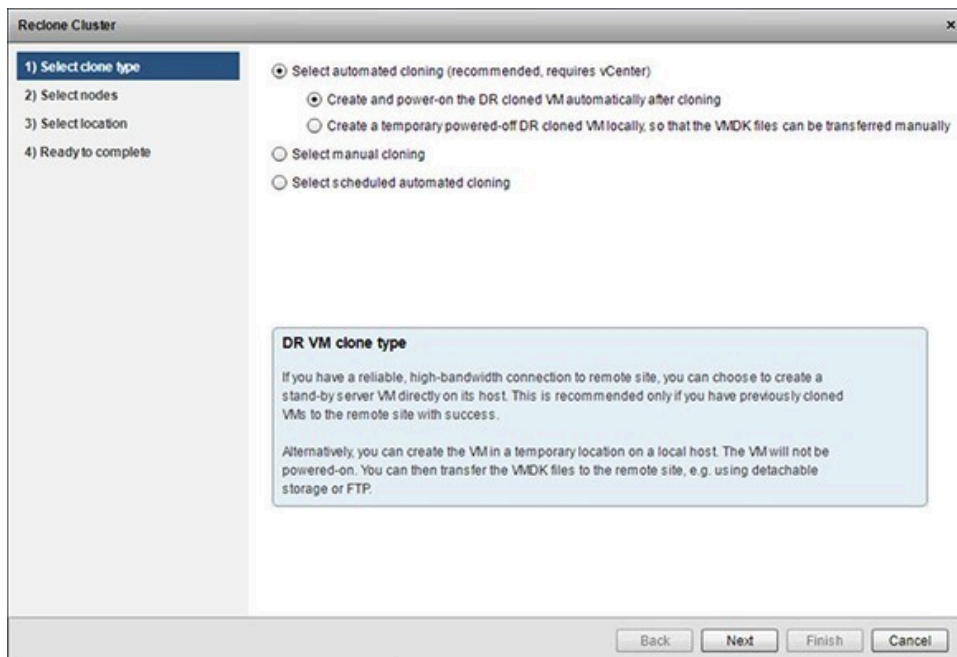
- Automated Recloning
- Manual Recloning
- Scheduled Recloning

Automated Recloning

The Automated Recloning task is completely handled by either vCenter Server or vCenter Converter.

This option is available when choosing the *Select automated cloning (recommended, requires vCenter)* option when configuring the clone type.

Figure 2-42. Select automated cloning



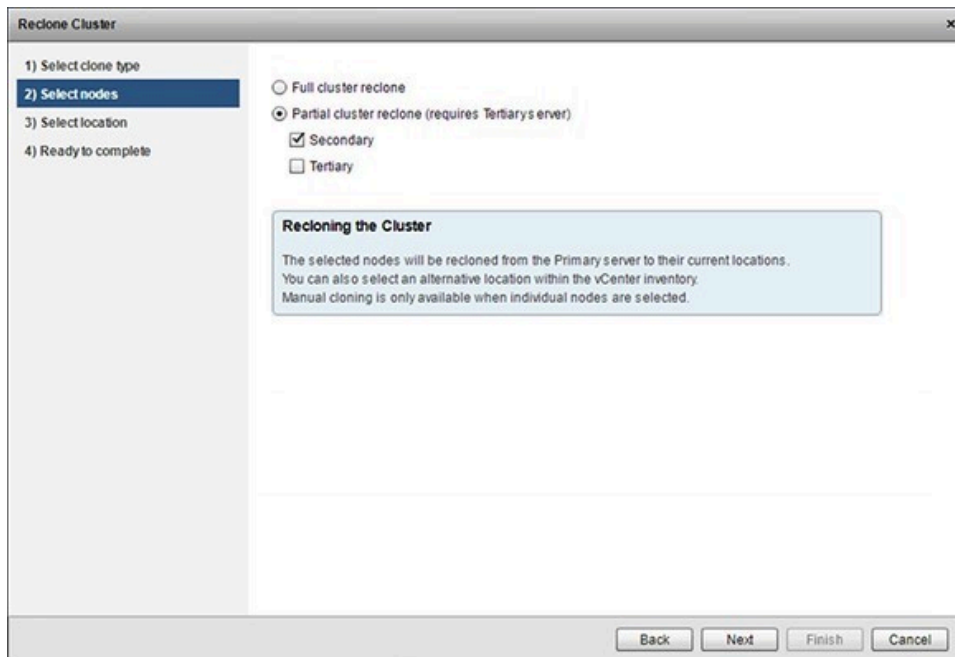
The cloning task is completely handled in an automated manner by either vCenter Server or vCenter Converter, providing that the vCenter Server connection, including the Default Host and vCenter Converter connection, if needed, have been properly configured.

The automated recloning task can handle the clones using two different approaches:

- **Create and power-on the DR cloned VM automatically after cloning.** This option is suitable when a high-bandwidth connection is available to the target clone host. It is recommended to make sure that you can successfully clone a VM to the remote host before using this option.
- **Create a temporary powered-off DR cloned VM locally, so that the VMDK files can be transferred manually.** This option is suitable when the cloned VM will be moved manually to the new host (for example using detachable storage or FTP to transfer the VMDK files).

Once the automated cloning option is chosen, the nodes to be cloned are available for selection in the *Select nodes* page.

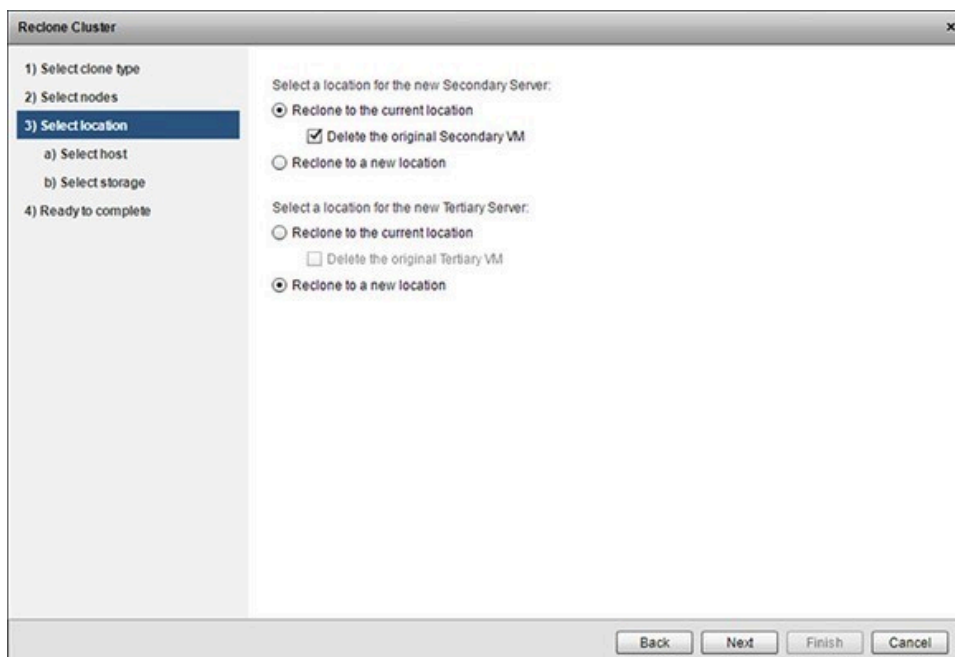
Figure 2-43. Select nodes



A **full cluster reclone** will clone the available passive nodes. If the cluster has a Tertiary node besides its Secondary, the **partial cluster reclone** option is available. This allows the recloning of only the Secondary or Tertiary nodes, or both, if selected so.

The selected nodes can be cloned to the same location as the original nodes or to a new host configured in the *Select location* page.

Figure 2-44. Select location

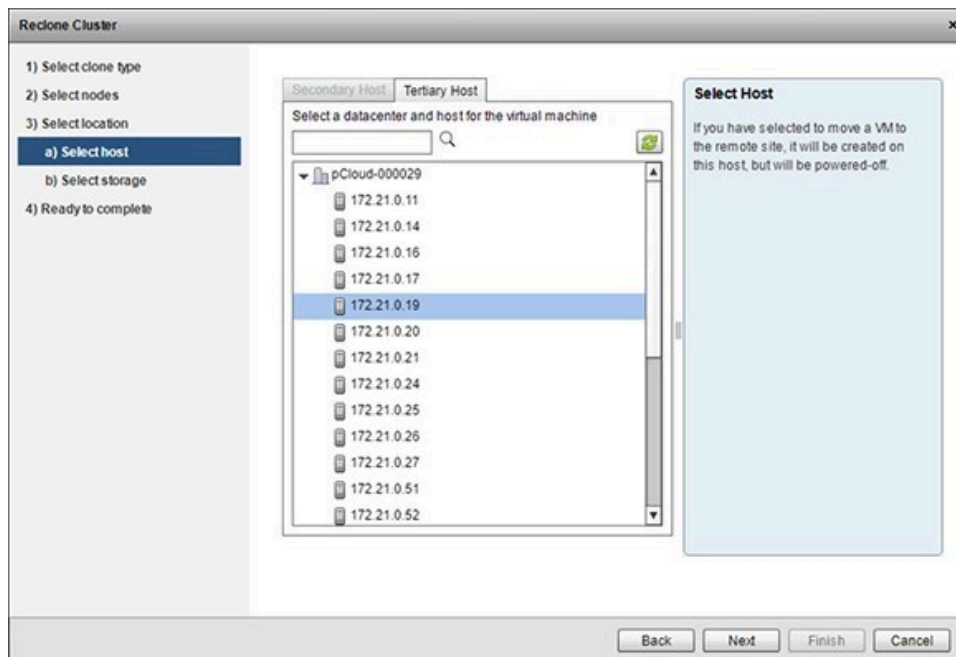


The location for each selected node can be configured as follows:

- **Reclone to the current location.** Uses the current location of the selected node as clone host. The Delete the original VM option allows the replacement of the selected node with its new clone.

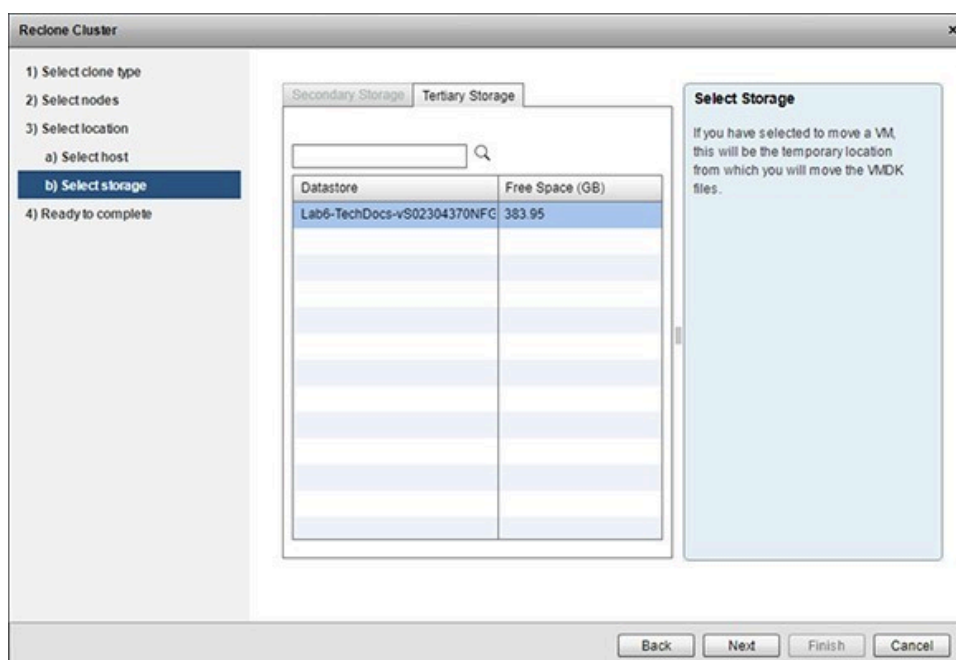
- **Reclone to a new location.** Allows the selection of a new host and storage for the node's clone. When choosing this option, the **Select host** and **Select storage** sub-sections will be available:
 - The *Select host* page allows you to choose a target host through VMware vCenter Server, where the new clone will be created. The clone will be powered-off after creation. The host selection is available for each node marked for recloning.

Figure 2-45. Select host



- The *Select storage* page allows you to select the storage location for each node. When creating powered-off clones for manual transfer, the VMDK files of the cloned VMs will be stored here.

Figure 2-46. Select storage



Note

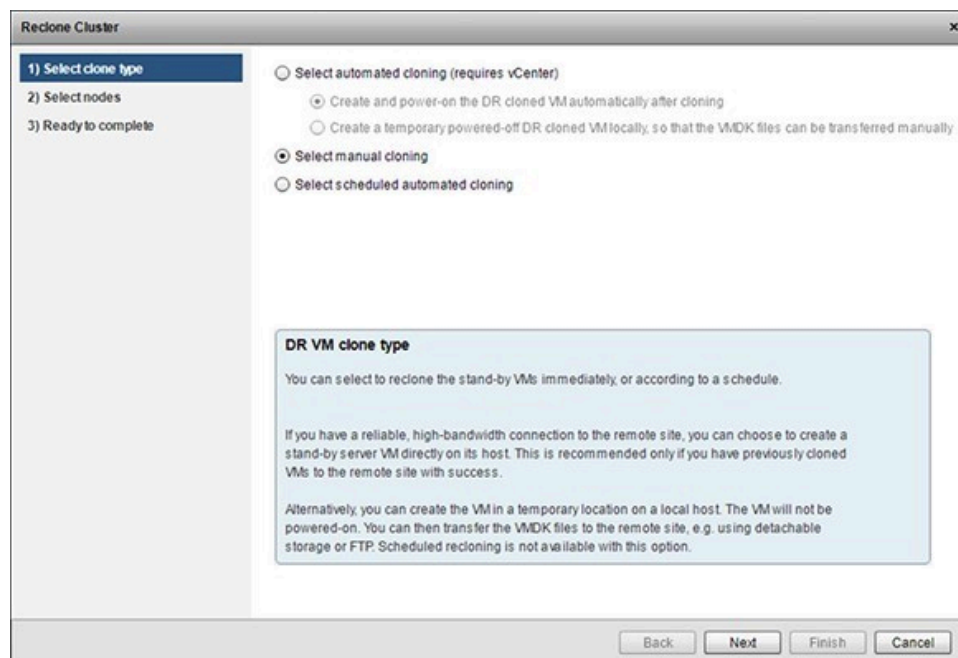
If the new location has not been configured for the recloning job or the EMS is not able to retrieve the original reclone target location from the Secondary or Tertiary nodes, the VMware vCenter Server Default Host will be used for the target reclone location. The Default Host needs to be configured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

The *Ready to complete* page shows the summary of the recloning task. The automated recloning task will be executed immediately after its configuration, as soon as the cluster is in a ready state.

Manual Recloning

This cloning task is manually executed by the user. The manual recloning task can be executed at any time after its configuration, when initiated by the user.

Figure 2-47. Reclone manual cloning



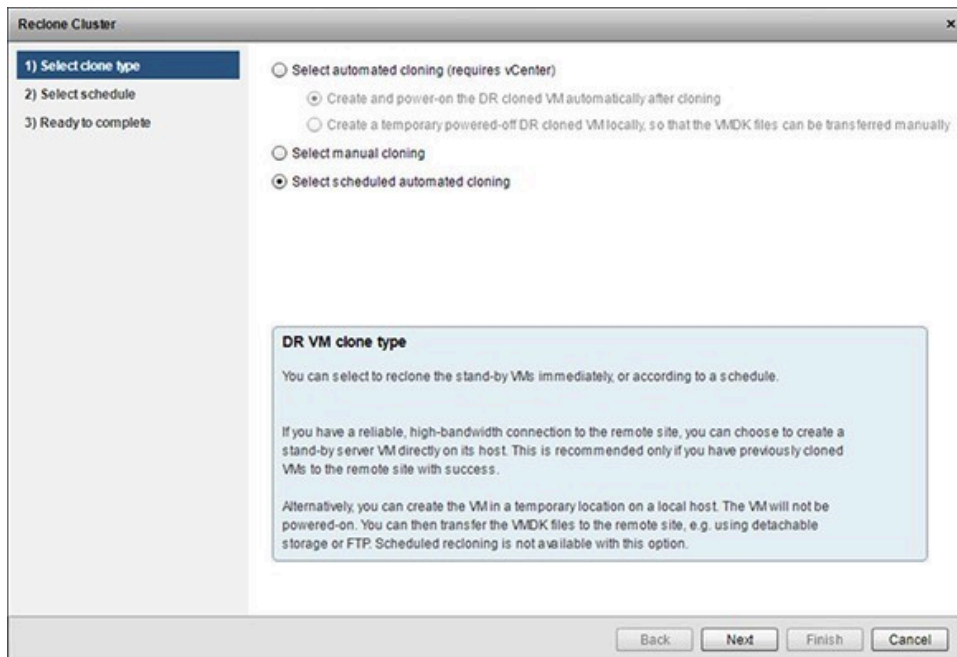
Once the manual cloning option is chosen, the nodes to be cloned are available for selection in the *Select nodes* page. Just like in Automated reclone, a **full cluster reclone** means that all available passive nodes will be cloned (manually, by the user). If the cluster has a Tertiary node besides its Secondary, the **partial cluster reclone** option is available. This allows the recloning of only the Secondary or Tertiary nodes, or both, if selected so.

The *Ready to complete* page shows the summary of the recloning task. The task will not be executed unless initiated by the user.

Scheduled Recloning

Allows the scheduling of a recloning task that executes in the same manner as the Automated recloning option, but at a specified time and using the specified repetition pattern.

Figure 2-48. Select scheduled automated cloning

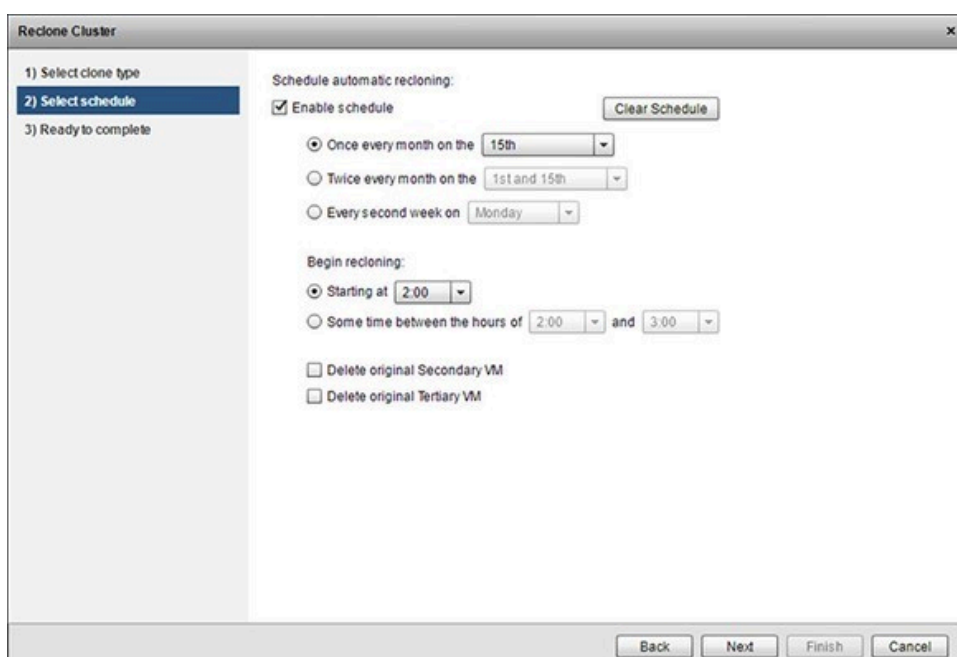


Note

Scheduled recloning may not be available for Engine clusters upgraded from older versions until at least one automated reclone operation is performed before using the Scheduled recloning feature.

After selecting the Scheduled automated cloning option, the Select schedule page offers the possibility to enable the schedule and configure it.

Figure 2-49. Select schedule



The automated recloning execution date can be programmed as follows:

- **Once every month.** The reclone task is executed in the selected day of every month.
- **Twice every month.** The reclone task is executed in the two selected days of every 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.** The reclone task can be scheduled to run once in two weeks, on the selected day. This allows a greater flexibility for scheduling two weeks reclone intervals without taking in consideration the length of the months.

The time of the execution set using the following options:

- **Starting at.** The reclone task is executed at the selected hour during the planned execution day(s). The Engine will wait for the cluster to be in the ready state within the span of the selected hour. If the cluster becomes ready for recloning within the time frame, the reclone task will be executed. If the cluster does not enter the ready state within the 60 minutes time frame, the reclone task will not be executed.
- **Some time between the hours.** The reclone task is executed in the time frame between the selected hours, as soon as the cluster is in the ready state. If the cluster does not enter the ready state within the specified time frame, the reclone task will not be executed. This option allows the configuration of a time frame bigger than 60 minutes in which the cluster can be ready for recloning. The original passive nodes can be deleted after cloning if the **Delete original VM** option is selected.

A schedule can be reset (or cleared) using the **Clear Schedule** button available in the top side of the page.

The *Ready to complete* page shows the summary of the recloning schedule and the reclone task that will be performed.

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.

1.4.2. Manage

The **Manage** drop-down provides key management abilities such as to Discover Protected Servers, Add a Protected Server, Remove the Selected Server, and Download the Advanced Management Client.

1.4.2.1. Discover Protected Servers

About this task

Neverfail Continuity Engine Management Service provides the ability to perform discovery to identify all Neverfail Engine Clusters.

To discover protected servers:

Procedure

1. From the **Management > Manage** drop-down pane, click **Discover Protected Servers**.

The *Discover Server* dialog is displayed.

Figure 2-50. Discover Protected Servers dialog

Server	Result
--------	--------

2. Identify the IP address range to search by adding a beginning and ending IP address in the *Begin* and *End* fields.

Neverfail recommends leaving the Port Number field with the default port unless the default port is in use by another application and a custom port has been configured.

3. Add a username and password used to connect to Neverfail Engine in the *Username* and *Password* fields.

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

4. Click **Search** to run Neverfail Engine server discovery.

The Neverfail Continuity Engine Management Service displays all Neverfail Continuity Engine clusters discovered. Discovered items will be added automatically to the Protected Servers pane in the background.

5. Click **OK** or **Cancel** to dismiss the Discover Protected Servers dialog.

1.4.2.2. Add a Protected Server

About this task

Neverfail Continuity Engine Management Service allows you to add individual protected servers which may be part of a cluster.

Procedure

1. Click **Add a Protected Server** in the **Management > Manage** drop-down pane to add a server.

The *Add Server* dialog is displayed.

Figure 2-51. Add Server dialog

The image shows a Windows-style dialog box titled "Add Server". Inside the dialog, there is a section titled "Add a protected server to be managed". Below this title, there is a label "Enter the hostname (or Public IP address) and port number". This is followed by two input fields: "Host" with the value "192.168.0.5" and "Port Number" with the value "9727". Below these fields is another label "Enter the credentials for connecting to the server". Underneath this is a note: "Domain accounts should use the syntax username@domain". There are two more input fields: "Username" with the value "Administrator" and "Password" which is masked with asterisks. At the bottom right of the dialog are two buttons: "OK" and "Cancel".

2. Enter the hostname or IP address of server to be added in the Host field.

Neverfail Continuity Engine Management Service recommends leaving the Port Number field with the default port unless the default port is in use by another application and a custom port has been configured.

3. Click **OK** to add the Neverfail cluster.

The Neverfail Continuity Engine Management Service adds the Neverfail Engine cluster to the Protected Servers pane of the Neverfail Continuity Engine Management Service Summary page.

1.4.2.3. Remove the Selected Server

About this task

The Neverfail Continuity Engine Management Service provides the ability to remove specific Neverfail servers from the Neverfail Continuity Engine Management Service *Protected Servers* pane.

Procedure

1. Select the server to be removed from *Protected Servers* pane of the Neverfail Continuity Engine Management Service.
2. Select **Remove the Selected Server** in the **Management > Manage** drop-down pane.

The *Remove Server* dialog is displayed. You are prompted to verify that you want to remove the selected server from management by the Neverfail Continuity Engine Management Service.

Figure 2-52. *Remove Server* dialog



3. Click **OK**.

The intended Neverfail Engine server is removed from the Neverfail Continuity Engine Management Service Protected Servers pane.

1.4.2.4. Download the Advanced Management Client


About this task

The *Download the Advanced Management Client* feature is used to download the Advanced Management Client (Client Tools) to a workstation or server for remote management of Neverfail Engine.

Procedure

1. Select the **Download Advanced Management Client** feature.

Figure 2-53. *Download Advanced Management Client*



2. Select a target location for the downloaded file using the dialog navigation features.
3. Click **Save**.

1.4.2.5. Open or View Support Ticket

About this task

This feature allows you to open a new support ticket or view the details of an already opened one.

To open a new support ticket or view an existing one, use the **Open or view support ticket** option in the **Management > Manage** menu.

Procedure

1. Click the Open or view support ticket option.

The *Neverfail support portal* will open in a new browser tab.

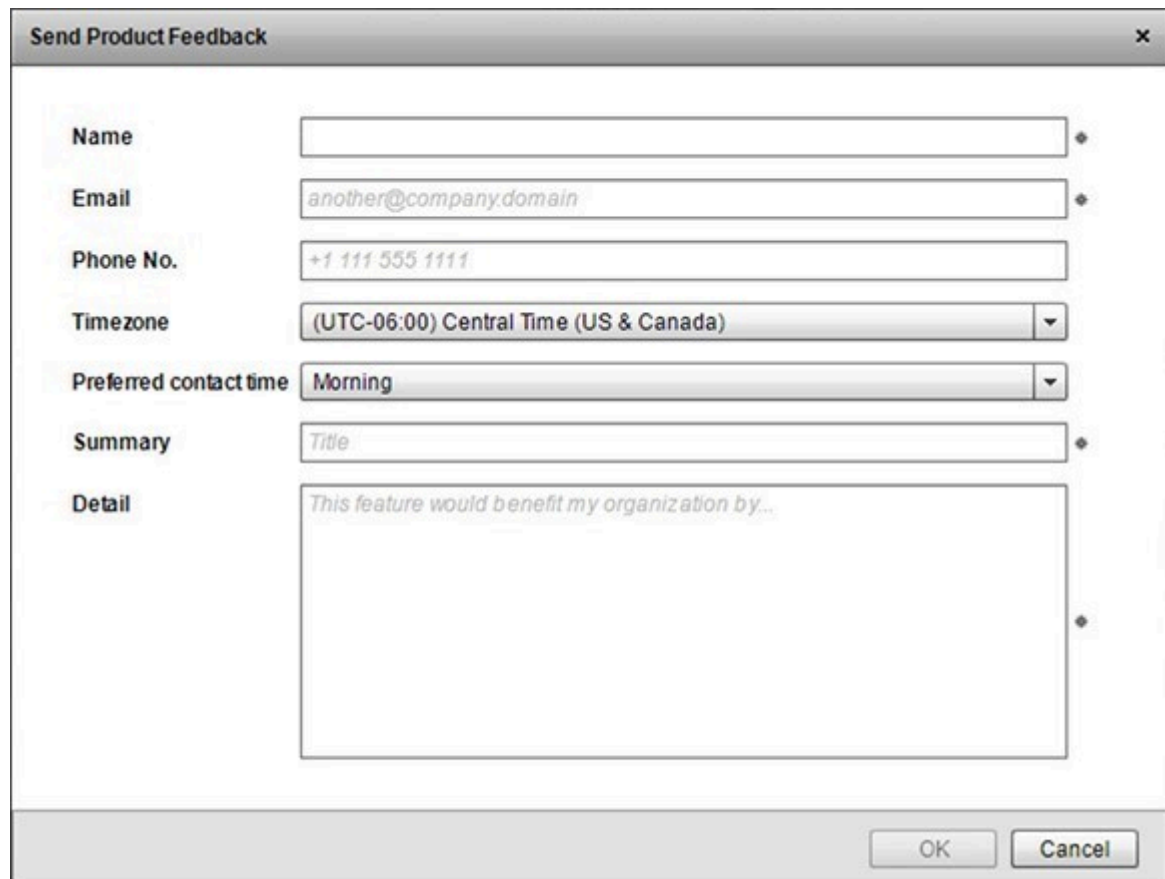
2. Log in the support portal using the appropriate credentials. Once inside the support portal, follow the normal procedure for creating tickets and viewing the status of existing ones.

1.4.2.6. Send Product Feedback

About this task

This feature allows you to send your Continuity Engine feedback to Neverfail.

Figure 2-54. Send Product Feedback



The image shows a 'Send Product Feedback' dialog box with a title bar and a close button. It contains several input fields: 'Name' (empty), 'Email' (filled with 'another@company.domain'), 'Phone No.' (filled with '+1 111 555 1111'), 'Timezone' (dropdown menu showing '(UTC-06:00) Central Time (US & Canada)'), 'Preferred contact time' (dropdown menu showing 'Morning'), 'Summary' (text field with placeholder 'Title'), and 'Detail' (large text area with placeholder 'This feature would benefit my organization by...'). At the bottom right are 'OK' and 'Cancel' buttons.

To send your feedback for the Neverfail Continuity Engine product, click the **Give product feedback** option from the **Management > Manage** menu.

Procedure

1. Click the **Give product feedback** menu option.

The *Send Product Feedback* dialog will open on screen.

2. Enter the following details in the dialog fields.
 - Name (required)
 - Email (required)
 - Phone number
 - Time zone
 - Preferred contact time
 - Summary (required): a short, one line description of the feedback.
 - Details (required): your detailed feedback regarding the Neverfail Continuity Engine product.
3. Click **OK** to send your feedback to Neverfail.

1.4.3. Integrate

Neverfail Continuity Engine Management Service allows you to easily integrate some VMware vCenter functionality directly from the Neverfail Continuity Engine Management Service user interface.

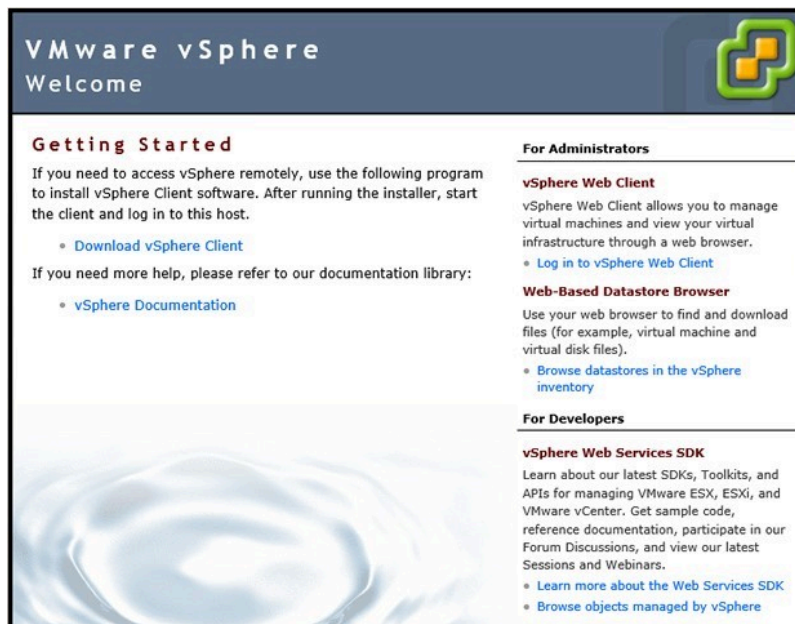
1.4.3.1. Log in to VMware vSphere Client

Neverfail Continuity Engine Management Service provides the ability to log in to the VMware vSphere Client directly from Neverfail Continuity Engine Management Service to manage VMware resources.

To log in to VMware vSphere Client:

Using the *Neverfail Continuity Engine Management Service* user interface, select **Log in to VMware vSphere Client**. A browser is launched providing access to the VMware vSphere Client.

Figure 2-55. VMware vSphere



1.4.3.2. Create VMware SRM Plan Step for Selected Server

Before you begin

- The Neverfail Continuity Engine Management Service installed on vCenter Server in the Recovery and Protected Sites
- Microsoft PowerShell 2.0 installed on all SRM servers that will run command files, for example the SRM Servers in the Recovery and Protected sites
- The PowerShell Execution Policy must be set to RemoteSigned on all SRM Servers, use the following PowerShell command:

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

About this task

This feature works to extend capabilities of VMware's Site Recovery Manager (SRM). While SRM provides the ability to failover virtual servers to a secondary site, this feature integrates Neverfail Engine physical or virtual servers into the failover process as a natural step in the SRM Site Recovery Plan executed by SRM. It works by allowing the administrator to create an SRM Step that can be added to the SRM Site Recovery Plan thereby allowing servers protected by Neverfail Engine to participate in failover of servers protected by Site Recovery Manager.

Procedure

1. Launch the Neverfail Continuity Engine Management Service user interface.
2. Select a Neverfail Engine server in the left pane to be added to the SRM Site Recovery Plan.

Important

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

3. Click the **Management > Integrate > Create VMware SRM Plan Step for Selected Server** button.

The *Create a Plan Step for VMware vCenter Site Recovery Manager* dialog is displayed.

Figure 2-56. Create SRM Plan Step

Create a Plan Step for VMware vCenter Site Recovery Manager

Create a script to initiate a switch-over of NFV8.abcd.local as part of an SRM recovery plan

Requires Powershell V2 on the SRM server and permission for powershell scripts to run locally without signing. For servers which are members of Business Application Groups, all members of a group will failover or switchover together. It is recommended to add only the 'First to switch' server of a group to the SRM plan.

✓ Authentication token generated for switch-over of NFV8.abcd.local

1) Choose which server the script will make active. This depends on which server is located on the site for which you are creating a plan. In order to make the server active on either site, you will require two scripts - one for each option.

☐ Make Primary server active ☒ Make Secondary (or Tertiary) server active

2) If you want the plan to wait for the server to become active, enter the number of seconds. Otherwise, enter 0.

Maximum time to wait:

3) Enter alternate IP addresses by which the SRM server can reach the server when passive. Multiples are separated by commas.

Alternate IP addresses:

4) If you want to log script output to a file on the SRM server, enter the path here otherwise leave blank. Recommended for SRM 5.0

Log file for command:

5) The script should be saved and copied to the SRM server on the same site as the server being made active. For SRM 5.0, the scripts must have identical names and locations on each SRM server. Use the Save As... button to save it as a batch file.

6) Paste this command into the recovery plan in the SRM client, ensuring it matches where you have placed the script on the SRM server.

`c:\windows\system32\cmd.exe /c c:\nf_make_active_NFV8.abcd.local.bat`

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

If the SRM Plan Step is being created on the site where the Primary server is located, select **Make Primary Server Active**. If the SRM Plan Step is being created on the site where the Secondary server is located, select **Make Secondary server active**.

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

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

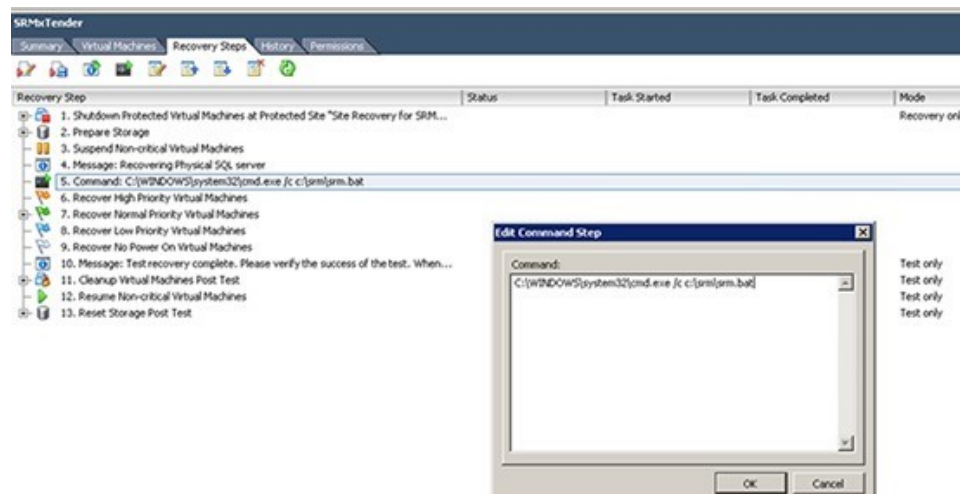
6. Alternate IP addresses are configured on each server in the Neverfail pair so that SRM can switch the servers even when the Protected Site cannot be contacted, for example in times of disaster. Enter the Alternate IP address that will be used by SRM to contact the Neverfail Engine server in the *Alternate IP addresses* field, separate multiple IP addresses with a comma. These IP addresses are typically added to the servers as *Management IP Addresses*.
7. If you want to log the script output to a file on the SRM server, enter a path in the *Log file for command* field (recommended for SRM 5.0), otherwise, leave the field blank.
8. Generate two scripts using the SRMXtender Plug-in.
 - Generate one script with *Make Primary Server Active* selected.
 - Generate one script with *Make Secondary Server Active* selected.

9. The scripts should be saved as *.bat* files with each being saved to a file share on the SRM server in the same site as the server being made active. Click the **Save As** button to save the script as a *.bat* file.

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

10. Launch the VMware vSphere Web Client and connect to the Recovery vCenter Server.
11. Navigate to **Home > Solutions and Applications > Site Recovery Manager** and select the intended **Recovery Plan**.
12. Select the *Recovery Steps* tab.

Figure 2-57. SRM Edit Command Step



13. Add a *Command Step* at the desired point in the Recovery Plan, for example before the *Recover High Priority Machines Step* if the applications running on these servers depend upon the physical server.
14. In the *Add Command Step* dialog enter: `C:\WINDOWS\system32\cmd.exe /c`
`<path_to_saved_file>\<file_name>.bat`
`<path_to_saved_file>` is the path where you have saved the `\<file_name>.bat` file at step 9.
15. Click **OK**.
Repeat the step creation process for each Neverfail pair that is to participate in the Site Recovery Plan.

1.4.4. License

The Neverfail Continuity Engine Management Service user interface provides the ability to license your Neverfail Continuity Engine cluster using a simple wizard.

1.4.4.1. Configure an Internet Proxy Server for Licensing

About this task

For organizations that use an Internet Proxy, the Configure Internet Proxy Settings dialog provides the ability to configure settings for the proxy to allow Neverfail Engine licensing to successfully complete.

To configure for use with an internet proxy:

Procedure

1. Provide the hostname or IP address of the proxy, the port number, and if required account credentials.

Figure 2-58. Configure Internet Proxy Settings

Configure Internet Proxy Settings

An Internet connection is required from Neverfail CE Management Service when you are selecting licenses to apply. If you require an Internet proxy, enter the details below.

☒ Use a proxy server

Host Name or IP Address: 192.168.1.211

Port Number: 8081

☒ Use the following credentials:

User Name: Administrator

Password: *****

OK Cancel

1.4.4.2. License the Selected Server

About this task

Licensing is performed via the Neverfail Continuity Engine Management Service.

Note

Automated licensing of Neverfail Engine requires use of the internet. If your organization uses an internet proxy, configure proxy information in the **Management -> License > Configure an Internet proxy server for licensing** dialog.

Procedure

1. To add a license for Neverfail Engine, navigate to the *Management* drop-down and click on **License > License the Selected Server**.

Figure 2-59. Apply License page

Apply License

1) Enter extranet credentials
 2) Accept EULA
 3) Select license
 4) Ready to apply license
 5) Apply license

Enter your email address. For the automated licensing option below, use your Neverfail account

engine@neverfail.com ☐ Remember email address

☒ Apply a license from your Neverfail account (requires Internet connection)

Enter the password

☐ Manually enter a license key

Enter the license key

Licensing Engine on CE-primary (Signature: 4ZUH4Q62)

Thank you for your interest. If you have not already purchased a license, please contact Neverfail. If you have purchased a license, please provide your credentials to access your licenses.

Proxy settings can be configured if a direct Internet connection is not available to Neverfail CE Management Service. See Management>License... If you have no Internet connection or need to reset your extranet password, please contact Neverfail support or email support@neverfail.com.

[Contact Neverfail support.](#)

EULAAcceptance: -

Back Next Finish Cancel

- If there is an Internet connection from the Neverfail Continuity Engine Management Service, select *Apply a License from your Neverfail account*, enter your Neverfail account password and press **Next**.
- If there is no Internet connection from the Neverfail Continuity Engine Management Service, you can obtain a license key from Neverfail. Select *Manually enter a license key*, enter the key and press **Next**.

Figure 2-60. Manual License Entry

Apply License

1) Enter extranet credentials
2) Accept EULA
3) Select license
4) Ready to apply license
5) Apply license

Enter your email address. For the automated licensing option below, use your Neverfail account
 ☐ Remember email address

☐ Apply a license from your Neverfail account (requires Internet connection)
 Enter the password

☒ Manually enter a license key
 Enter the license key

Licensing Engine on CE-primary (Signature: 4ZUH4Q62)
 Thank you for your interest. If you have not already purchased a license, please contact Neverfail. If you have purchased a license, please provide your credentials to access your licenses.
 Proxy settings can be configured if a direct Internet connection is not available to Neverfail CE Management Service. See Management>License... If you have no Internet connection or need to reset your extranet password, please contact Neverfail support or email support@neverfail.com.
[Contact Neverfail support.](#)

EULAAcceptance: -

2. In the Accept EULA step, read the EULA content and accept the terms of the agreement using the **I Accept Terms of the License Agreement** checkbox.

- If the *Apply a License from your Neverfail account* option is used, once the license agreement is accepted and the **Next** button is clicked, the *Apply License* wizard will continue with step 3.
- If the *Manually enter a license key* option is used, once the license agreement is accepted and the **Next** button is clicked, the *Apply License* wizard will continue with step 5.

Figure 2-61. Accept EULA

Apply License

1) Enter extranet credentials
2) Accept EULA
3) Select license
4) Ready to apply license
5) Apply license

NEVERFAIL END-USER LICENSE AGREEMENT

This End-User License and Support Agreement ("EULA" or "Agreement"), made and entered into as of the Effective Date, is a legal agreement between you, either an individual or an entity ("Licensee") and (a) Neverfail LLC, if Licensee resides in the United States ("Neverfail") or (b) Neverfail Ltd. if the Licensee resides outside the United States ("Neverfail"). This Agreement sets forth the terms and conditions under which Neverfail licenses certain of its software products and provides related support services to Licensee. Licensee acknowledges that Licensee has read Neverfail's Support Services Agreement ("SSA"), made available to Licensee on Neverfail's website, www.neverfail.com.

IMPORTANT-READ CAREFULLY: BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, LICENSEE AGREES TO BE BOUND BY THE TERMS OF THIS END USER LICENSE AGREEMENT ("EULA") AND THE SSA. IF LICENSEE DOES NOT AGREE TO THE TERMS OF THIS EULA, DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE.

1. DEFINITIONS

1.1 "Confidential Information" means all non-public information provided by or relating to a Discloser or its affiliates and includes, without limitation, either party's source code, customer lists, products, product roadmaps, financial information, business information and marketing strategies disclosed in written or other tangible form (including on magnetic media) or by electronic, oral, visual or other means.

1.2 "Delivery Date" means the date the Software License Key is provided to Licensee to allow Licensee to download the Licensed Software for electronic delivery.

1.3 "Designated Computer(s)" means the server(s) or processor(s) on which the Licensed Software will be installed and that will be configured not to exceed the Quantity specified in each Purchase Order.

1.4 "Discloser" means a party, including its Representatives, that discloses Confidential Information to

☒ I Accept Terms of the License Agreement

EULAAcceptance: -

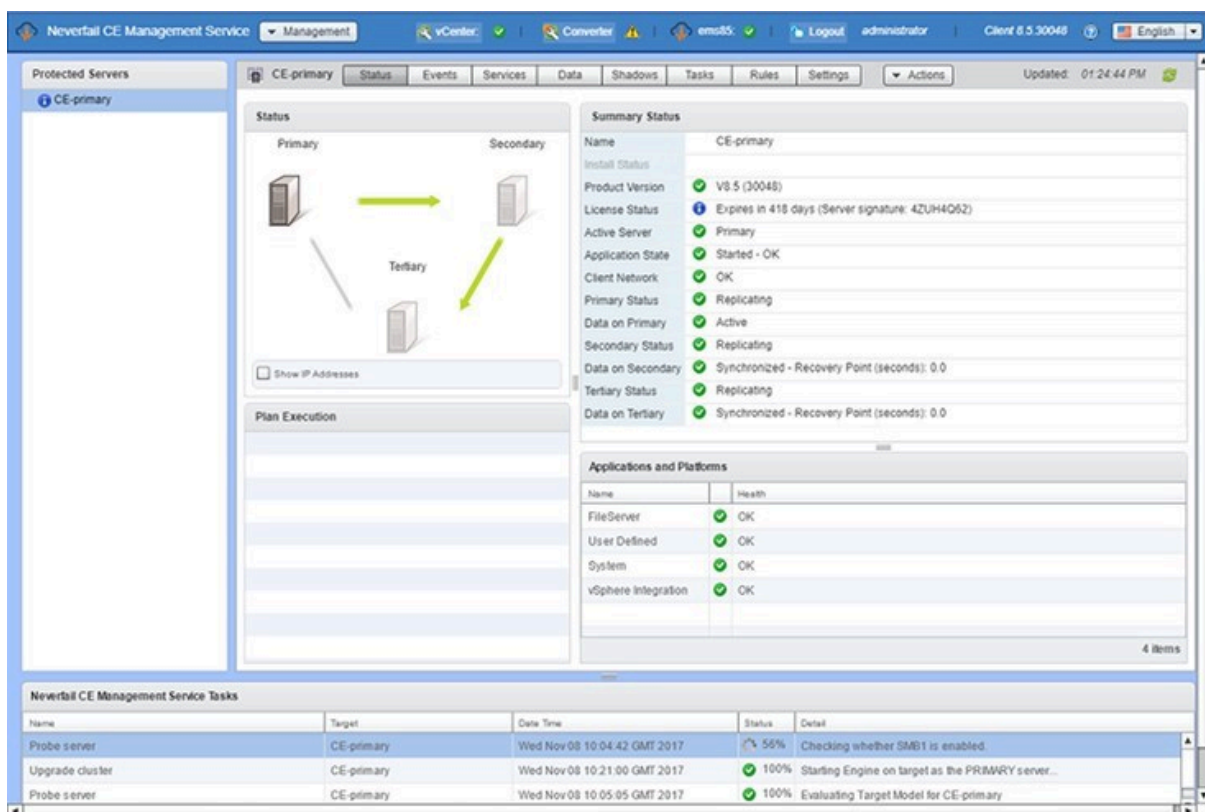
3. In the *Select License* step, from the table of licenses, select the license to apply based on the features required. Licenses already used for the selected cluster are shown as **Applied**. Click **Next**.
4. Review the *Ready to Complete* summary information and click **Next**.
5. On the *Apply License* step, click **Finish**.

1.5. Summary

The *Summary Page* contains multiple panes that provide the current status of the server, the version of the cluster, and details about licensing of the cluster.

The Neverfail Continuity Engine Management Service identifies the current active server and provides the status of Replication, the Application State, the File System State, and the Client Network State of servers in the cluster.

Figure 2-62. Summary Page



The Neverfail Continuity Engine Management Service supports multiple language translations for its user interface elements and presented information. The top-right drop-down menu allows you to change the language of the Engine Management Service without reloading the user interface.

Figure 2-63. Change Language Menu

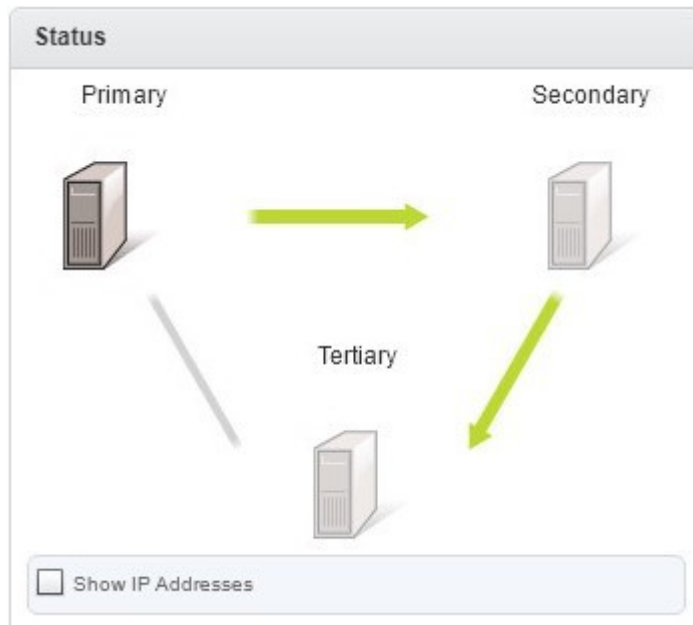


1.6. Status

The *Status* pane provides a view of the currently selected server pair or trio.

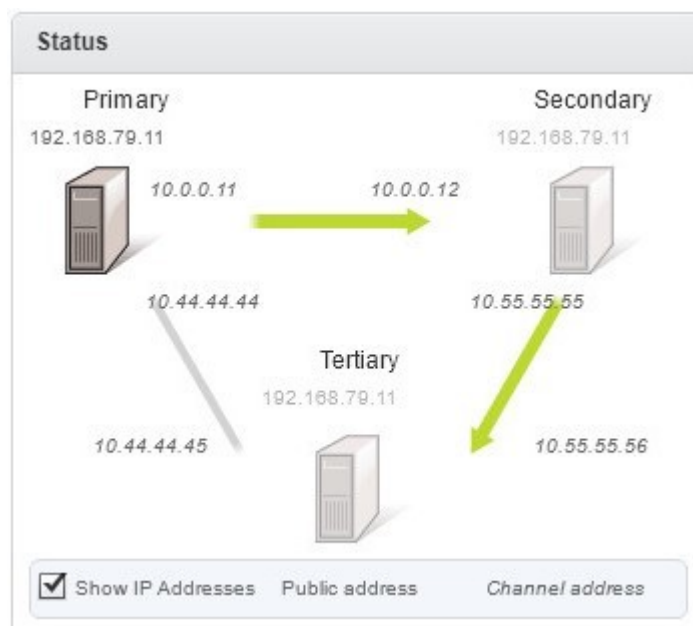
The *Status* pane displays a graphic representation of the currently selected cluster and what the cluster is doing. Additionally, it displays which of the servers are active, the status of replication, and the direction of replication (for example in a pair, Primary to Secondary or Secondary to Primary).

Figure 2-64. Status Pane




The *Status* pane also allows you to see the cluster IP schema by checking the **Show IP Addresses** check-box.

Figure 2-65. Displaying the IP Schema



Hovering the mouse cursor above any node of the cluster in the *Status* pane will display the IP details for the hovered node.

Figure 2-66. The IP Details on Cursor Hover



1.6.1. Summary Status

The *Summary Status* pane provides a status of all operations currently being performed on the server cluster.

The *Summary Status* pane displays the status of replication, synchronization, the application and network state, license status, and the installed version of Neverfail Engine.

Figure 2-67. Summary Status pane

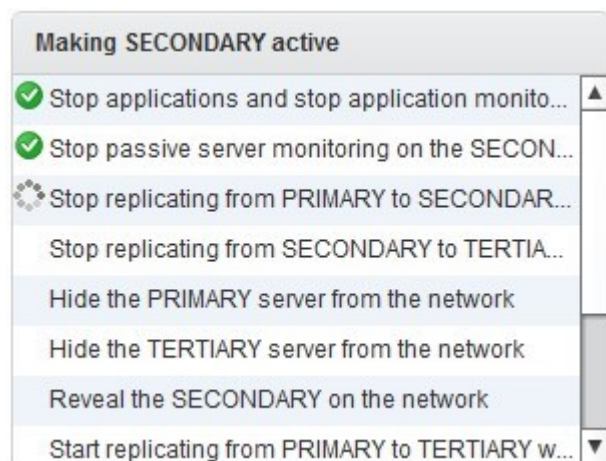
Summary Status	
Name	CE-primary
Install Status	
Product Version	✓ V8.5 (30033)
License Status	i Expires in 424 days (Server signature: 4ZUH4Q62)
Active Server	✓ Primary
Application State	✓ Started - OK
Client Network	✓ OK
Primary Status	✓ Replicating
Data on Primary	✓ Active
Secondary Status	✓ Replicating
Data on Secondary	✓ Synchronized - Recovery Point (seconds): 0.0
Tertiary Status	✓ Replicating
Data on Tertiary	✓ Synchronized - Recovery Point (seconds): 0.0

1.6.2. Plan Execution

The *Plan Execution* pane displays plans being executed by Neverfail Engine.

Plans are sequences of actions required to perform functions such as switch-over or installing a new plug-in. Plans can be executed in response to user action (such as Make Active) or automatically (such as failover). The *Plan Execution* pane will display the progress of the plan as it is executed. Once the plan is complete, it is removed from the *Plan Execution* pane.


Figure 2-68. Plan Execution pane



1.6.3. Applications and Platforms

The *Applications and Platforms* pane displays the currently installed protected applications and their status. It also shows the health status of platforms such as the OS and hardware.

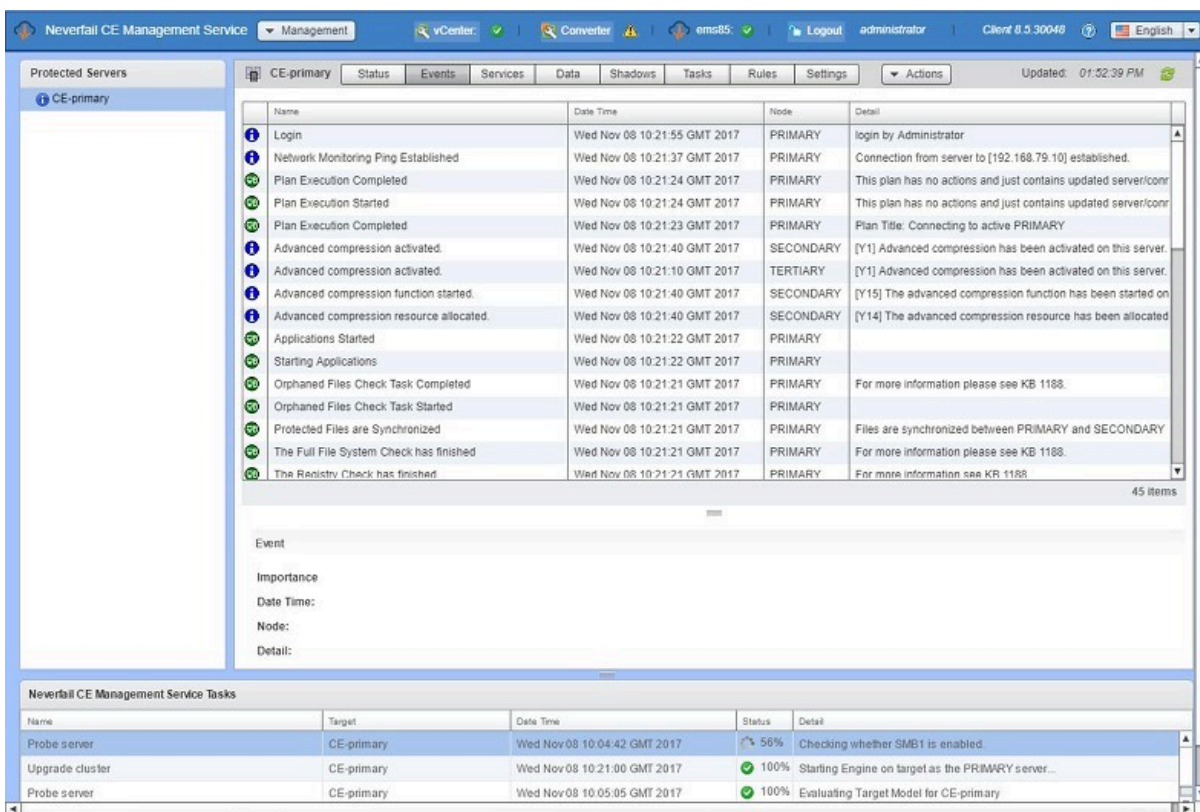
Figure 2-69. Applications and Platforms

Applications and Platforms		
FileServer		OK - OK
System		OK - OK
User Defined		OK, Finished 'DNSupdate (SECONDARY)' in 47400ms with status Completed with exit cod

1.7. Events

The events that Neverfail Engine logs are listed chronologically (by default) on the *Events* page, the most recent event appears at the top of the list with older events sequentially below it.

Figure 2-70. Events page







Name	Date Time	Node	Detail
Login	Wed Nov 08 10:21:55 GMT 2017	PRIMARY	login by Administrator
Network Monitoring Ping Established	Wed Nov 08 10:21:37 GMT 2017	PRIMARY	Connection from server to [192.168.79.10] established.
Plan Execution Completed	Wed Nov 08 10:21:24 GMT 2017	PRIMARY	This plan has no actions and just contains updated server/comr
Plan Execution Started	Wed Nov 08 10:21:24 GMT 2017	PRIMARY	This plan has no actions and just contains updated server/comr
Plan Execution Completed	Wed Nov 08 10:21:23 GMT 2017	PRIMARY	Plan Title: Connecting to active PRIMARY
Advanced compression activated.	Wed Nov 08 10:21:40 GMT 2017	SECONDARY	[Y1] Advanced compression has been activated on this server.
Advanced compression activated.	Wed Nov 08 10:21:10 GMT 2017	TERTIARY	[Y1] Advanced compression has been activated on this server.
Advanced compression function started.	Wed Nov 08 10:21:40 GMT 2017	SECONDARY	[Y15] The advanced compression function has been started on
Advanced compression resource allocated.	Wed Nov 08 10:21:40 GMT 2017	SECONDARY	[Y14] The advanced compression resource has been allocated
Applications Started	Wed Nov 08 10:21:22 GMT 2017	PRIMARY	
Starting Applications	Wed Nov 08 10:21:22 GMT 2017	PRIMARY	
Orphaned Files Check Task Completed	Wed Nov 08 10:21:21 GMT 2017	PRIMARY	For more information please see KB 1188.
Orphaned Files Check Task Started	Wed Nov 08 10:21:21 GMT 2017	PRIMARY	
Protected Files are Synchronized	Wed Nov 08 10:21:21 GMT 2017	PRIMARY	Files are synchronized between PRIMARY and SECONDARY
The Full File System Check has finished	Wed Nov 08 10:21:21 GMT 2017	PRIMARY	For more information please see KB 1188.
The Registry Check has finished	Wed Nov 08 10:21:21 GMT 2017	PRIMARY	For more information see KB 1188

Name	Target	Date Time	Status	Detail
Probe server	CE-primary	Wed Nov 08 10:04:42 GMT 2017	56%	Checking whether SMB1 is enabled.
Upgrade cluster	CE-primary	Wed Nov 08 10:21:00 GMT 2017	100%	Starting Engine on target as the PRIMARY server...
Probe server	CE-primary	Wed Nov 08 10:05:05 GMT 2017	100%	Evaluating Target Model for CE-primary

The events listed in the Event page 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, the full detail of the entry can be found in the lower portion of the pane when an event is selected.

There are four categories of importance of events that Neverfail Engine is configured to log:

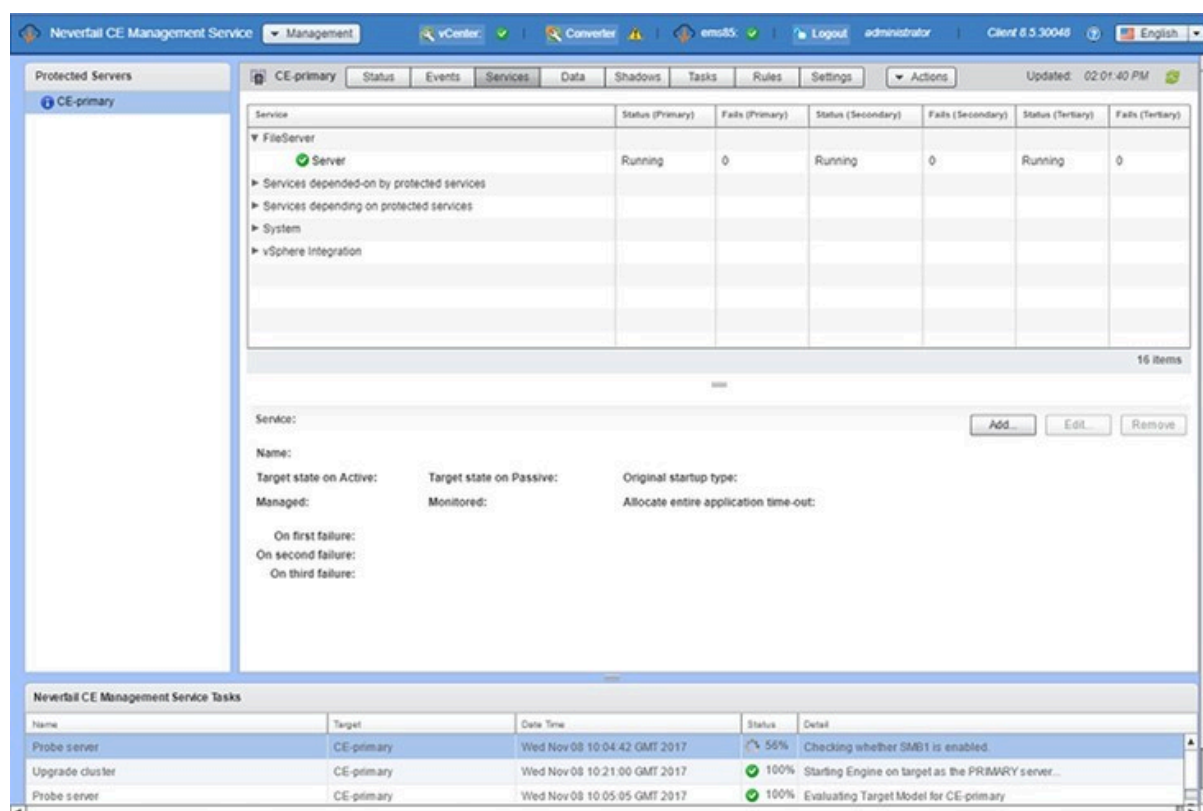
Icon	Definition
	These are critical errors within the underlying operation of Neverfail Engine and can be considered critical to the operation of the system.
	Warnings are generated where the system finds discrepancies within the Neverfail Engine operational environment that are not deemed critical to the operation of the system.

Icon	Definition
	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.

1.8. Services

The status of all protected services is displayed on the *Services* page. The status shows both the target and actual state for all servers in the cluster and the Failure Counts for each server.

Figure 2-71. Services page



The target state of protected services can be specified for the active and passive server(s), and is typically *Running* on the active and *Stopped* on the passive(s). Services depending on protected services are managed (for example, started and stopped) by Neverfail Engine but not monitored (for example, not restarted if stopped by some external agency). Services upon which protected services depend are monitored (for example, restarted if stopped) but not managed (for example, not stopped if protected applications are stopped).

1.8.1. Add a Service

About this task

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

To add a service:

Procedure

1. Select the Service tab and then click **Add** at the lower right of the pane.

Figure 2-72. Add Service

The screenshot shows the 'Add Service Protection' dialog box. It has a title bar 'Add Service Protection'. Inside, there's a 'Service:' dropdown menu with 'Schedule' selected. Below that are two dropdowns: 'Target state on Active server:' with 'Running' selected, and 'Target state on Passive server:' with 'Stopped' selected. There are two checked checkboxes: 'Monitor that service is in target state' and 'Start and stop service when starting and stopping protected applications'. Below these are three dropdowns for failure actions: 'On first failure:' with 'Log Warning', 'On second failure:' with 'Recover Service', and 'On third failure:' with 'RestartApplications'. At the bottom right are 'OK' and 'Cancel' buttons.

2. Select the service and set the *Target State on Active server* and *Target State on Passive server* values. Normally, the *Target State on Active server* is set to *Running* and the *Target State on Passive server* is set to *Stopped*. User defined services configured with a target state of *Running* on both active and passive servers do not stop when **Stop Applications** is clicked.
3. To make Neverfail Engine monitor the state of the service, select the **Monitor State** check box. To let Neverfail Engine manage the starting and stopping of the service, select the check box. Neverfail Engine also lets you assign three sequential tasks to perform in the event of failure. Task options include the following:
 - Restart Applications – Restarts the protected application.
 - Switchover – Initiates an automatic failover to the currently passive server.
 - Recover Service – Restarts the service.
 - Log Warning – Adds an entry to the logs.
 - A User Defined task, created in the Tasks page, as a Rule Action task type.
 - vSphere Integration\RestartVM – Cleanly shuts down and restarts the Windows OS on the target VM.
 - vSphere Integration\ TriggerMigrateVM – Depending on the parameters specified it can be vMotion, enhanced vMotion or storage vMotion.
 - vSphere Integration\ TriggerMigrateVMandRestartApplications – Same as TriggerMigrateVM + application restart.

- vSphere Integration\ TriggervSphereHaVmReset – Communicates with vCenter Server to reset the virtual machine, but does so using the vSphere HA App Monitoring mechanism. This is potentially more robust, but requires the VM to be on an vSphere HA cluster with Integrate with vSphere HA Application Monitoring enabled in the VmAdaptor plug-in settings.

Rule Action tasks are additional user defined tasks previously created by the user and must be created on the active Neverfail Continuity Engine server

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

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

1.8.2. Edit a Service

About this task

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

Note

Only user defined services can be configured regarding the target state, Monitor State, and Manage Starting and Stopping. The plug-in defined services cannot be edited in this sense. Only their recovery actions can be edited.

Procedure

1. Click the **Edit** button at the lower portion of the pane.

The *Edit Service Protection* dialog appears, which provides a subset of same options available when a new service is added.

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

Figure 2-73. Edit Service Protection

Edit Service Protection

Service: Server (LANMANSERVER)

Target state on Active server: Running

Target state on Passive server: Stopped

☒ Monitor that service is in target state

☒ Start and stop service when starting and stopping protected applications

On first failure: Recover Service

On second failure: Restart Applications

On third failure: Switchover

☒ Allocate entire application time-out when recovering service

OK Cancel

3. To unprotect a User Defined service and stop monitoring the service, click on the Services tab. Select the service and click **Edit**.
4. Clear the *Start and stop service when starting and stopping protected applications* check box, and then click **OK**.

1.8.3. Configure Service Recovery Options for Protected Services

About this task

Neverfail Continuity Engine Management Service provides the ability to configure the Service Recovery Options for services that are protected.

Procedure

1. Navigate to the Services page.
2. Click the **Edit** button. Select the action to take for the 1st, 2nd, and 3rd instance of failure. Click **OK**.

Figure 2-74. Edit Service Protection

The screenshot shows the 'Edit Service Protection' dialog box. At the top, the 'Service:' field is set to 'Server (LANMANSERVER)'. Below this, there are two dropdown menus: 'Target state on Active server:' set to 'Running' and 'Target state on Passive server:' set to 'Stopped'. Underneath these are two checked checkboxes: 'Monitor that service is in target state' and 'Start and stop service when starting and stopping protected applications'. Further down are three dropdown menus for failure actions: 'On first failure:' set to 'Recover Service', 'On second failure:' set to 'Restart Applications', and 'On third failure:' set to 'Switchover'. At the bottom left, there is a checked checkbox 'Allocate entire application time-out when recovering service'. At the bottom right are 'OK' and 'Cancel' buttons.

1.8.4. Remove a Service

About this task

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

Note

Only user defined services can be removed. Plug-in defined services can not be removed.

Procedure

1. Select the user defined service to be removed and click **Remove** at the lower portion of the pane. The user defined service is removed from the list of protected services.

1.9. Data

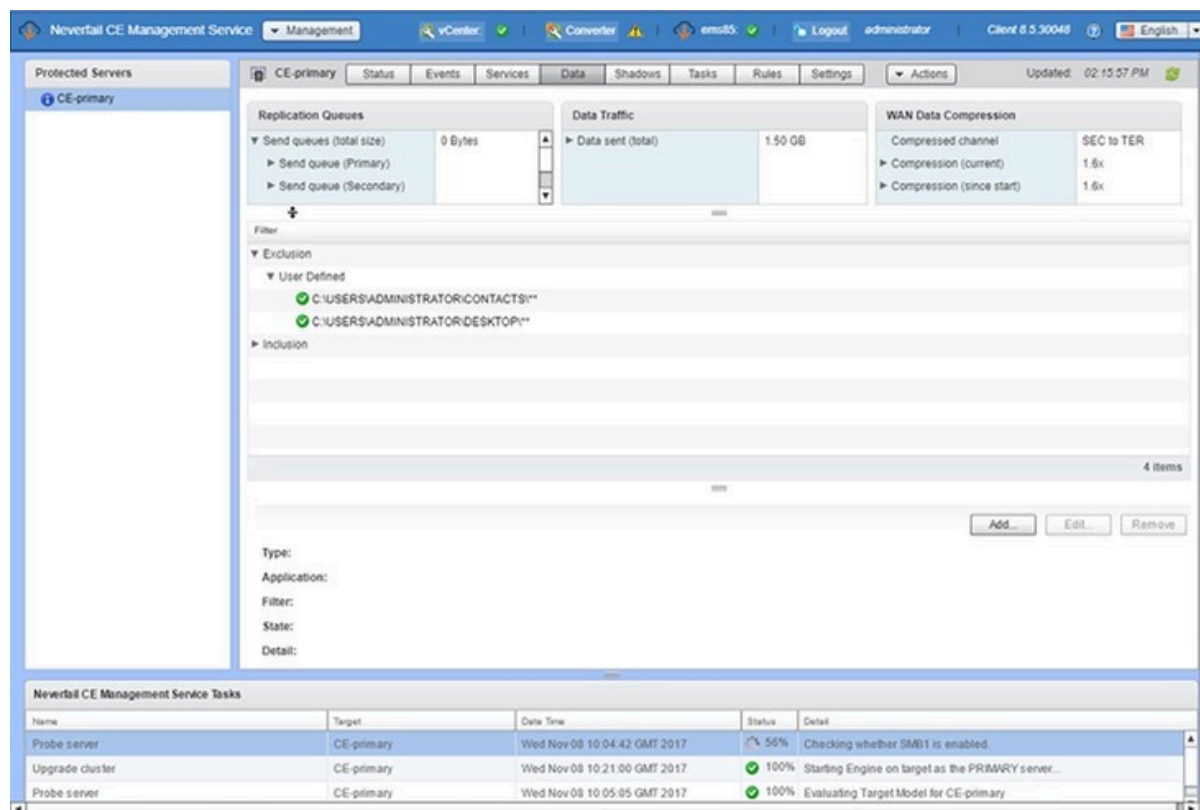
Neverfail Continuity Engine can protect many permutations or combinations of file structures on the active server by the use of custom inclusion and exclusion filters configured by the administrator.

Note

The Neverfail Continuity 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 Queues**.

Figure 2-75. Data page



The *Replication Queues* pane – The statistics of the connection with regards to the data sent by either server and the size of the active server's send queue and passive server's receive queue are displayed.

The *Data Traffic* pane – The Data Traffic displays the volume of data that has been transmitted across the wire from the active server to the passive server.

The *WAN Data Compression* pane – Neverfail Continuity Engine offers WAN Compression as an optional feature to assist in transferring data fast over a WAN. When included in your Neverfail Engine license, WAN Compression can be configured through the **Settings** page. The **Data** page provides a quickly accessible status on the current state of WAN operations, identifies the compressed channel, and displays the amount of compression that is being applied currently and since the start.

1.9.1. Add File Filters

About this task

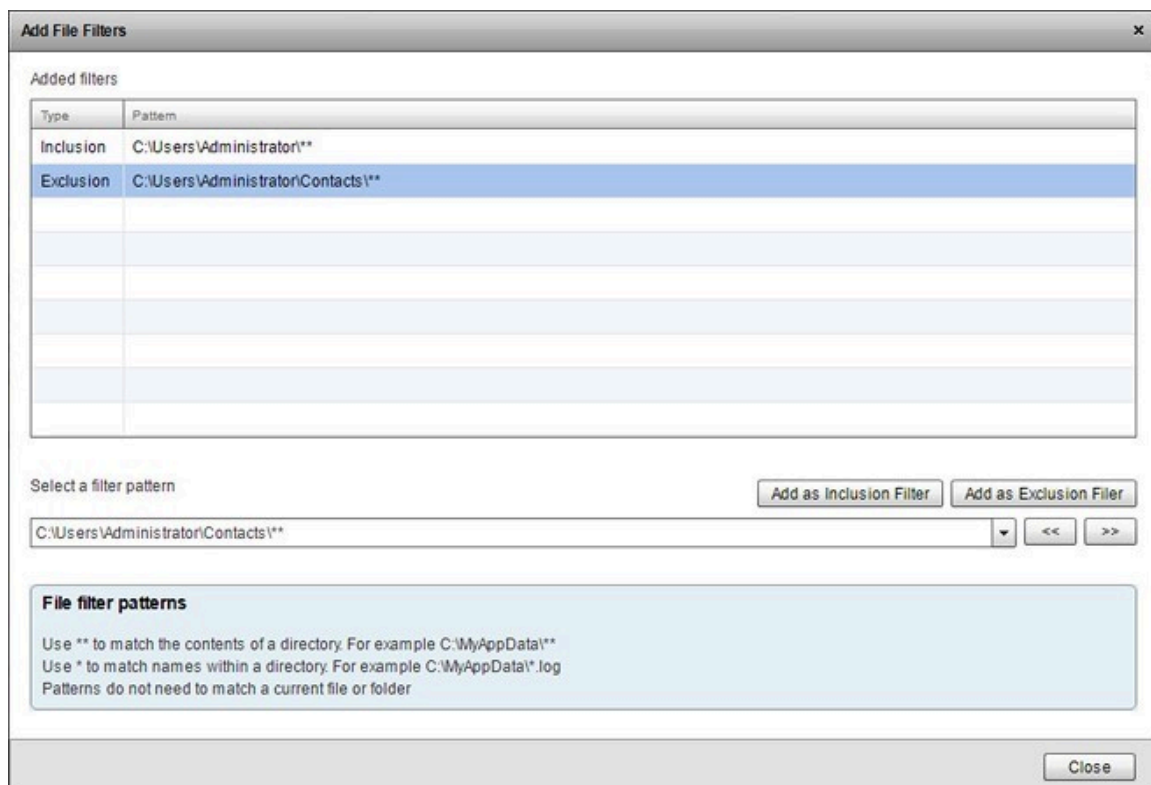
Administrators can add filters to include additional files or folders in the protected set or to exclude a subset of files or folders within the protected set.

To add a user defined File Filter to add or exclude files from the protected set, perform the following steps:

Procedure

1. Click the **Add** button to open the *Add File Filters* dialog.

Figure 2-76. Add Filter



2. Filters to protect user defined files and folders are defined by typing the complete path and pattern or by specifying a pattern containing wildcards.

The filter pattern field will autocomplete any path once the user starts typing. The autocomplete feature searches the Primary node drives for directories in order to suggest the most probable path.

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

Note

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

3. Click **Add as Inclusion Filter** to add the files matching the pattern to the protected files set or **Add as Exclusion Filter** to exclude the files matching the pattern from the protected files set.

1.9.2. Edit Filters

About this task


User defined Inclusion/Exclusion filters can be edited to enable/disable the filter using the Neverfail Continuity Engine Management Service.

To Edit a user defined Inclusion/Exclusion Filter:

Procedure

1. Select the filter and click the **Edit** button located under the filters pane on the *Data* page.

Figure 2-77. Edit Inclusion Filter



2. Edit the value in the *New Filter* text box by typing over the current file filter definition or select to enable/disable the filter.
3. Click **OK**.

Plug-in defined filters can only be edited to enable/disable the filter.

The file filter is changed and becomes active.

1.9.3. Remove Filters

About this task

To Remove a user defined filter:

Note

Plug-in filters can not be removed.

Procedure

1. To remove an Inclusion filter or Exclusion filter, select the filter in the Filter pane and click **Remove**.

1.10. Shadows

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

1.10.1. Best Practices for Using Volume Shadow Copy Service & DRM

About this task

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

Note

Decide which volume to use for storing Shadow Copies before using DRM because you must delete any existing shadow copies before you can change the storage volume. Neverfail recommends that a separate volume be allocated for storing shadow copies. Do not use a volume to store both Neverfail Engine protected data and unprotected, regularly updated data.

Procedure

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

VSS is also used by the Shadow Copies of Shared Folders (SCSF) feature of Windows 2008R2, and consequently, some of the following recommendations are based on Microsoft™ Best Practices for SCSF.

For example: do not write backups of data (even temporarily) to a volume that contains Neverfail Engine protected files, as that increases the space required for snapshots.

2. In accordance with the following guidelines from Microsoft: Select a separate volume on another disk as the storage area for shadow copies. Select a storage area on a volume that is not shadow copied. Using a separate volume on another disk provides two advantages. First, it eliminates the possibility that high I/O load causes deletion of shadow copies. Second, this configuration provides better performance.
3. Be sure to allocate enough space for the retained shadow copies.

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

Note

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

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

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

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

5. Perform a trial-rollback.

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

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

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

Note the following:

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

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

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

If DRM detects the deletion of any expected Shadow Copies, this is noted in the Neverfail Engine *Event Log*. This is an indication that VSS reached its limit of available space or number of Shadow Copies. If many Shadow Copies are automatically discarded, consider adding more storage, or reconfiguring your schedule to create and maintain fewer shadow copies.

1.10.2. Configure Shadow Creation Options

About this task

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

Note

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

Procedure

1. **Create a shadow every...**

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

2. **Create a shadow on the Active once per day at...**

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

Note

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

Whether a shadow is actually taken is also controlled by *Only between the hours:* and *Only on the days:*, if either of these are configured, then a shadow is taken only within the days/hours defined

by them. The following two options limit the number of shadows taken during periods when the data is not changing.

3. Only between the hours...

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

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

4. Only on the days...

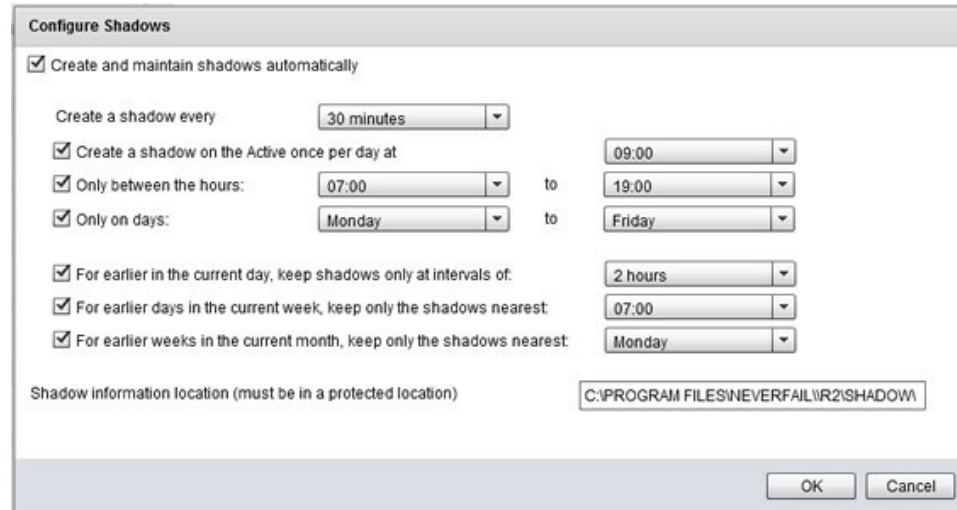
When the check box is selected, the range defined by the two drop down lists is applied to the automatic creation of shadows either on the passive server(s) (as controlled by *Create a shadow every:*) or active server (as controlled by *Create a shadow on the Active once per day at:*).

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

Note

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

Figure 2-78. Shadow Creation Options



1.10.3. Configure the Shadow Copy Schedule

About this task

DRM can create and delete shadow copies automatically according to a configurable schedule. The aim of the schedule is to provide a balance between providing a fine-granularity of rollback points-in-time on the one hand, and conserving disk space and number of shadow copies on the other. To achieve this balance, the available configuration options reflect the observation that recent events generally are

of more interest and value than older ones. For example, the default schedule maintains one shadow from every day of the last week, and one shadow from every week of the last month.

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

Procedure

1. Navigate to the **Shadows** page and click **Configure**.

The *Configure Shadow Schedule* dialog appears.

Figure 2-79. *Configure Shadow Schedule*

Configure Shadows

☒ Create and maintain shadows automatically

Create a shadow every

☒ Create a shadow on the Active once per day at

☒ Only between the hours: to

☒ Only on days: to

☒ For earlier in the current day, keep shadows only at intervals of:

☒ For earlier days in the current week, keep only the shadows nearest:

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

Shadow information location (must be in a protected location)

OK Cancel

2. Select the **Create and maintain shadows automatically** check box.

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

Note

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

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

The Volume Shadow Copy Service (VSS) component of Windows 2008/2012, may automatically delete old shadows because of lack of disk space even when the Create and maintain shadows automatically check box is not selected.

1.10.4. Configure Shadow Keep Options

About this task

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

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

Procedure

1. For earlier in the current day, keep shadows only at an interval of...

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

2. For earlier days in the current week, keep only the shadow nearest...

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

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

3. For earlier weeks in the current month, keep only the shadows nearest...

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

To calculate the “nearest”, an hour is required. The calculation attempts to use the selected time from *For earlier days in the current week, keep only the shadow nearest:* if it is selected, otherwise the *Only between the hours* start time is used if it is selected, finally, when neither of these options are configured, *Midnight* is used.

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

Figure 2-80. Shadow Keep Options

Configure Shadows

☒ Create and maintain shadows automatically

Create a shadow every

☒ Create a shadow on the Active once per day at

☒ Only between the hours: to

☒ Only on days: to

☒ For earlier in the current day, keep shadows only at intervals of:

☒ For earlier days in the current week, keep only the shadows nearest:

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

Shadow information location (must be in a protected location)

OK Cancel

1.10.5. Manually Create Shadow Copies

About this task

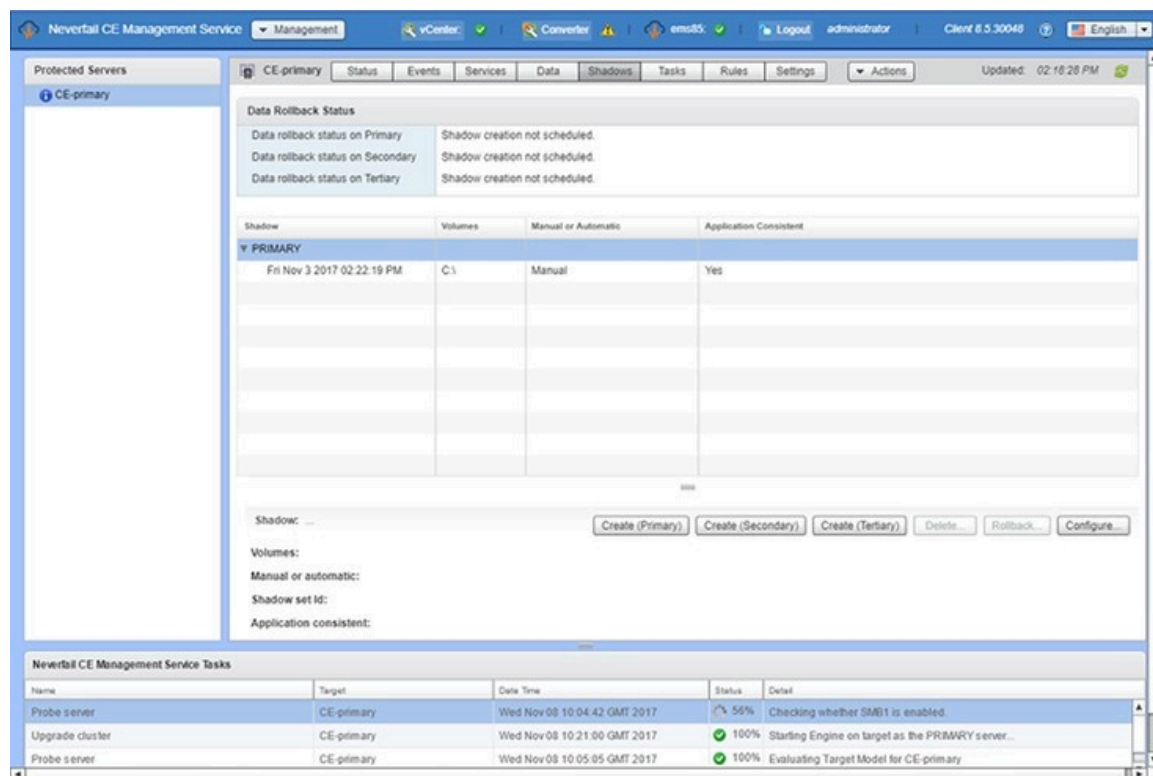
Shadow Copies can be created manually using the steps below:

Procedure

1. In the *Shadow* pane of the *Shadows* page, click **Create (Primary)**, **Create (Secondary)** or if present, **Create (Tertiary)**.

A Shadow Copy is created on the selected node.

Figure 2-81. Create Shadow Dialog



1.10.6. Delete a Shadow Copy

About this task

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

Procedure

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

1.10.7. Roll Back Protected Data to a Previous Shadow Copy

About this task

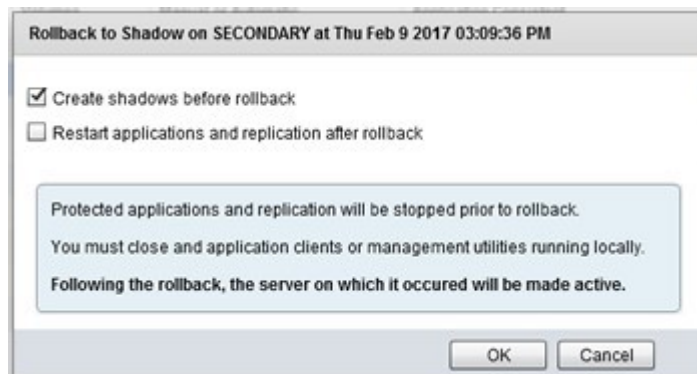
Should the need arise to roll data back to a previous point in time, perform the following:

Procedure

1. Go to the *Shadow* pane of the *Shadows* page and select an existing Shadow from the Primary, Secondary, or Tertiary server list and click **Rollback**.

2. A dialog is presented allowing you to create a shadow immediately before the rollback, and select whether to restart applications and replication after the rollback.

Figure 2-82. Rollback to Shadow dialog



Note

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

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

3. Click **OK**.

A confirmation dialog is presented.

4. Click **Yes**.

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

If the rollback fails, the reason for the failure is shown in the status display. This may be because a particular file set of files or registry key cannot be accessed. For example, a file may be locked because the application is inadvertently running on the server performing the rollback, or permissions may prevent the SYSTEM account from updating. Rectify the problem and try performing the rollback again.

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

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

1.11. Tasks

Tasks are actions which are required for automated application management.

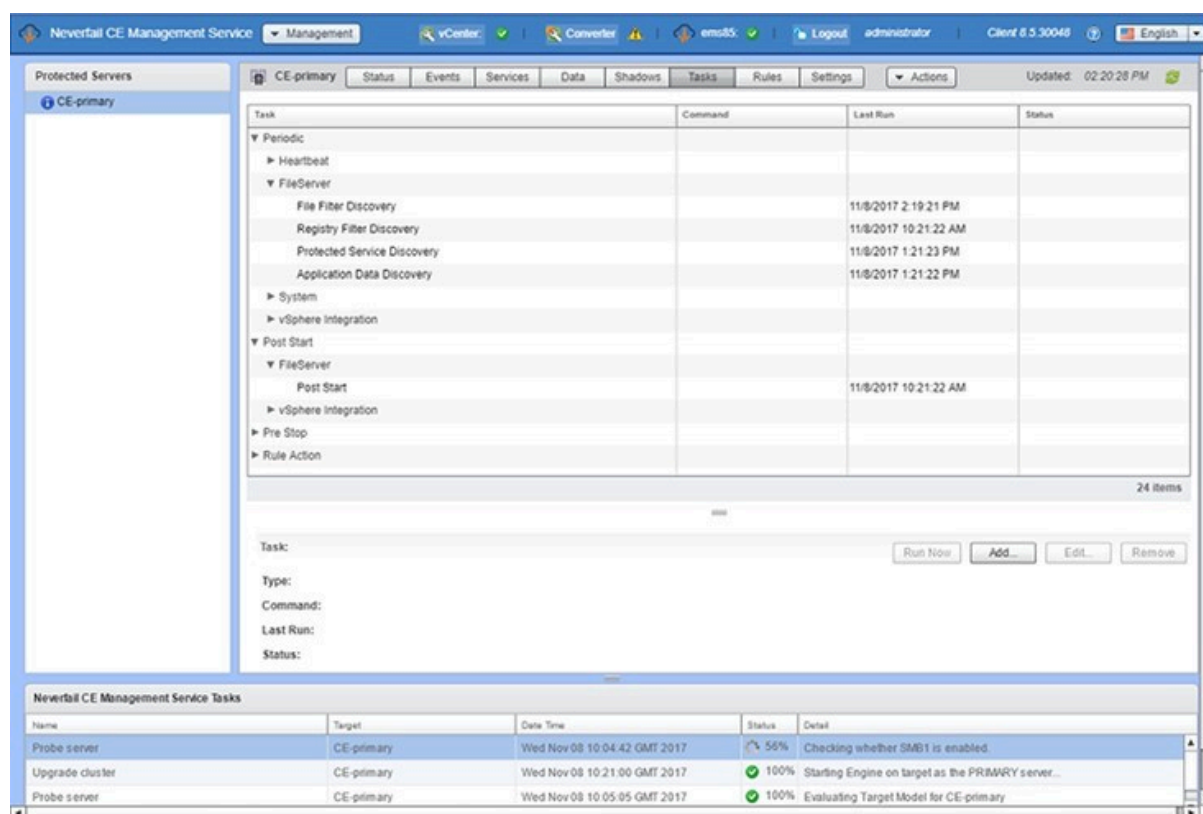
Task types are determined by when the tasks are run, and include the following:

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

Tasks can be defined and implemented by plug-ins or by the user, or they can be built-in tasks defined by Neverfail 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 Continuity Engine Management Service Tasks page provides a list of tasks and associated status information, as well as features to quickly manage tasks.

Figure 2-83. Tasks page



1.11.1. Run Now

About this task

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

Procedure

1. Select an existing task and click **Run Now** at the lower right of the pane.

The task runs. You can watch the Status column of the Task list for messages as the task runs to completion.

1.11.2. Add Task

About this task

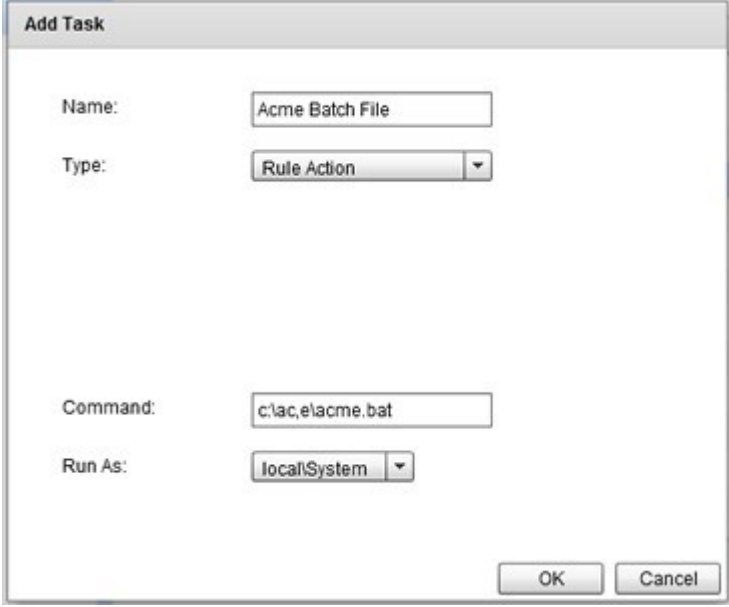
Tasks can be added from the Tasks page of the Neverfail Continuity Engine Management Service. To add a User Defined task:

Procedure

1. Click **Add** at the lower right of the pane.

The *Add Task* dialog appears.

Figure 2-84. Add Task

The image shows a Windows-style dialog box titled "Add Task". It contains four labeled fields: "Name:" with a text box containing "Acme Batch File"; "Type:" with a drop-down menu showing "Rule Action"; "Command:" with a text box containing "c:\acme\acme.bat"; and "Run As:" with a drop-down menu showing "localSystem". At the bottom right are "OK" and "Cancel" buttons.

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

This is required only for Network Configuration tasks.

5. In the *Command* text box, type in the path or browse to the script, .bat file, or command for the task to perform.

When the Command entry requires specific user credentials, you must select that user from the Run As drop-down list.

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

1.11.3. Edit Task

About this task


You can edit the interval, a command, or disable an existing task. To edit a task:

Procedure

1. Click **Edit** at the lower right of the pane.

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

Figure 2-85. *Edit Task*



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

1.11.4. Remove Task

About this task

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

Note

Only user defined tasks can be removed. Plug-in task removal will be vetoed.

Procedure

1. Select an existing task click **Remove** at the lower right of the pane.

A confirmation message appears.

2. Click **Yes** to remove the task, or click **No** to close the message without removing the task.

1.12. Rules

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

Rules use the following control and decision criteria for evaluation:

- Name: (the name of the rule).
- Enabled: (whether the rule is enabled or not).
- Condition: (the condition being evaluated).
- Status: (the current status of the rule being evaluation)
- Triggered: (the condition fails to meet configured parameters resulting in initiation of a duration count)
- Triggered Count: (a count of the number of times the rule has failed)
- Duration: (the length of time the condition exists before triggering the failure action).
- Interval: (the length of time between failure actions).
- First Failure: (action to take upon first failure) The default is set to Log Warning.
- Second Failure: (action to take upon second failure) The default is set to Log Warning.
- Third Failure: (action to take upon third failure) The default is set to Log Warning.

Figure 2-86. Rules page

The screenshot displays the 'Rules' page in the Neverfail CE Management Service interface. The left sidebar shows 'Protected Servers' with 'CE-primary' selected. The main area is divided into tabs: Status, Events, Services, Data, Shadows, Tasks, Rules, Settings, and Actions. The 'Rules' tab is active, showing a table of rules for 'CE-primary'. The rules are categorized under 'FileServer' and 'System', with 'Disk' rules expanded. Each rule entry includes a status icon (green checkmark), a name, a condition, a duration, a current status, a triggered count, and a trigger count. Below the table, there is a 'Rule:' section with fields for Condition, Duration, Status, Triggered, Trigger Count, and On First/Second/Third Failure. At the bottom, the 'Neverfail CE Management Service Tasks' table is visible, showing tasks like 'Upgrade cluster', 'Probe server', and 'Upgrade cluster' with their targets, dates, and statuses.

Rule	Condition	Duration	Status	Triggered	Trigger Count
Free Disk Space	Free disk space < 10 %	600 s	85		0
Free Disk Space On Drive(s)	Free disk space on dr...	600 s	All drives OK		0
Disks Writable	Disk(s) Writable C:		All disks are writable		0
Disk IO	Disk Usage: Time > ...	600 s	OK		0
Disk Reads Per Sec	Disk Reads / sec > 0	1500 s			0
Disk Writes Per Sec	Disk Writes / sec > 0	1500 s			0
Disk Queue Length	Current Disk Queue L...	1800 s	0		0
Disk Avg Secs Per Read	Average Seconds / R...	1800 s			0

Name	Target	Date Time	Status	Detail
Upgrade cluster	CE-primary	Wed Nov 08 10:21:00 GMT 2017	100%	Starting Engine on target as the PRIMARY server...
Probe server	CE-primary	Wed Nov 08 10:05:05 GMT 2017	100%	Evaluating Target Model for CE-primary
Upgrade cluster	CE-primary	Wed Nov 08 10:21:00 GMT 2017	100%	Starting Engine on target as the PRIMARY server

1.12.1. Check a Rule Condition

To check a rule condition, select the rule in the *Rules* page and click **Check Now** on the lower right portion of the page.

Neverfail Engine immediately checks the rule conditions of the current configuration against the attributes of the system and application.

1.12.2. Edit a Rule

About this task

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

To Edit a rule:

Procedure

1. To edit a rule, select the rule in the *Rules* list.
2. Click **Edit** at the lower right of the page.

The *Edit Rule* dialog appears.

Figure 2-87. Edit Rule dialog

Edit Rule

Name: Free Disk Space

☒ Enabled

Condition: Free disk space < 10 %

Duration: 600 seconds

Interval: 60 seconds

On First Failure: HeartbeatLog Warning

On Second Failure: HeartbeatLog Warning

On Third Failure: HeartbeatLog Warning

OK Cancel

Use this dialog to Enable or Disable a Rule, set the specific options for the Rule, and to assign tasks to perform *On First Failure*, *On Second Failure*, and *On Third Failure*. The following tasks can be assigned in the event of a failure:

- **Recover Service** – Restarts the service.
- **Restart Applications** – Restarts the protected application.
- **Log Warning** – Adds an entry to the logs.
- **Switchover** – Initiates a switchover to the currently passive server.
- **Rule Action** – Executes the command or script previously defined as a *Rule Action* task.

If the installed servers are in a virtual to virtual configuration, the following additional tasks are available as a result of the vSphere Integration Plug-in.

- **vSphere Integration\RestartVM** — Cleanly shuts down and restarts the Windows OS on the target VM
- **vSphere Integration\ TriggerMigrateVM** — Depending on the parameters specified it can be vMotion, enhanced vMotion or storage vMotion
- **vSphereIntegration\TriggerMigrateVMandRestartApplication** — Same as TriggerMigrateVM + application restart
- **vSphere Integration\ TriggervSphereHaVmReset** — Hard Reset of the VM implemented by integration with VMware HA

Note

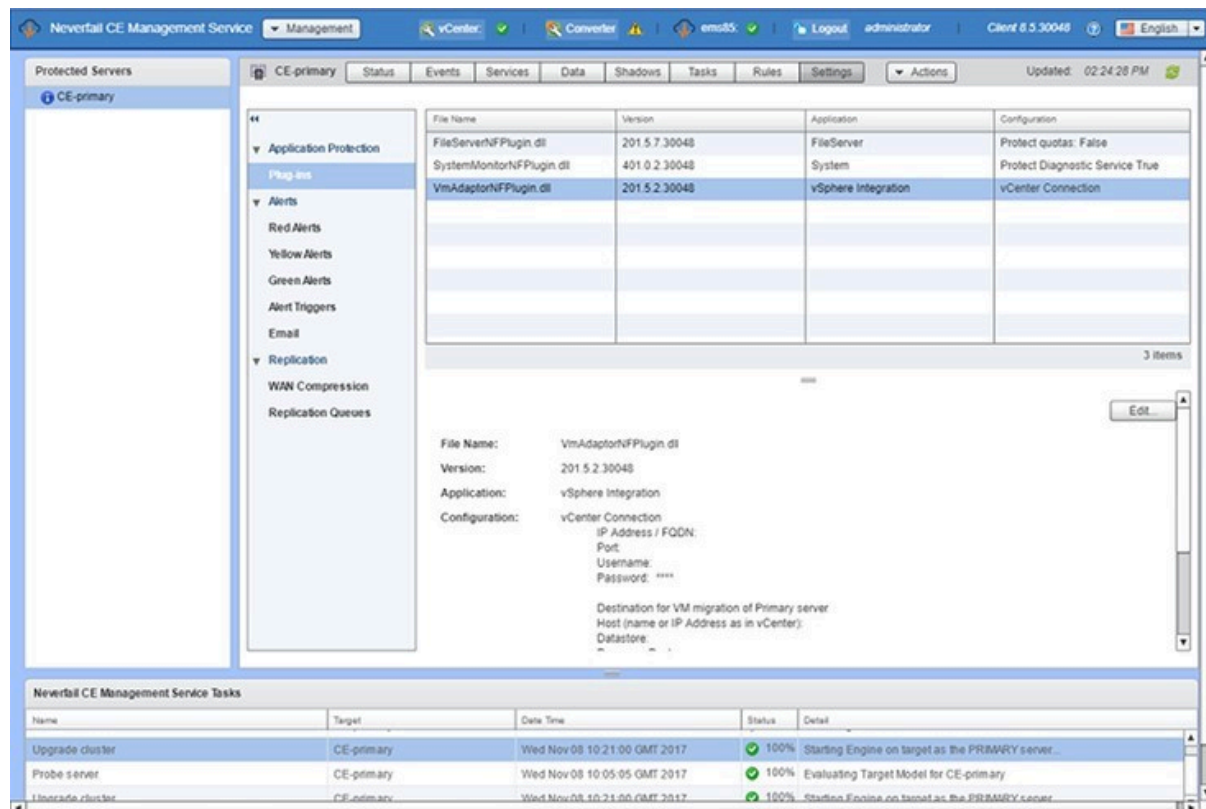
This option requires vSphere HA Application monitoring for the cluster and VM.

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

1.13. Settings

The *Settings* page contains features to configure Plug-ins, Alerts, Email, WAN Compression and Replication Queue settings.

Figure 2-88. Settings page



1.13.1. Configure Plug-ins

About this task

The Neverfail Continuity Engine Management Service allows you to edit the configuration of user installed plug-ins.


To edit an existing plug-in, select *Plug-ins* in the left pane and then select the intended Plug-in from the *Plug-ins* list and perform the following steps:

Procedure

1. Click the **Edit** button on the right side of the *Plug-in Detail* pane.

The *Edit Plug-in* dialog appears.

Figure 2-89. Edit Plug-in dialog



Note

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

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

1.13.2. Alert Settings

About this task

The *Settings* page lets you configure the Neverfail Engine server to send predefined alerts to remote Neverfail Engine administrators via email. The process for adding recipients is the same for all three trigger levels.

Procedure

1. Select the type of alert (Red, Yellow, and Green) in the left pane resulting in the *Alert Settings* pane displaying for the selected alert.
2. Click the **Edit** button in the upper right portion of the *Alert Settings* pane.

Figure 2-90. Alert Settings

Edit Red Alert Settings

Send Mail ☒

☒ Always
☐ Once
☐ Once Per Days

Mail Recipients admin@acme.bat

Mail Subject Neverfail Continuity Engine Red Alert from \${EventHostName} (\${EventHostId}): \${EventName}

Mail Content Neverfail Continuity Engine Red Alert: \${EventName}. This happened at \${EventTime} on the \${EventHostId} \${EventHostName} while

Run Command ☐

Command

OK Cancel

3. Select the *Send mail* check box.
4. Select how many times to send the email (*Always*, *Once*, or *Once per [user configurable time period]*).
5. Enter a recipient's fully qualified email address into the *Mail Recipients* text box. Add additional recipients separated by a semi-colon.
6. Repeat step 4 to until all recipients have been added.
7. The Subject and Content of the alert emails for all three alerts can be adjusted to suit the environment. Neverfail recommends using the pre-configured content and adding customized content as needed.

When Send mail is selected, there are three alternatives:

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

1.13.3. Using WScript to Issue Alert Notifications

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

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

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

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

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

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

1.13.4. Alert Triggers

Select *Alert Triggers* under Alerts in the left pane of the *Settings* page to view the currently configured alert triggers.

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

To modify the current configuration, click the **Edit** button in the upper left portion of the *Alert Triggers* pane. Each alert can be re-configured to trigger as a red, yellow, or green alert or no alert by selecting or clearing the appropriate check boxes. After the alert trigger levels are defined, click **OK** to save the configuration.

Figure 2-91. Edit Alert Triggers

Event	Trigger Red Alert	Trigger Yellow Alert	Trigger Green Alert
▼ Application			
Task Error Output	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Stopping Applications	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Application Warning	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Service Status Info	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Autoswitch Requested	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Application Error	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Starting Applications	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Timeout in Starting/Stopping Applications	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
▼ Channel			
Neverfail Channel connection has been lost	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Advanced compression resource allocated.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
A channel has connected	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Standard compression interface initialized.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Advanced compression interface not initialized.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Standard compression not initialized.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Exception in advanced compression.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
There is no available disk space for queued file/registry update data	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Exception in standard compression.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Failed to establish the Neverfail Channel	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

OK Cancel

1.13.5. Email Settings

Neverfail Engine can alert the administrator or other personnel and route logs via email when an Alert condition exists. To configure this capability, in the *Settings* page, select *Email* in the left pane and click the **Edit** button in the upper right of the *Email Settings* pane.

Figure 2-92. Email Settings

Edit Email Settings	
Outgoing Mail Server for Primary Server	smtp1.mail.server
Outgoing Mail Server for Secondary Server	smtp2.mail.server
Outgoing Mail Server for Tertiary Server	
Send Mail As	user@mail.server
Mail Server Requires Authentication	<input checked="" type="checkbox"/>
Username	administrator
Password	*****
Enable SSL	<input checked="" type="checkbox"/>

OK Cancel

In the *Edit Email Settings* dialog, enter the Outgoing mail server (SMTP) of each server in the Cluster. Enter the mail server name using its fully qualified domain name. Next, configure the default *Send Mail* as email address. This can be customized but the email address used must be an email account authorized to send mail through the SMTP server.

Note

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

Where SMTP servers require authentication to accept and forward SMTP messages, select the *Mail Server requires authentication* check box and specify the credentials for an appropriate authenticated user account.

The **Enable SSL** option can be checked to use SSL for accessing the SMTP server.

Click **OK** to save the changes or click **Cancel** to close the dialog without making any changes.

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

1.13.6. Wan Compression

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

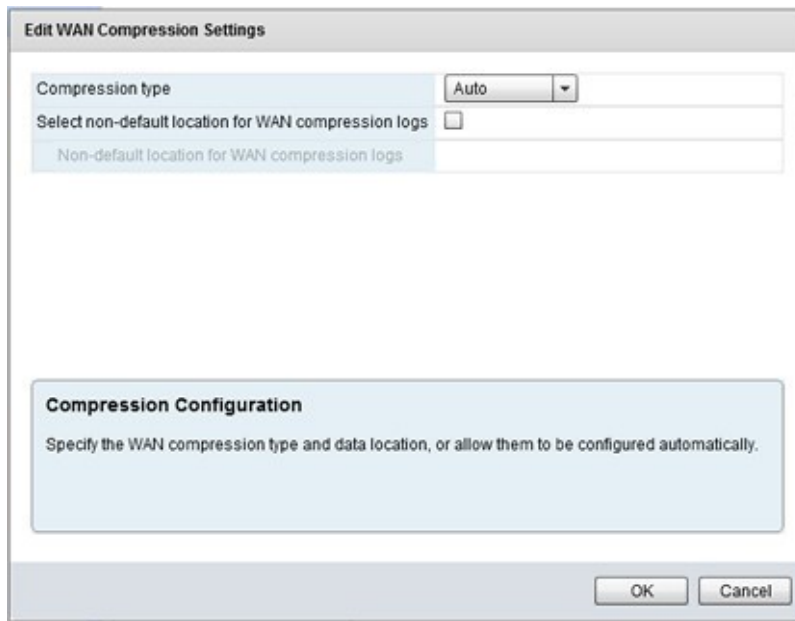
Note

Enabled compression type – Auto – is the recommended setting.

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

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

Figure 2-93. Edit WAN Compression dialog



Edit WAN Compression Settings

Compression type: Auto

Select non-default location for WAN compression logs: ☐

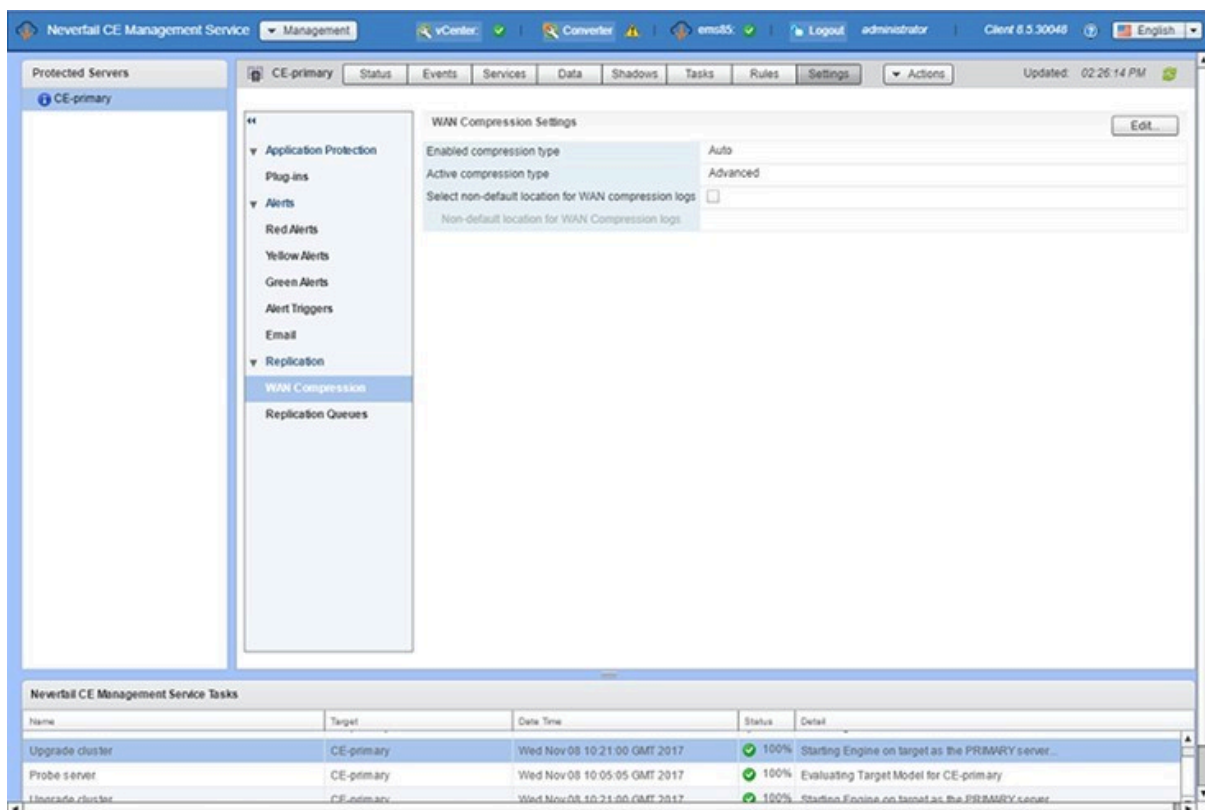
Non-default location for WAN compression logs:

Compression Configuration

Specify the WAN compression type and data location, or allow them to be configured automatically.

OK Cancel

Figure 2-94. WAN Compression page



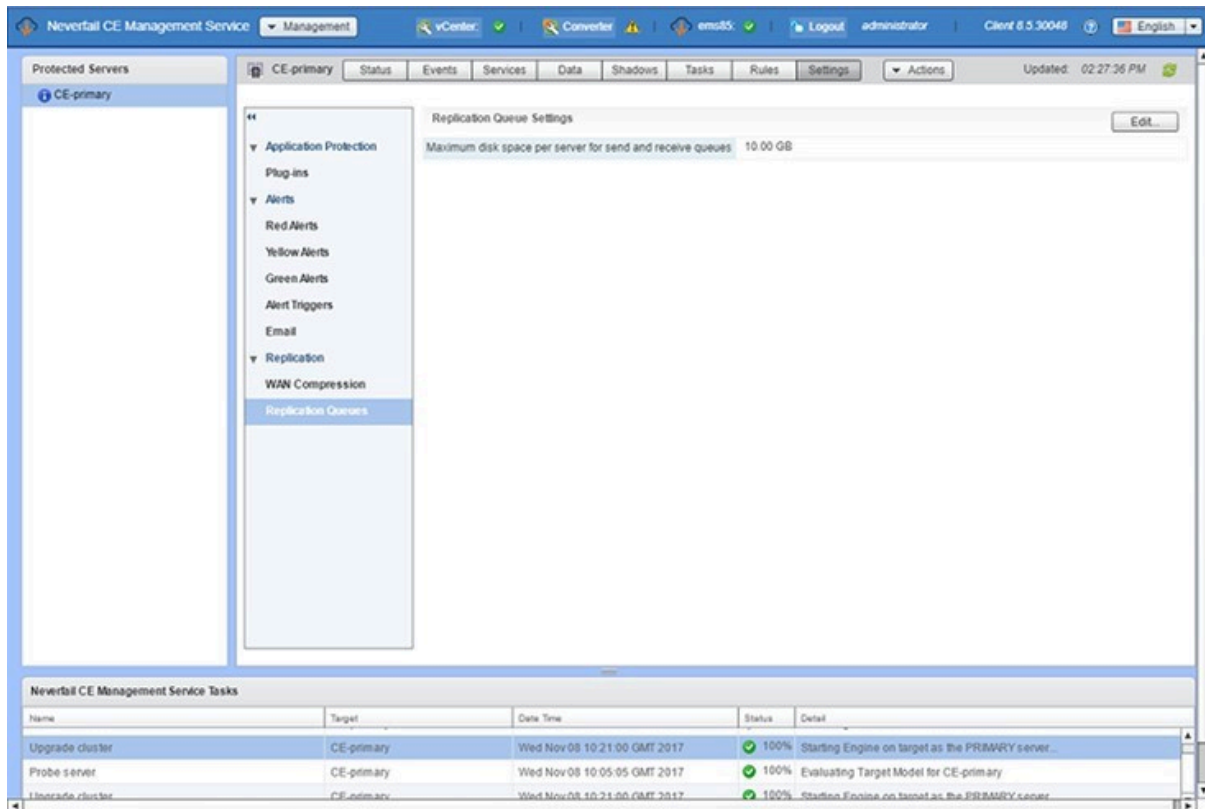
The screenshot shows the 'WAN Compression Settings' page for the 'CE-primary' server. The left sidebar contains a tree view with 'WAN Compression' selected under the 'Replication' category. The main content area displays the settings for WAN compression, including 'Enabled compression type' (Auto), 'Active compression type' (Advanced), and 'Select non-default location for WAN compression logs' (checkbox). The bottom of the interface shows a table of tasks.

Name	Target	Date Time	Status	Detail
Upgrade cluster	CE-primary	Wed Nov 08 10:21:00 GMT 2017	100%	Starting Engine on target as the PRIMARY server...
Probe server	CE-primary	Wed Nov 08 10:05:05 GMT 2017	100%	Evaluating Target Model for CE-primary
Upgrade cluster	CE-primary	Wed Nov 08 10:21:00 GMT 2017	100%	Starting Engine on target as the PRIMARY server...

1.13.7. Replication Queue Settings

The *Settings* page displays the size of the replication queues configured on each server in the cluster.

Figure 2-95. Configured Queue Size

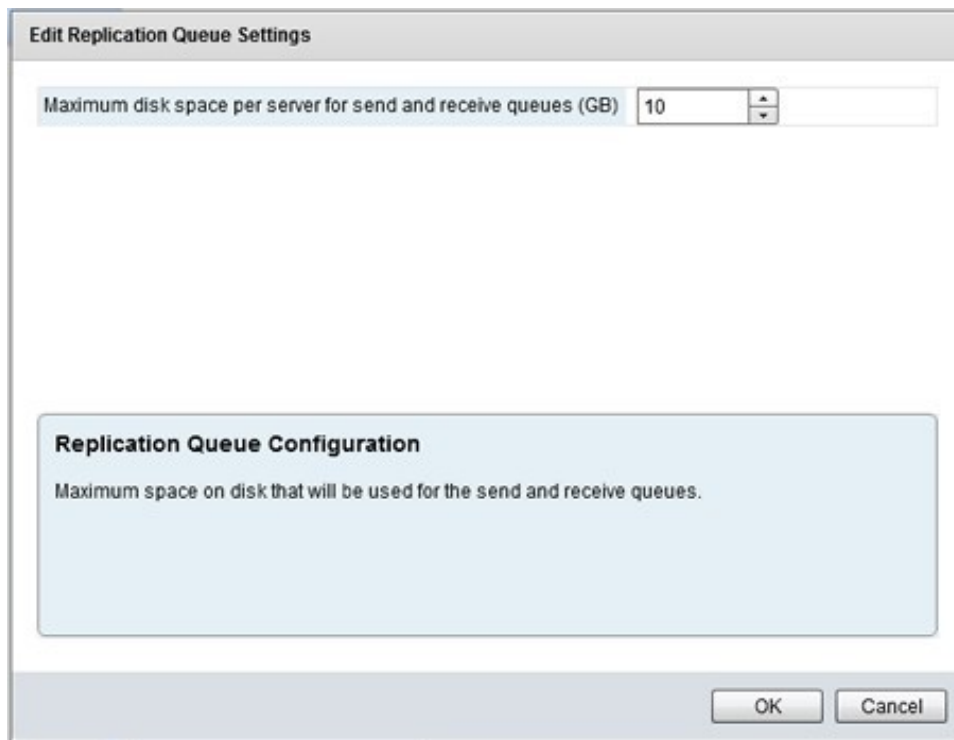


The Edit Replication Queue Settings dialog allows you to configure the maximum disk space per server for the Send and Receive queues on each server.

To configure the **maximum** disk space to be used for the Send and Receive queues:

1. Click the **Edit** button.
2. Enter the maximum disk space to reserve for the Send and Receive queue.
3. Click **OK**.

Figure 2-96. Edit Replication Queue Settings dialog

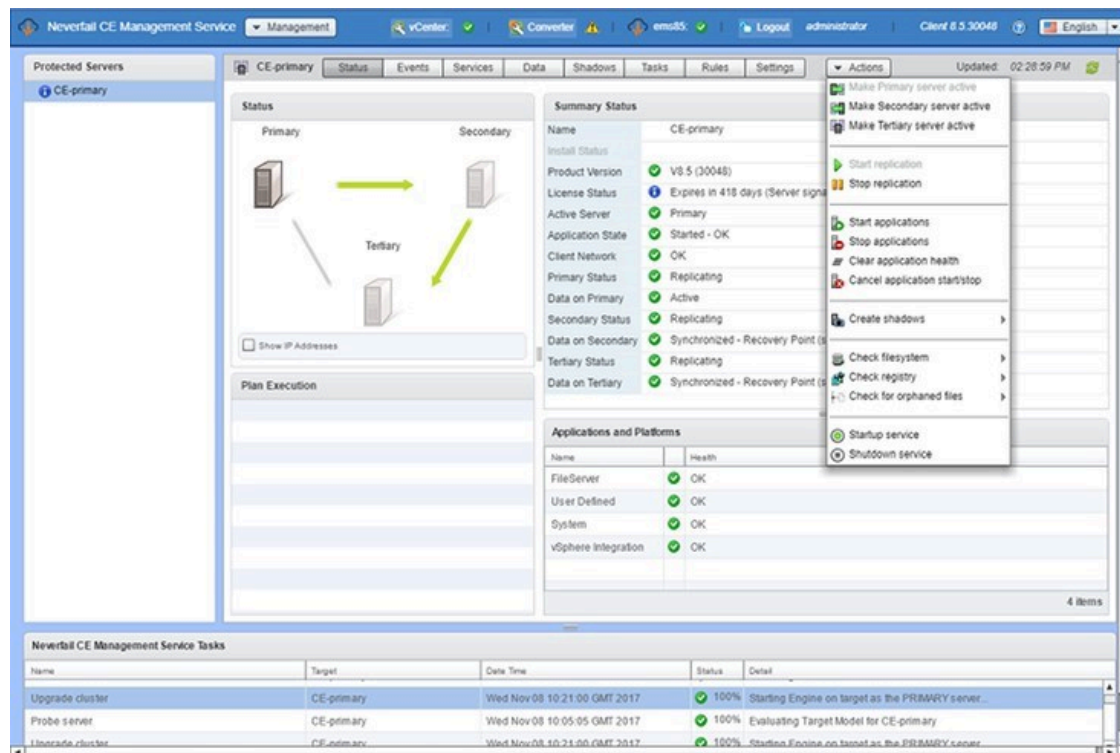


1.14. Actions

The *Actions* drop-down pane provides the ability to Control Neverfail Engine using the Neverfail Continuity Engine Management Service.

The Neverfail Continuity EngineManagement Service provides the ability to perform the main operations, comprising a Switchover, Start/Stop Replication, Start/Stop Applications, Cancel application start/stop, Create Shadows, Check file and registry system, and Startup/Shutdown of Neverfail Engine.

Figure 2-97. Actions drop-down pane



Perform a Switchover

- To make the Primary server of the Neverfail cluster active, click the **Make Primary Server Active** button. The *Make Primary Server Active* dialog asks you to verify that you want to make the Primary server active. Click **OK** to make the Primary Server Active.
- To make the Secondary server of the Neverfail cluster active, click the **Make Secondary Server Active** button. The *Make Secondary Server Active* dialog asks you to verify that you want to make the Secondary server active. Click **OK** to make the Secondary Server Active.
- To make the Tertiary server of the Neverfail cluster active, click the **Make Tertiary Server Active** button. The *Make Tertiary Server Active* dialog asks you to verify that you want to make the Tertiary server active. Click **OK** to make the Tertiary Server Active.


Start Replication

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

Stop Replication

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

Figure 2-98. Stop Replication



Start Applications

When protected applications are stopped, click the **Start Applications** to start the protected applications once again.

Stop Applications

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

Clear Application Health

To reset the health status displayed in the *Summary* pane, click the **Clear Application Health** button. The health status is reset to green.

Cancel Application Start/Stop

To stop the start or stop processes of the protected applications, click the **Cancel Application Start/Stop** button in the *Actions* menu. The Engine Management System checks every 10 seconds for cancellation. If a application start/stop cancellation has been requested, the Engine Management System behaves as if the start/stop operation has immediately timed out.

- Canceling an application start operation will put the application in the **Unmanaged - Unmonitored** state.
- Canceling an *application stop operation* will have additional behavior compared to canceling an application start operation. When an application stop is canceled and times out, then the plan is considered to have failed, and a recovery plan will be generated. The recovery plan will attempt to start applications again. This behavior is intended to attempt to maintain application service even following failed auto-switchovers.

Create Shadows

To manually create a shadow copy on a designated node, navigate to **Actions > Create Shadows** and then select the designated node, **Create (Primary)**, **Create (Secondary)** or if present, **Create (Tertiary)**.

Check File System, Registry System, or Check for Orphaned Files

To manually check the files system, registry, or for orphaned files, navigate to **Actions** drop-down and select the system to check and then select the designated node, for example **Check Primary file system**, **Check Secondary file system** or if present, **Check Tertiary file system**.

Startup Service

Neverfail Engine can be started by logging on to the Neverfail Continuity Engine Management Service and selecting **Startup Service** from the **Actions** drop-down. The *Startup Options* dialog is displayed. Select one or more servers in the Neverfail cluster to start. Click **OK** to start Neverfail Engine on the selected servers in the cluster.

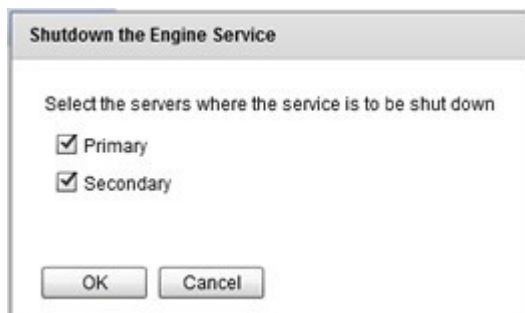
Figure 2-99. Startup Services



Shutdown Service

To shutdown Neverfail Engine, click **Shutdown Service** from the *Actions* button. The *Shutdown Options* dialog is displayed. Select one or more servers in the Neverfail cluster to shutdown. Click **OK** to stop Neverfail Engine on the selected servers in the cluster.

Figure 2-100. Shutdown



2. Managing Neverfail Continuity Engine Clusters

Neverfail Continuity 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

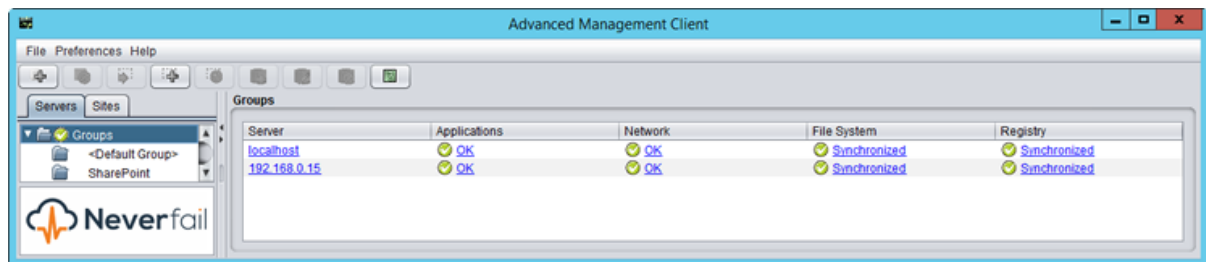
The controlling workstation must have Engine Management Service or Neverfail Advanced Management Client. The Advanced Client can be downloaded from Engine Management Service UI.

3. Review the Status of Neverfail Continuity Engine Clusters and Groups

Procedure

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.

Figure 2-101. Neverfail Engine Servers Overview page



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

2. Click on either:

Option	Description
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

4. Exit Neverfail Advanced Management Client

Procedure

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*.

5. Shutdown Windows with Neverfail Continuity 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.

6. Controlled Shutdown

About this task

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.

Procedure

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 on Windows Server 2008 and 2012 but is not configurable on Windows Server 2003.

Figure 2-102. Controlled Shutdown

Replication Configuration

Fast Check | **Controlled Shutdown** | Orphaned Files

Remain in sync, when the server shuts down (or is restarted) during specified hours.
Shutdown may be delayed up to a timeout while applications stop and replication data is sent.

Use Controlled Shutdown On

☒ Active Server
☐ Passive Server
(Passive Server shutdown will stop applications on Active Server)

Maximum Time To Shutdown

Primary: 50 secs
Secondary: 50 secs

Controlled Shutdown Schedule

	07:00	08:00	09:00	10:00	11:00	12:00	13:00	14:00	15:00	16:00
Sunday										
Monday										
Tuesday										
Wednesday										
Thursday										
Friday										
Saturday										

Controlled Shutdown

OK Cancel

- 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](#).

Chapter 3. Configuring Neverfail Continuity Engine

Related information

- [Configure Server Wizard](#)
- [Configure Machine Identity](#)
- [Configure Server Role](#)
- [Change the Client Connection Port](#)
- [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 Continuity Engine License Key](#)
- [Configure the Message Queue Logs](#)
- [Configure Maximum Disk Usage](#)

1. Configure Server Wizard

Before you begin

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).

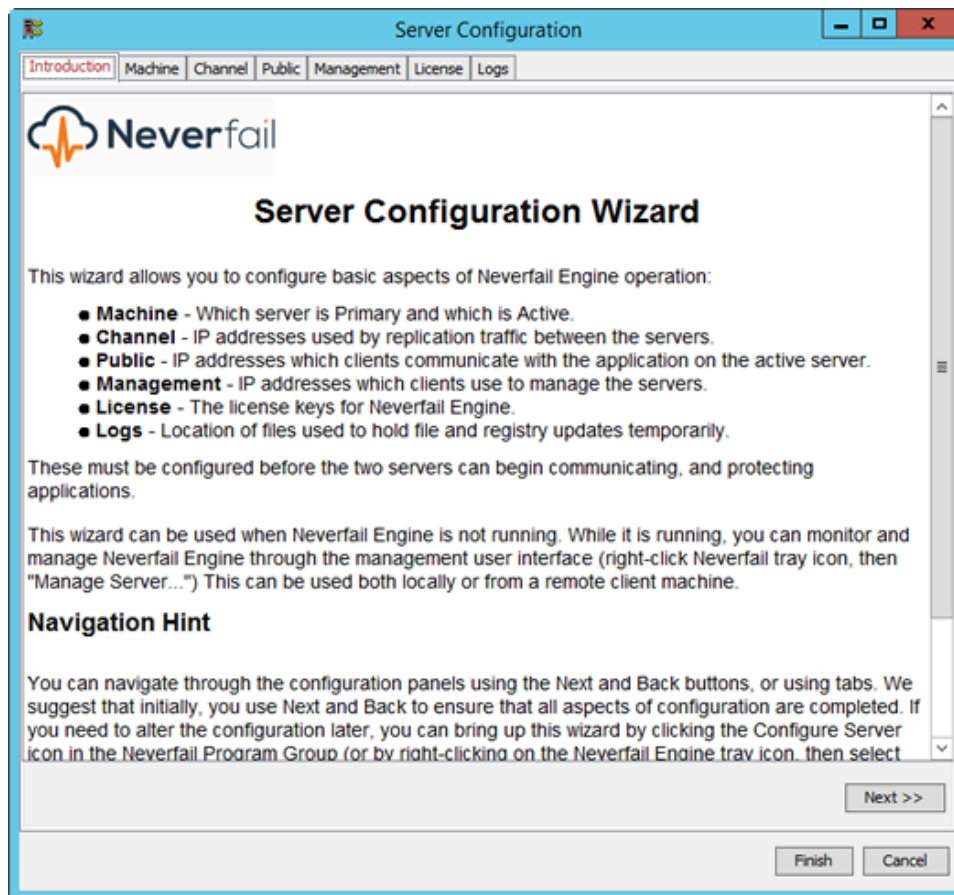
About this task

The Neverfail Continuity 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).

Procedure

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

Figure 3-1. Configure Server Wizard - Introduction Tab



2. Configure Machine Identity

Before you begin

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).

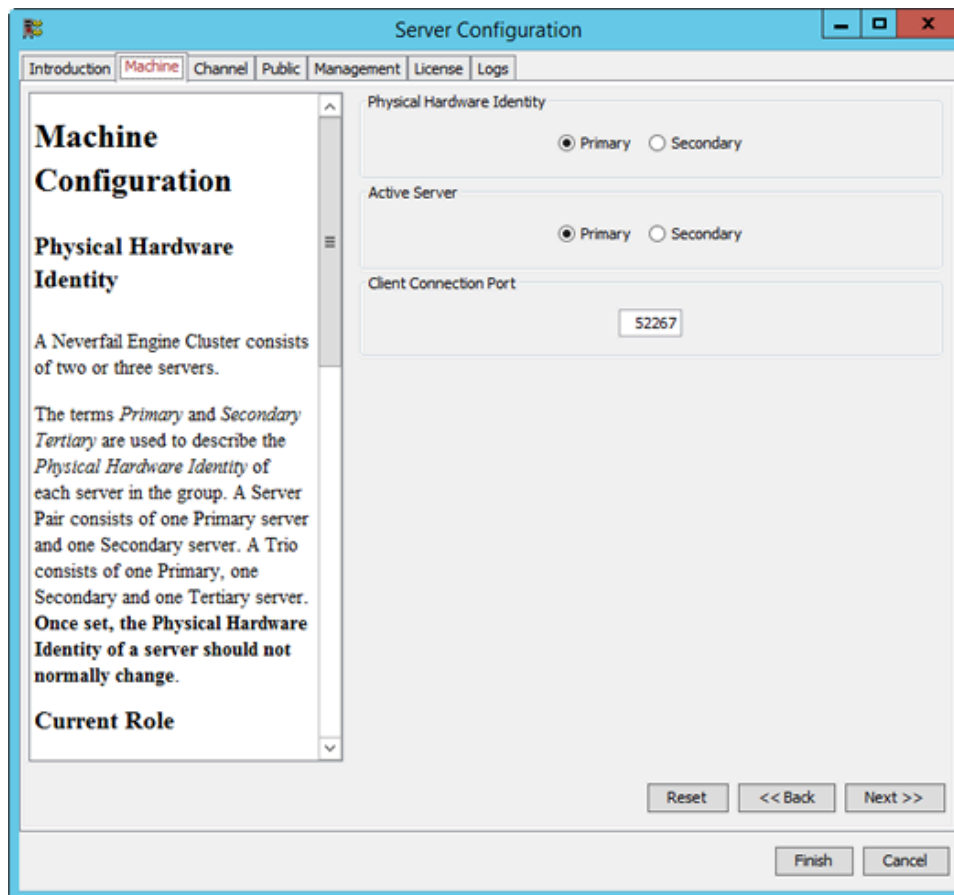
About this task

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).

Procedure

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**.

Figure 3-2. Configure Server Wizard - Machine Tab



3. Configure Server Role

Before you begin

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

About this task

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

Procedure

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**.

4. Change the Client Connection Port

About this task

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

Procedure

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.

5. Configure Channel IP Routing


About this task

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.

Procedure

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

Figure 3-3. Configure Server Wizard - Channel IP Routing



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.

6. Configure the Default Channel Port

About this task

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.

Procedure

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.

7. Configure Low Bandwidth Optimization

About this task

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.

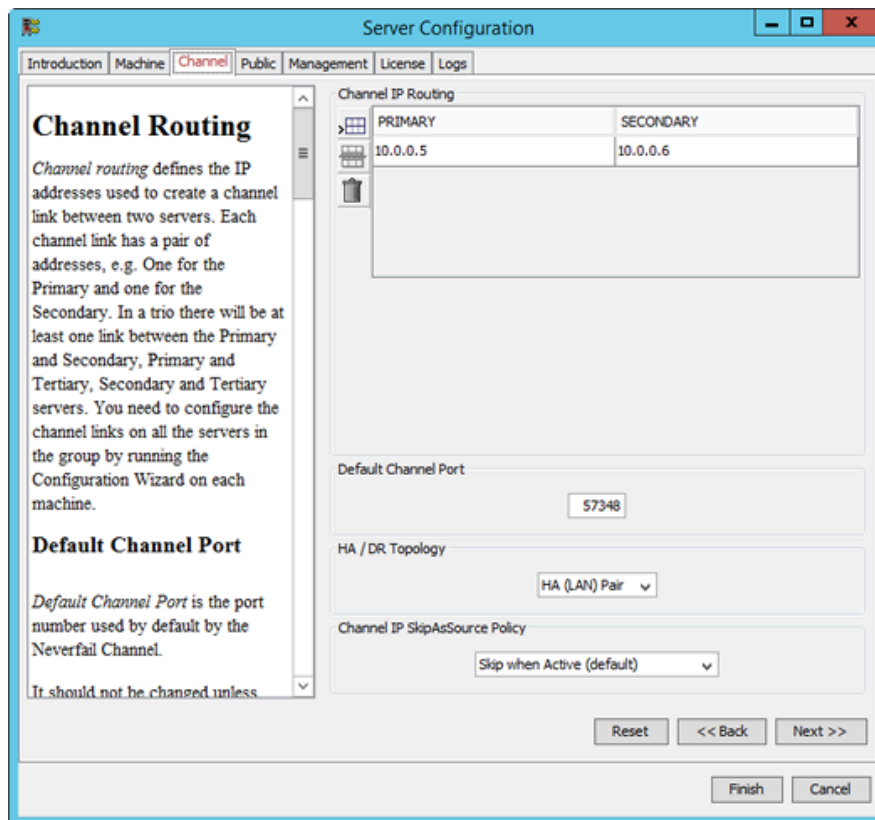
Note

The same HA/DR configuration must be set on all servers in the pair.

Procedure

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.

Figure 3-4. Configure Server Wizard - Channel tab



8. Configure Public IP Addressing

About this task

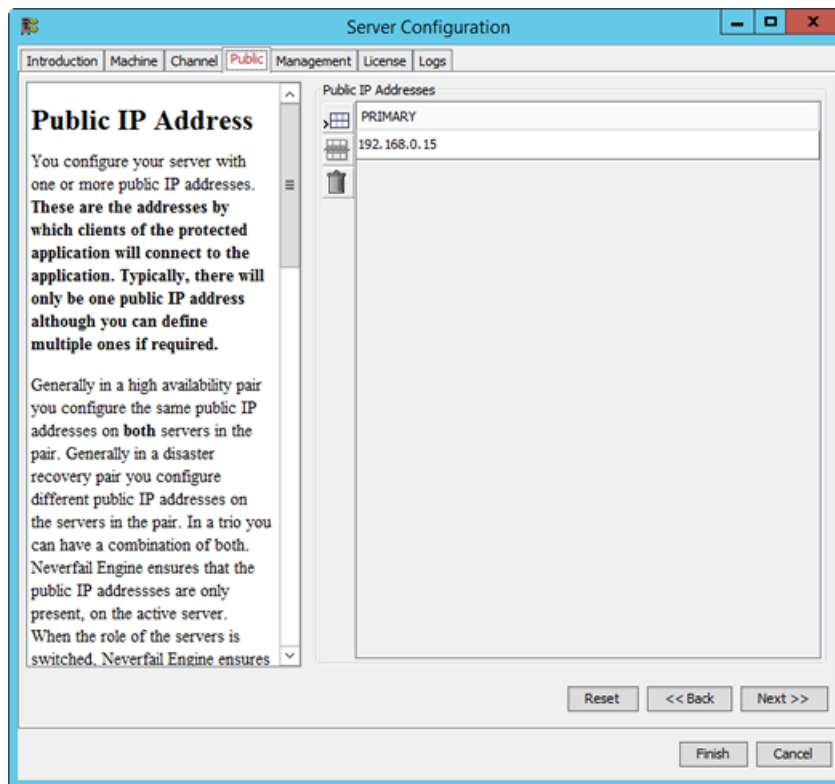
A typically configured Neverfail Engine Cluster uses only one Public IP address when deployed as a pair or on a LAN, but can be configured with more than one Public IP address. These are the addresses by which clients of the protected application connect to the application. Typical installations configure 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.

Procedure

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.

Figure 3-5. Configure Server Wizard - Public Tab

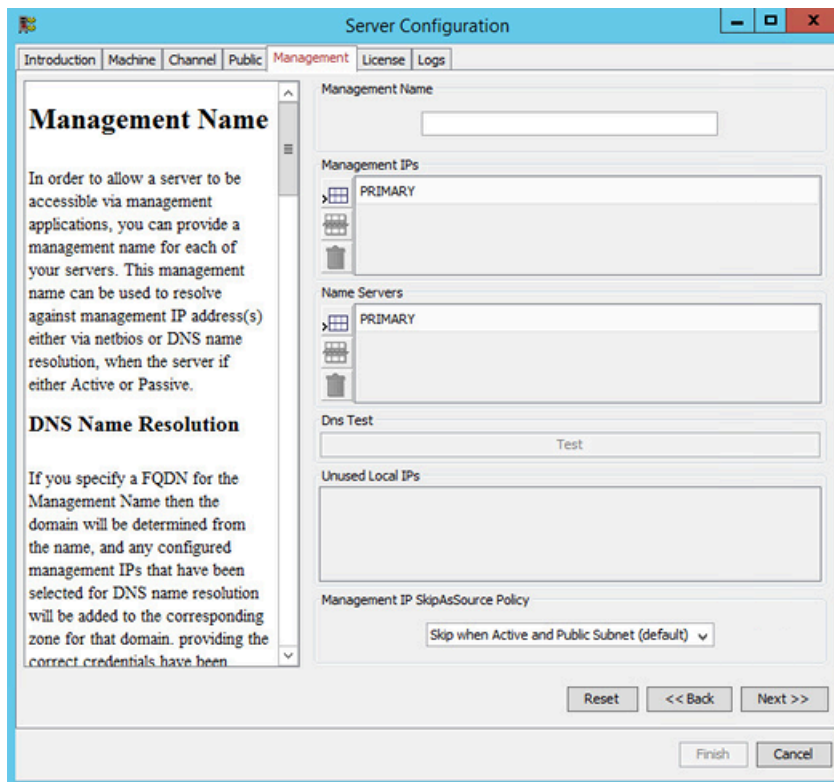


2. To add an address, double-click a row and manually type in the address or select one from a list of currently defined addresses. Click **Next** or **Finish** to accept changes.

9. 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.

Figure 3-6. Configure Server Wizard Management IP Addresses



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.

10. 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: [Deploying updates to passive servers using 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](#).

11. Add/Remove a Neverfail Continuity Engine License Key

About this task

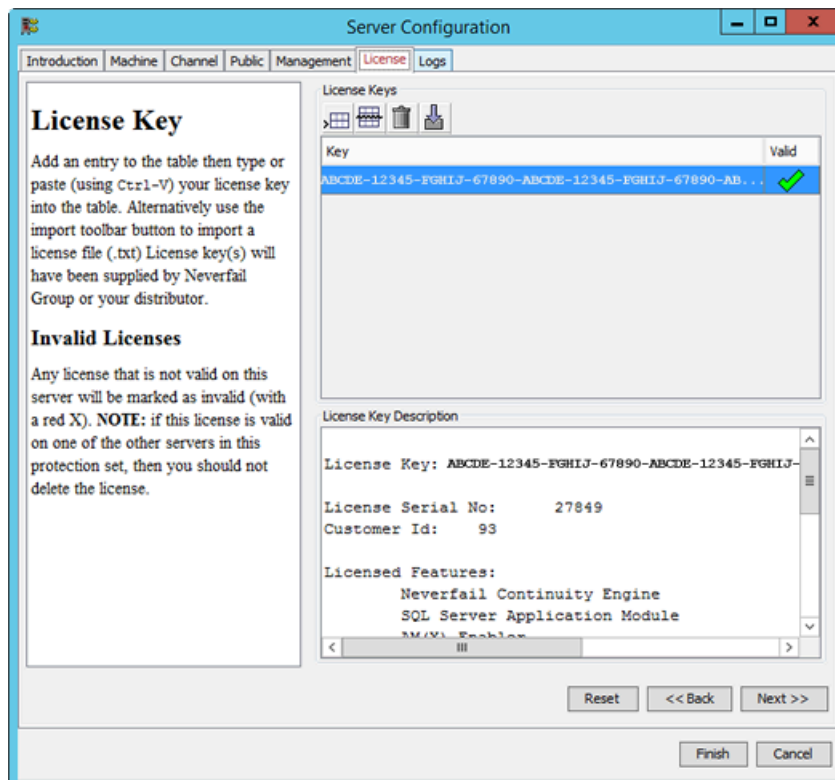
Neverfail recommends using the Engine Management Service user interface for licensing Neverfail Engine (see the Installation Guide).

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

Procedure

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

Figure 3-7. Configure Server Wizard - License Tab



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

12. Configure the Message Queue Logs

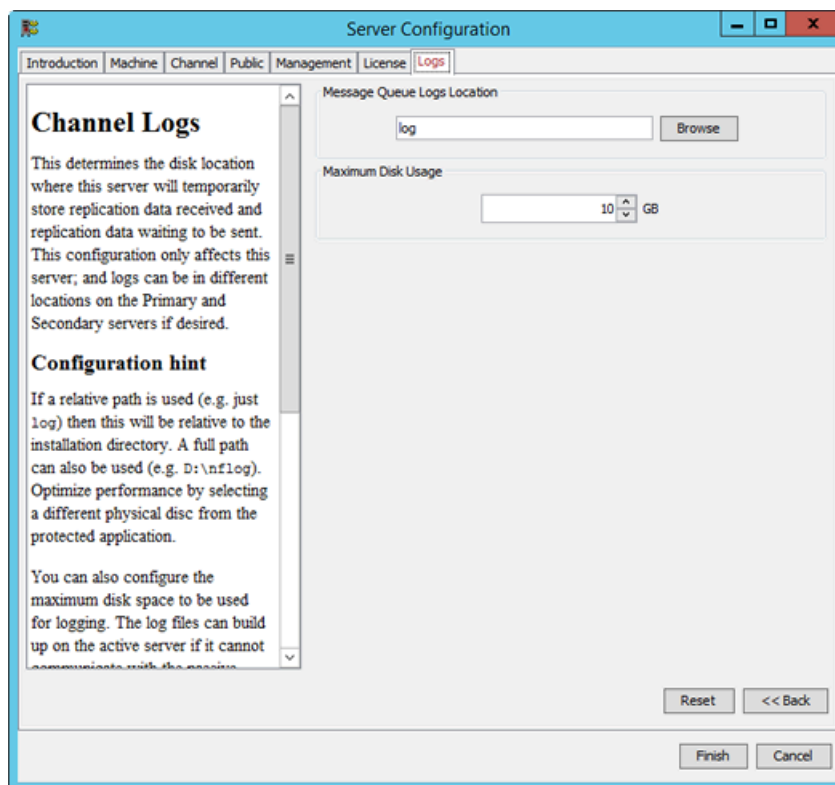
About this task

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.

Procedure

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.

Figure 3-8. *Configure Server Wizard - Logs Tab*



13. Configure Maximum Disk Usage

About this task

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

Procedure

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.

Part II. Management

Chapter 4. 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.

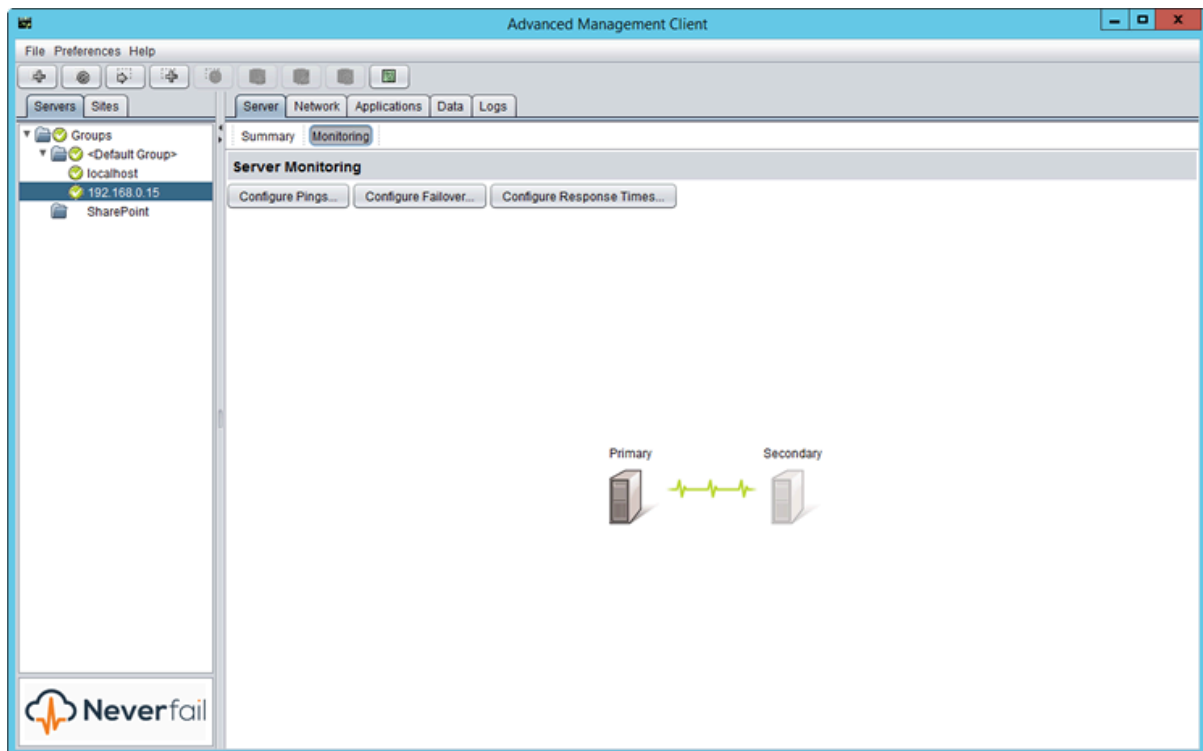
Related information

- [Monitoring the Status of Servers](#)
- [Configure Neverfail Continuity Engine Settings](#)
- [Forcing a Switchover](#)
- [Failover versus Switchover](#)
- [Split-brain Avoidance](#)

1. 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.

Figure 4-1. Server Monitoring page



2. Configure Neverfail Continuity Engine Settings

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

2.1. Configure Pings


About this task

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.

Procedure


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

Figure 4-2. Server Monitoring: Ping Configuration: Ping Settings Tab



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

Figure 4-3. Server Monitoring: Ping Configuration: Ping Routing Ta



2.2. Configure Failover

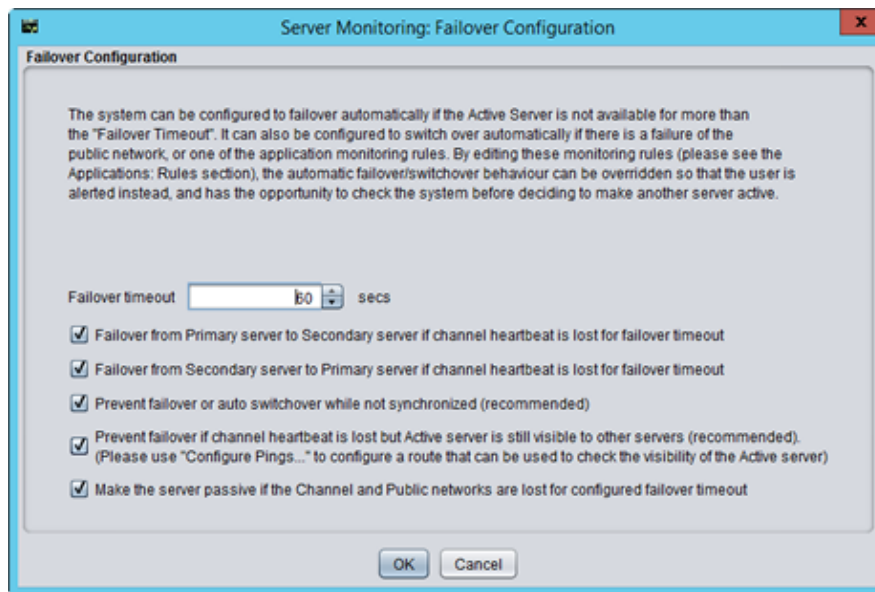
About this task

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.

Procedure

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

Figure 4-4. Server Monitoring: Failover Configuration



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

Note

For more information about configuring options for failover, see **Split-brain Avoidance**.

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

Note

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

2.3. Configure Response Times


About this task

Neverfail Engine also allows you to configure channel connection timeouts.

Procedure

1. Click **Configure Response Times** to open the *Server Monitoring: Response Times* dialog. The following options are available:
 - Time to wait following channel connection before starting replication
 - Time to wait following channel disconnection before stopping replication

Figure 4-5. Server Monitoring: Response Times



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

2.4. Common Administrative Tasks in Neverfail Continuity Engine

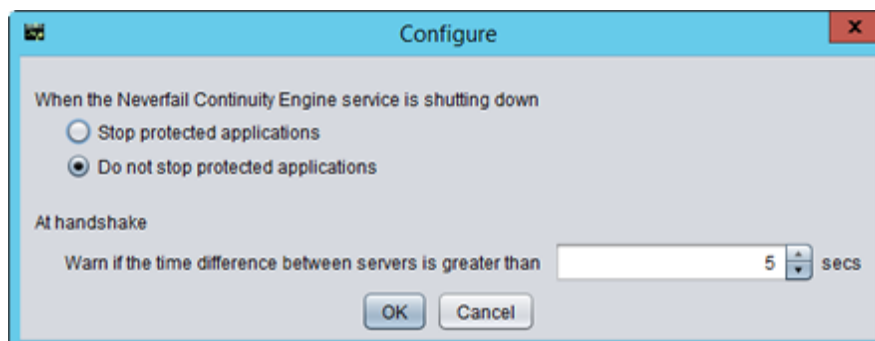
About this task

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

Procedure

1. Click to open the *Configure* dialog.

Figure 4-6. Configure (Shutdown)



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.

3. 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 rebooting the server.

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

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

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.

4. Failover versus Switchover

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.

4.1. Configuring Failover and Active Server Isolation

About this task

Neverfail Continuity 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**.

Procedure

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

Note

You must configure Management IP addresses on the 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.

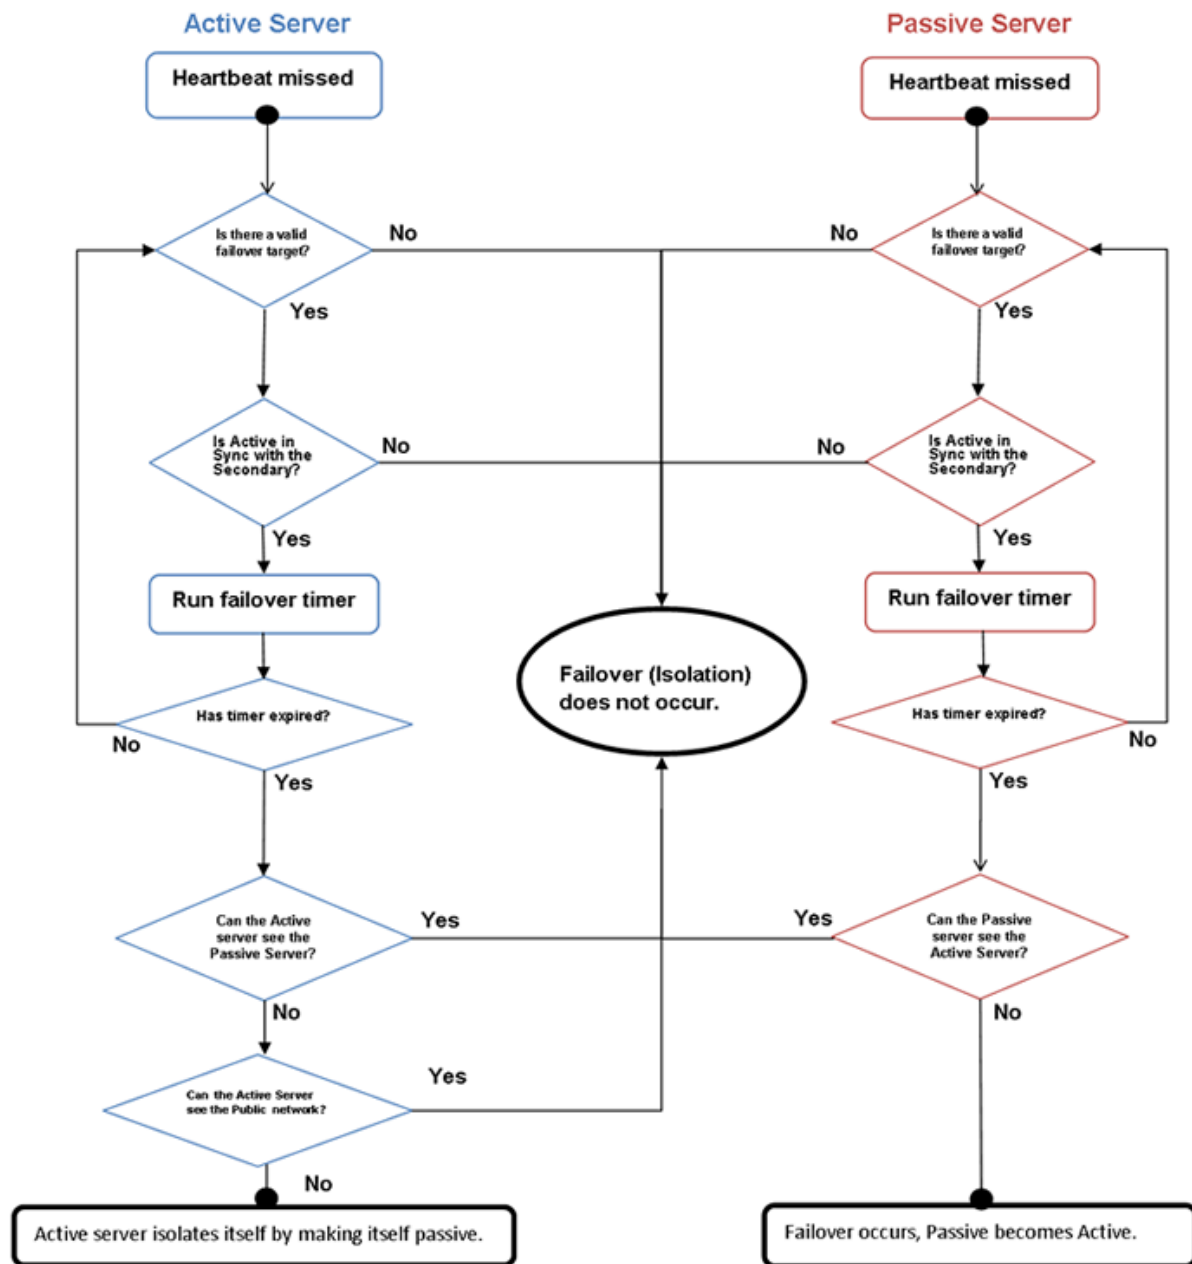
Active Server Isolation

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

Figure 4-7. Network Isolation Workflow Diagram



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.

4.2. Recover From a Failover

About this task

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.

Procedure

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:
 - a. 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.
 - b. Run the *Configure Server Wizard* and confirm that the server is set to Secondary and is active. Click **Finish** to accept the changes.

Note

If Neverfail 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*.

- c. 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.

5. Split-brain Avoidance

About this task

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.

Procedure

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

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

Note

You must configure Management IP addresses on the 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.

What to do next

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

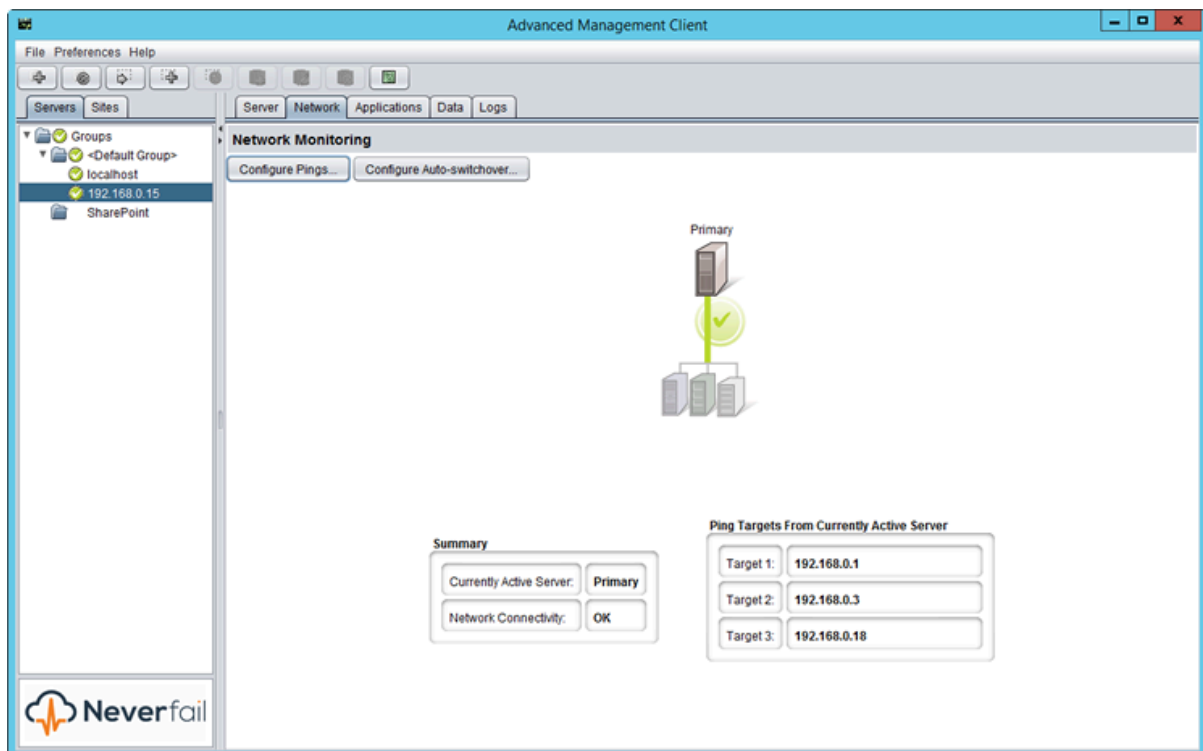
network reconnection occurs). The Management IP address can be added using the Configure Server Wizard.

Chapter 5. Network Protection

Neverfail Continuity 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.

Figure 5-1. Network Monitoring



Related information

- [Configure Public Network Monitoring](#)
- [Enabling Automatic Switchover in a WAN](#)
- [Setting Max Server Time Difference](#)

1. Configure Public Network Monitoring

About this task

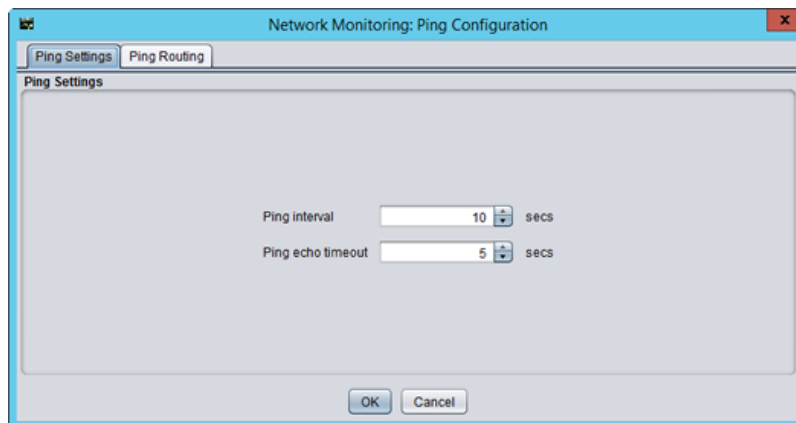
The Public network monitoring feature, previously discussed, is enabled by default during the installation of Neverfail Engine. This feature integrates the polling of the particular waypoints around the network through the active server's Public connection to ensure connectivity with the Public network 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:

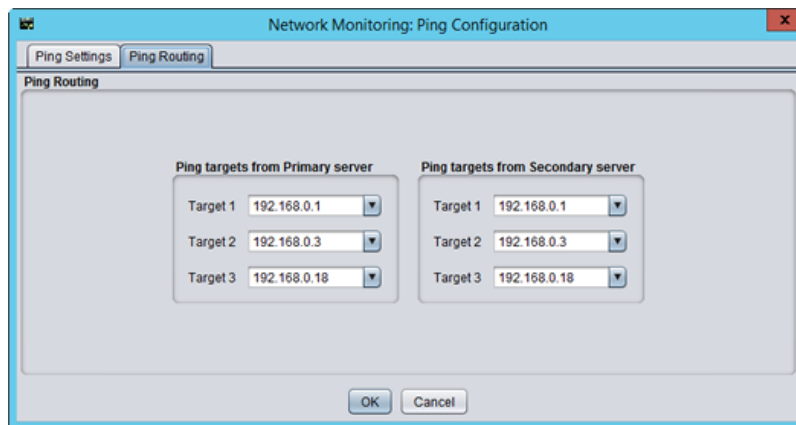
Procedure

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

Figure 5-2. Network Monitoring: Ping Configuration: Ping Settings



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

Figure 5-3. Network Monitoring: Ping Configuration: Ping Routing

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.

2. Enabling Automatic Switchover in a WAN

About this task


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

Procedure

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

Figure 5-4. WAN Auto-Switchover Configuration



3. Setting Max Server Time Difference

About this task

Neverfail Continuity 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.

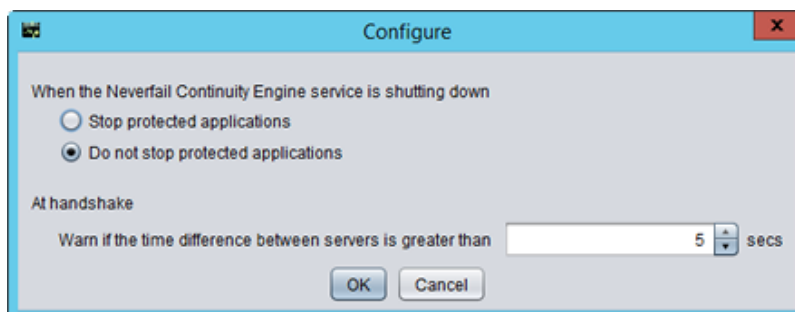
To set Max Server Time Difference:

Procedure

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

Figure 5-5. Server: Summary Configure dialog



Chapter 6. Application Protection

Related information

- [Applications Environment](#)
- [Applications: Summary](#)
- [Applications: Services](#)
- [Applications: Tasks](#)

1. 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.

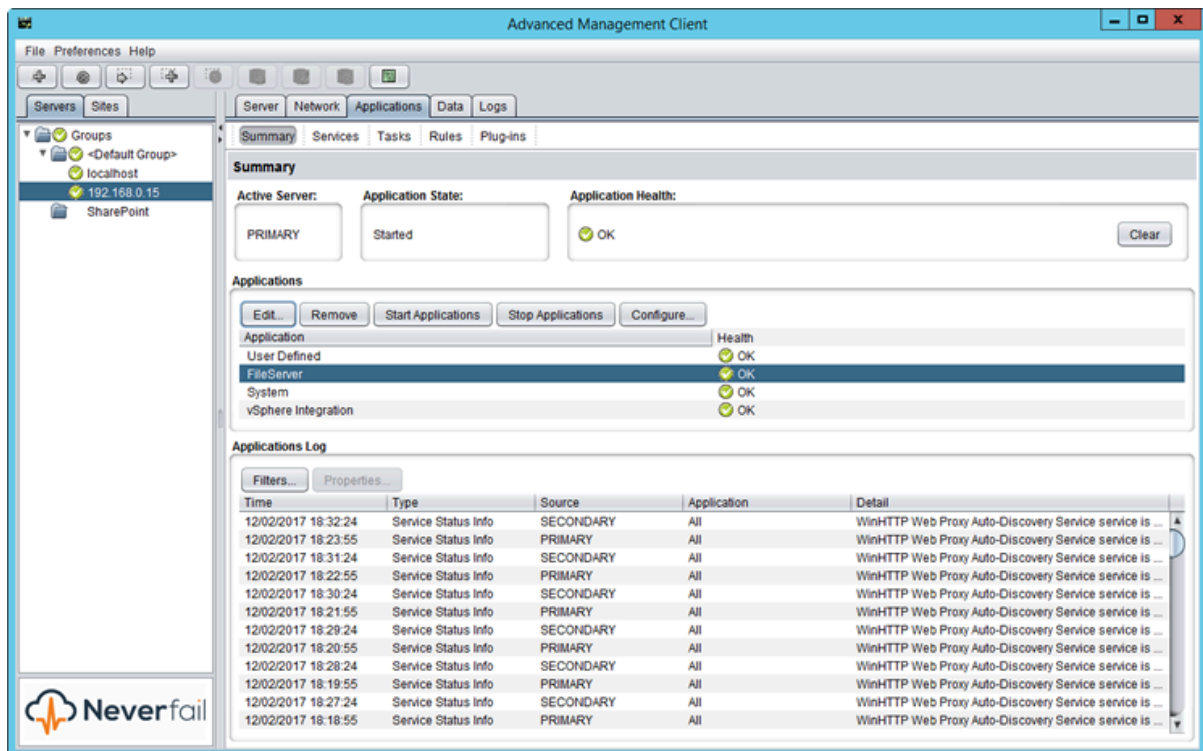
The AMFx also employs sponsorship for protected applications' files and services. With sponsorship, multiple plug-ins can share files or services. When removing a plug-in, sponsorship prevents removal of a shared file or service that is still required by a remaining plug-in.

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.

2. 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.

Figure 6-1. Applications: Summary



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

View Application Status

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


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.

Figure 6-2. Edit Application



Note

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

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

Remove an Application

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

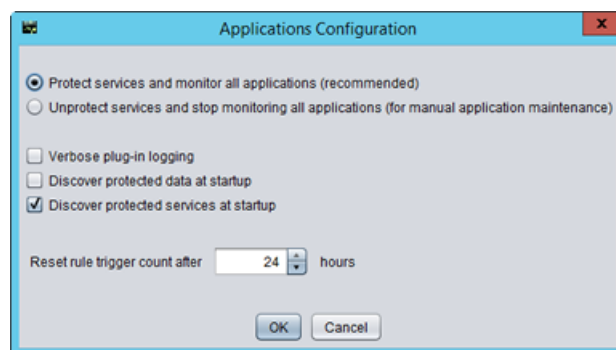
Configure Applications

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.

Figure 6-3. Applications Configuration



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

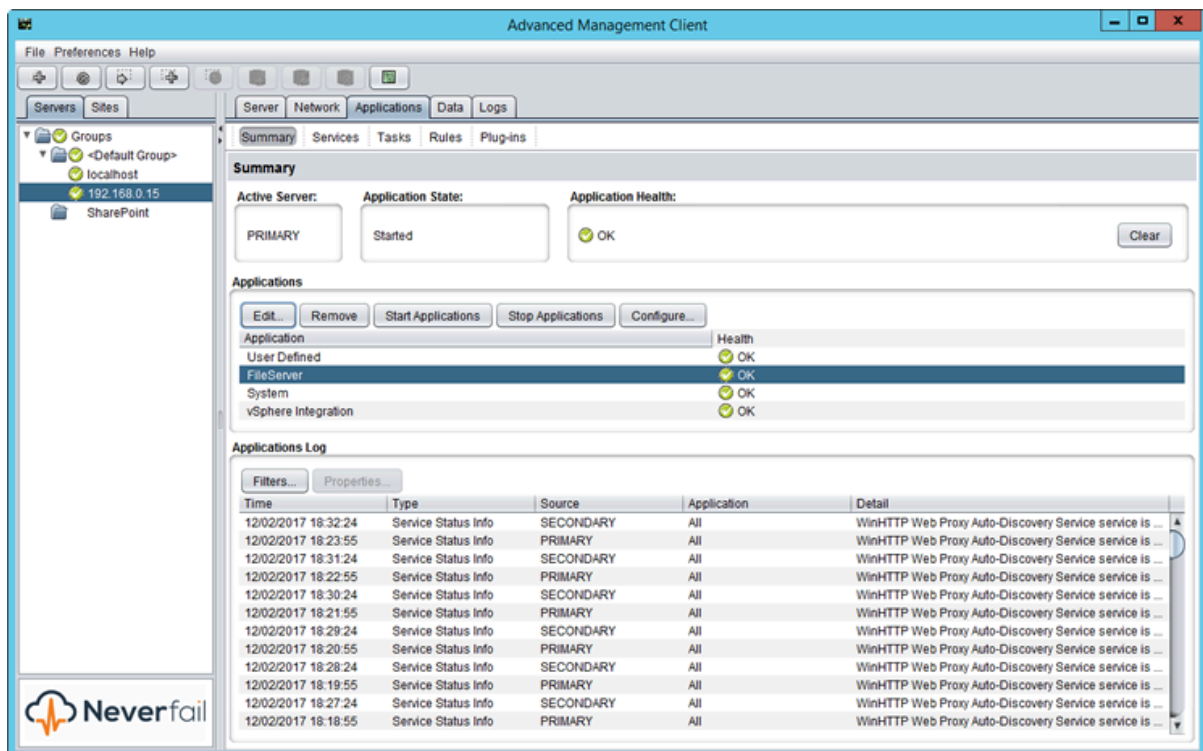
Optionally select any or all of the following:

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

View the Applications Log

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

Figure 6-4. Applications Log



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

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

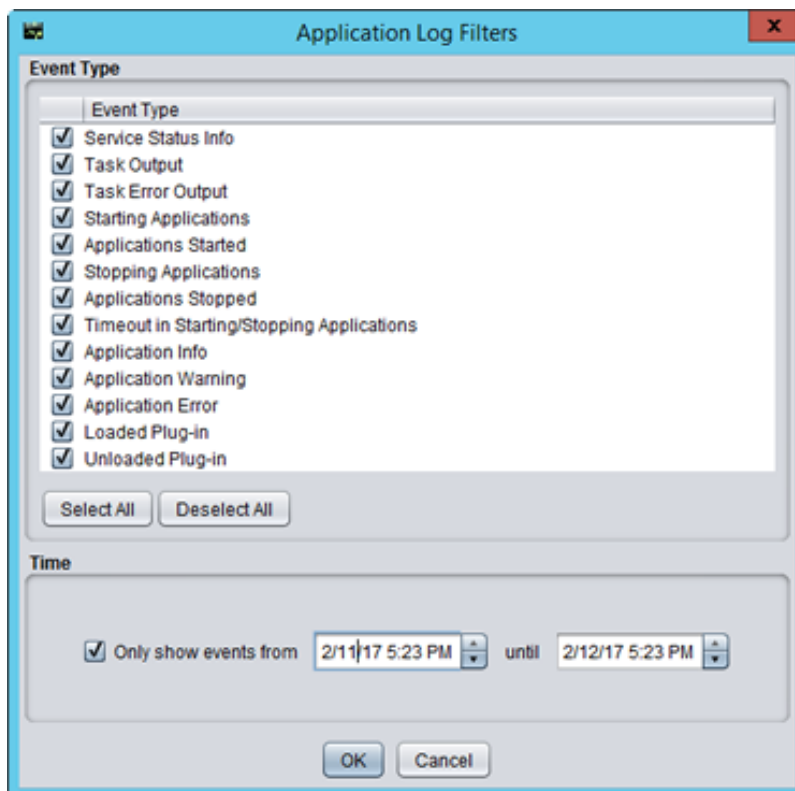
Filter Application Log Entries

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

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

The *Application Log Filters* dialog appears.

Figure 6-5. Application Log Filters

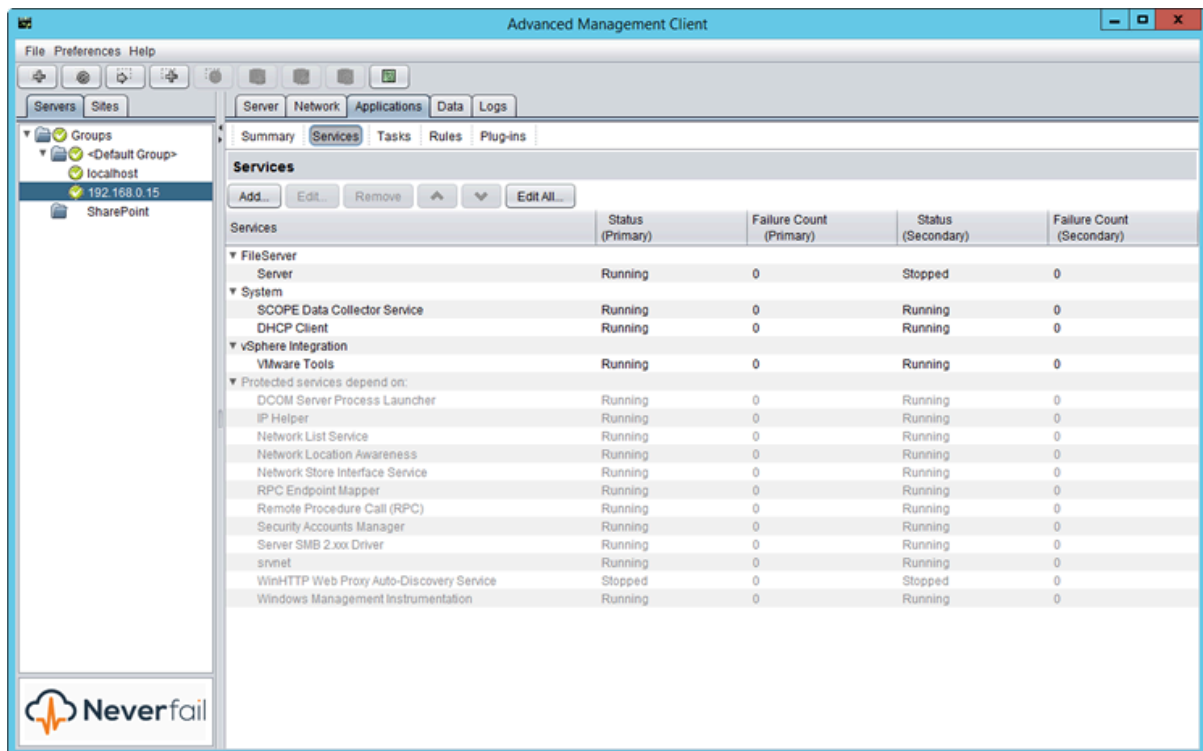


Use the check boxes (select to display or clear to hide) to filter *Application Log* entries by at least one *Event Type*. To display only entries within a particular time range, select the check box associated with *Only show events from* and type values into the two date/time text boxes or use the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

3. 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.

Figure 6-6. Application: Services

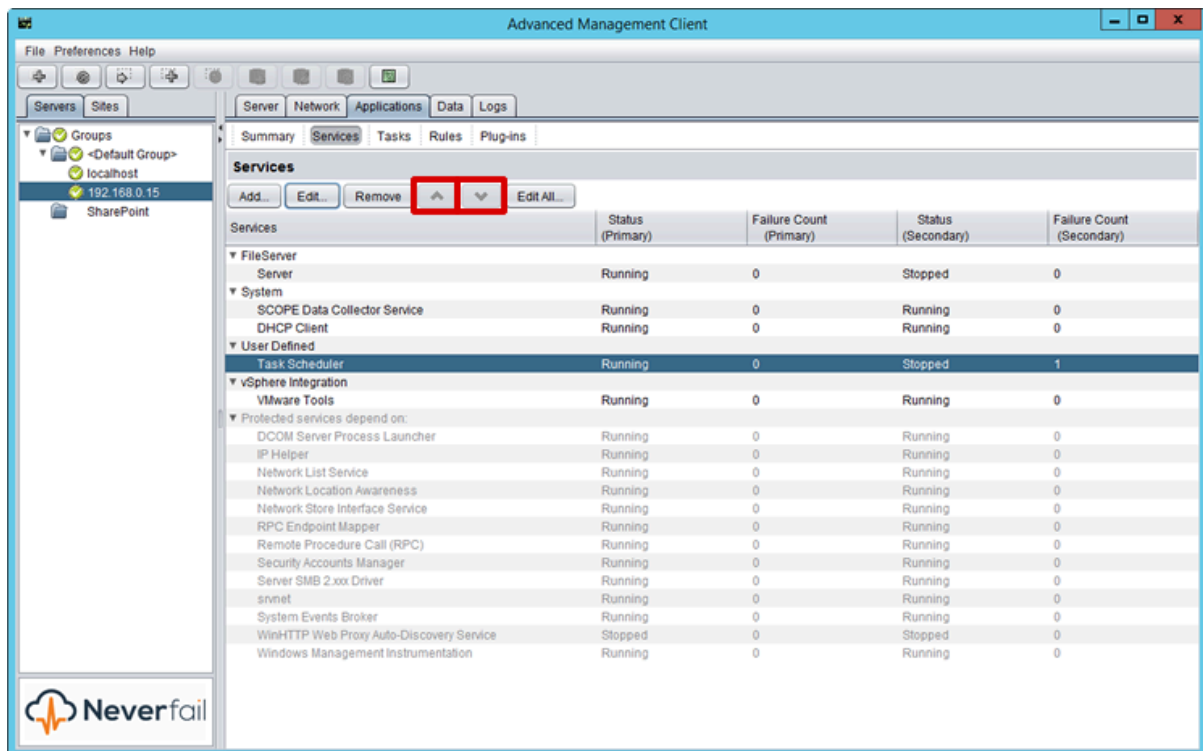


Change the Order of Services

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.

Figure 6-7. Change the Order of Services



4. Applications: Tasks

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

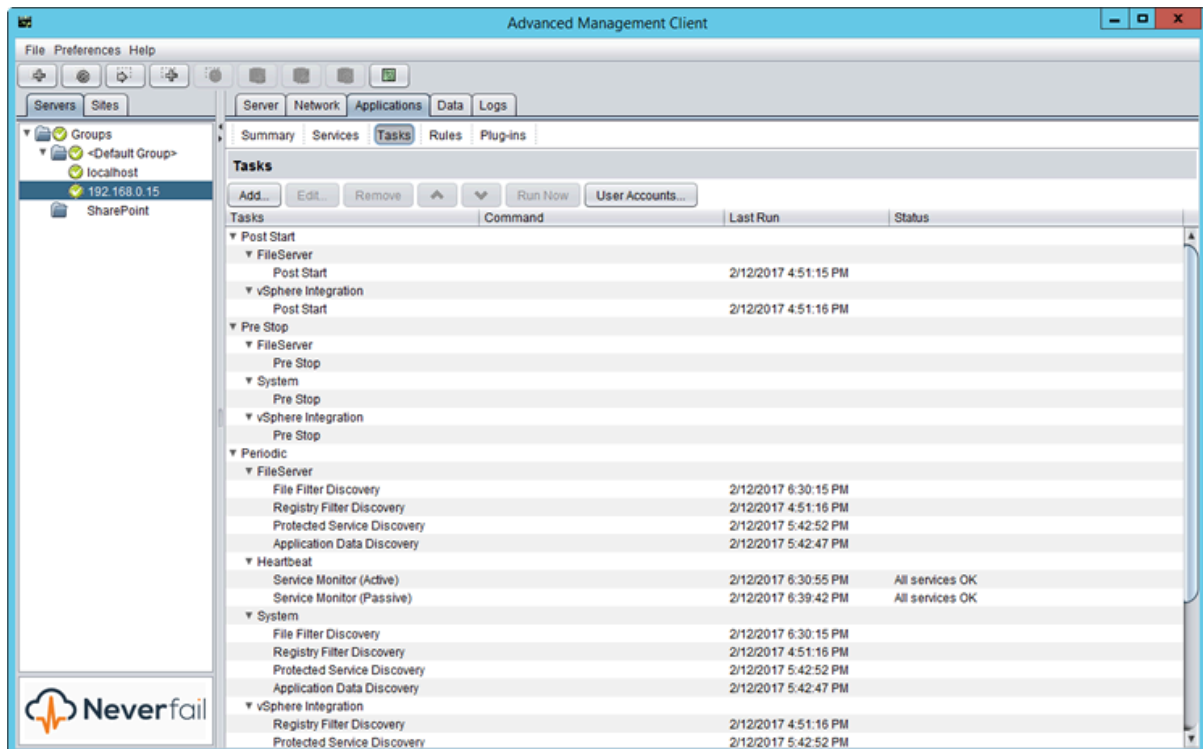
Task types are determined by when the tasks are run, and include the following:

- **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.

Figure 6-8. Applications: Tasks page



Change the Order of Tasks

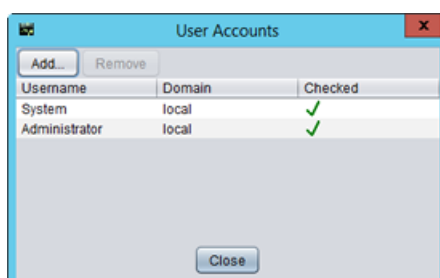
You can change the order of tasks using **Up** and **Down** arrows (near the top of the page or on the right-click menu) to change the order in which they appear in the list of tasks.

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.

Figure 6-9. User Accounts




- To add a user account:

1. Click **Add**.

The *Add User* dialog appears.

Figure 6-10. Add User Account



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

Note

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

- 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.

Chapter 7. Data Protection

Related information

- [Data: Replication](#)

1. Data: Replication

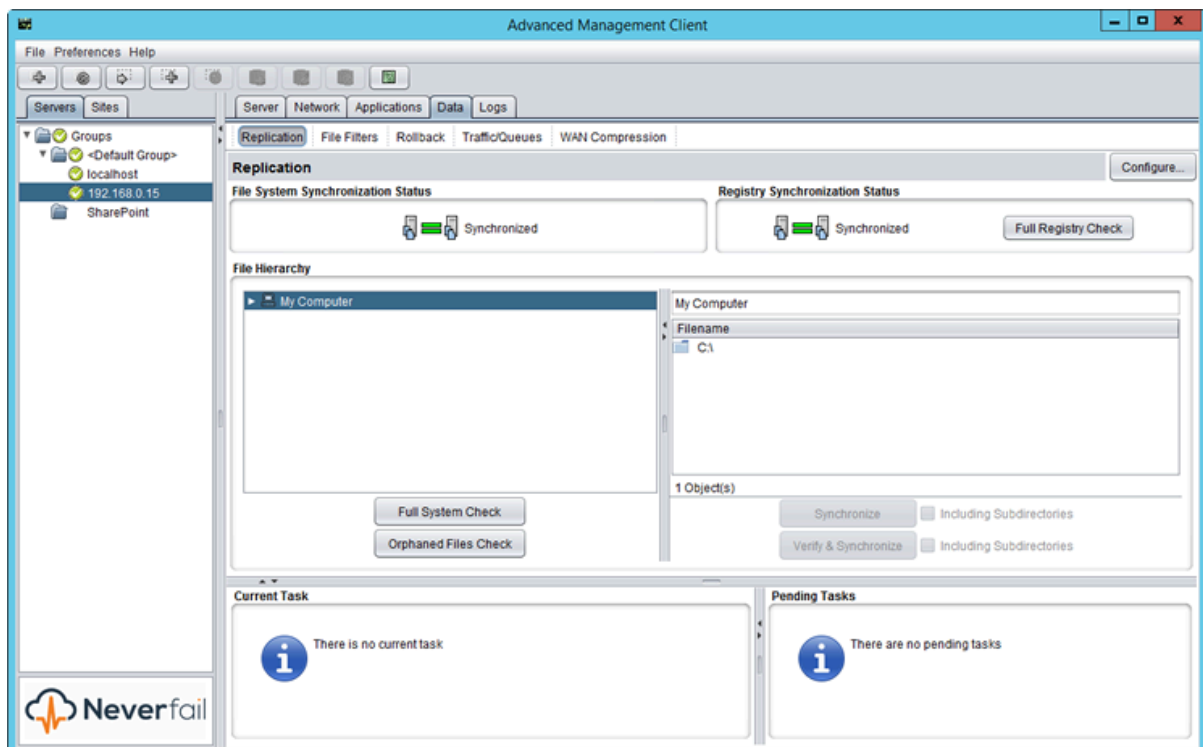
Neverfail Continuity Engine can protect many permutations or combinations of file structures on the active server by the use of custom inclusion and exclusion filters configured by the administrator.

Note

The Neverfail Continuity 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.

Figure 7-1. Data: Replication page



1.1. Initiate a Full System Check

About this task

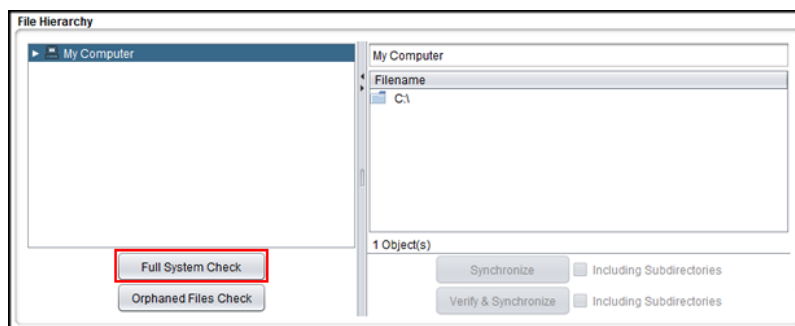
Certain system events, such as preceding a switchover or following a failover or split-brain syndrome, may require running a full system check to ensure that the entire protected file set is synchronized and verified. A full system check performs a block-level check identical to that performed during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

To initiate a full system check:

Procedure

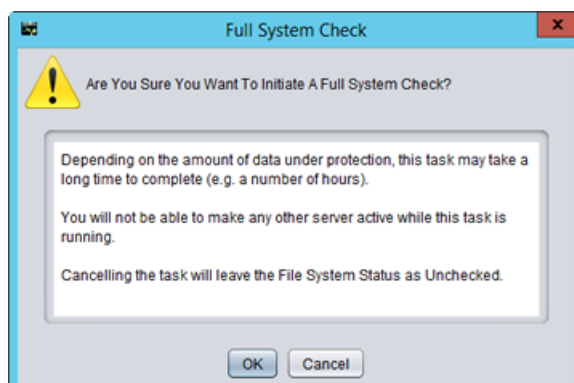
1. Click **Full System Check** in the left pane of the *File Hierarchy* pane.

Figure 7-2. Data: Replication File Hierarchy pane



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

Figure 7-3. Full System Check Caution Message



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

Note

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

1.2. Fast Check

About this task

The Fast Check process is used by Neverfail Continuity 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 initiate. Additionally, you can configure the length of time to wait for Fast Check to complete before starting applications.

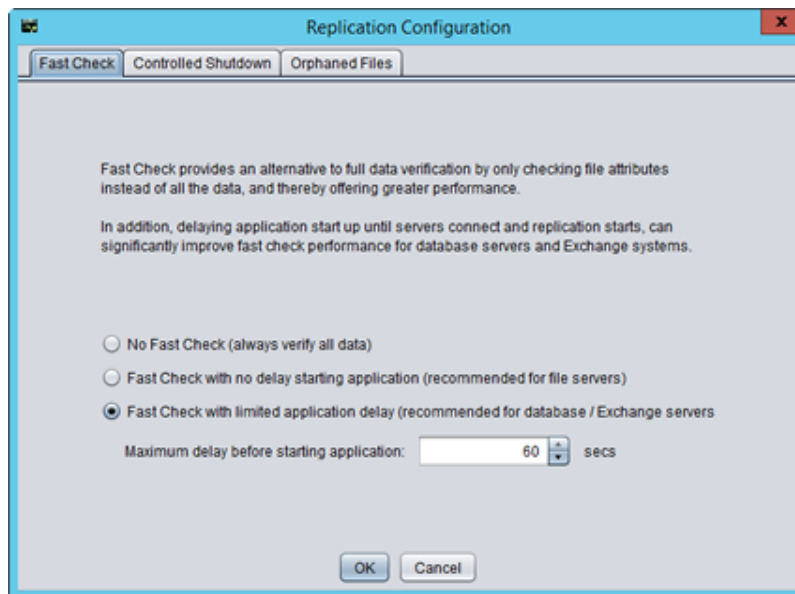
Fast Check is beneficial after a graceful shutdown where servers were synchronized before shutdown. Fast Check allows the server to check the file synchronization rapidly and start to service clients. If Fast Check detects files that are out-of-sync, it initiates the full verify and synchronization process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

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

Procedure

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**.

Figure 7-4. Configure Fast Check



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**.

1.3. Manually Initiate File Synchronization

About this task

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.

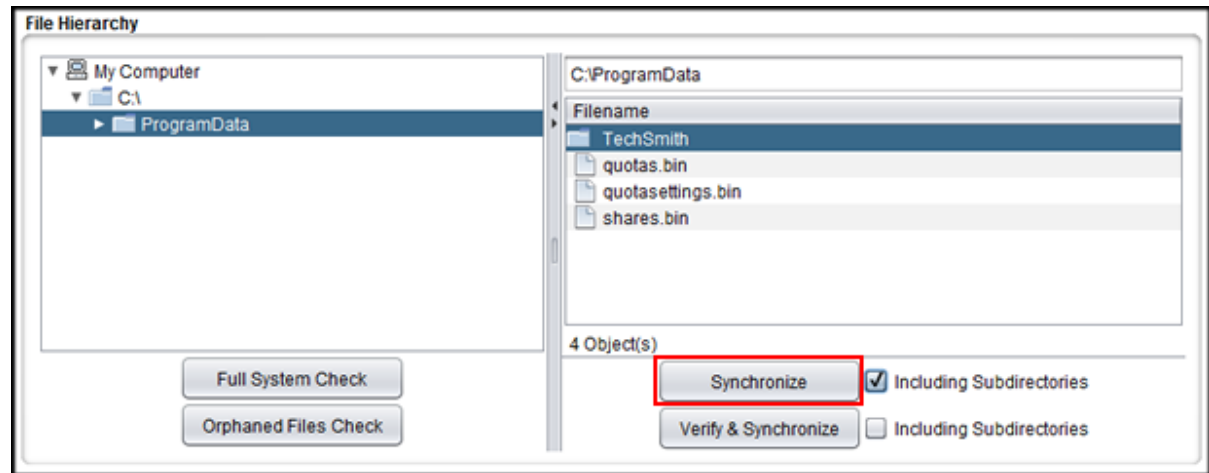
To manually re-synchronize:

Procedure

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

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

Figure 7-5. Manual Selection to initiate file synchronization



1.4. Manually Initiate Verify and Synchronize

About this task

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:

Procedure

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

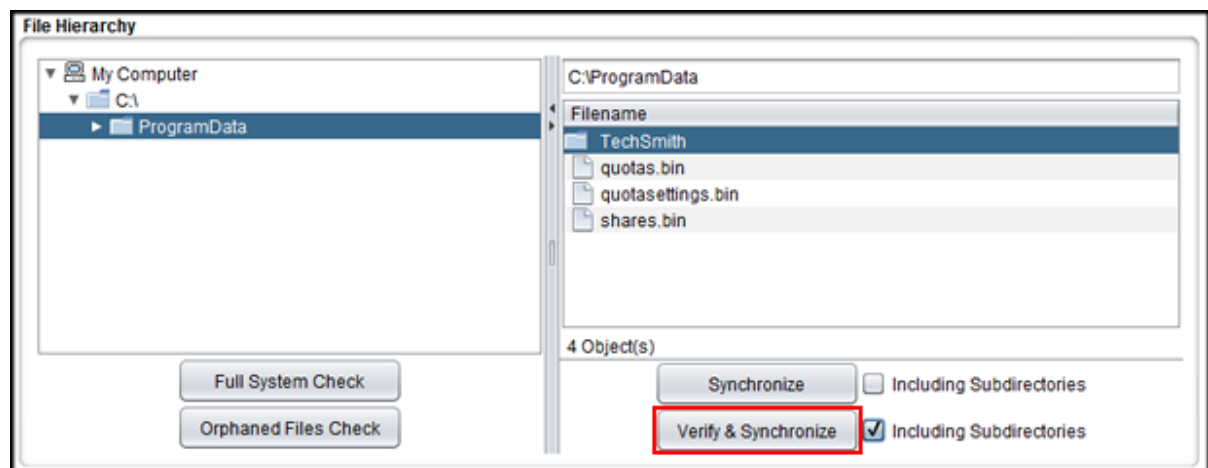
You also can right-click on a folder in the tree view (in the left pane of the *File Hierarchy* pane) to quickly select **Verify and Synchronize** from a menu. This option automatically includes subdirectories.

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

Note

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

Figure 7-6. Manual Selection to Initiate Verify and Synchronize



1.5. Orphaned Files Management

Neverfail Continuity 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

Filters for files, file types, or other wildcards.

Folder root filters

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

Exclusion file filters

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

Filters for files, file types, or other wildcards

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

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

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

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

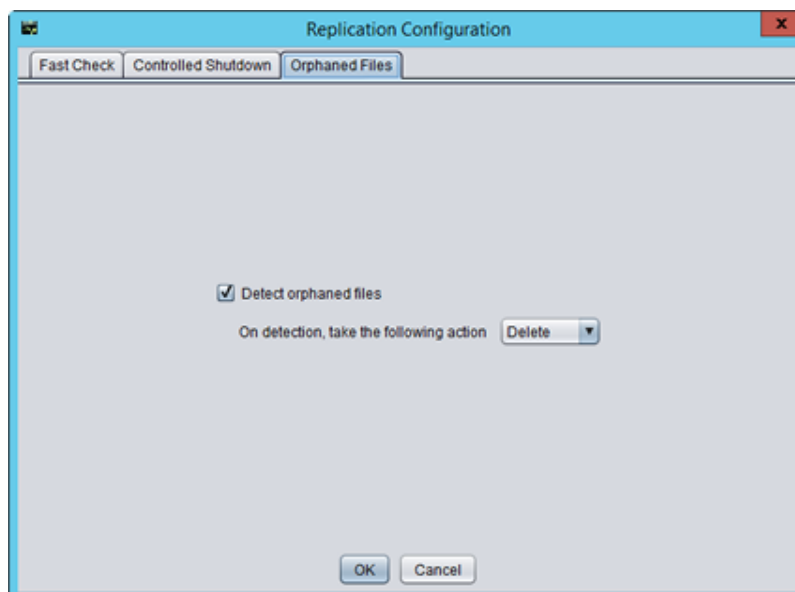
Configure Orphaned Files Check

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

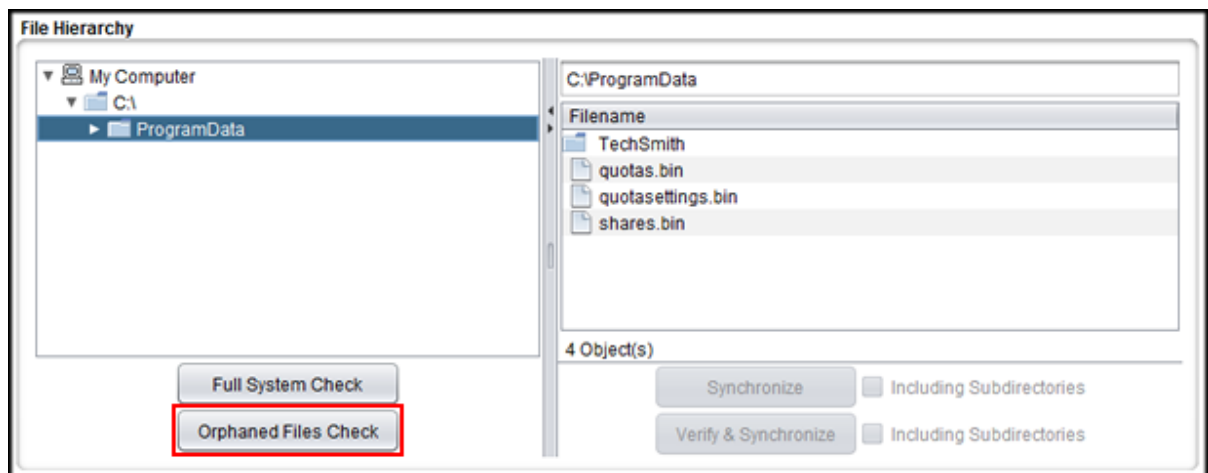
To Configure Orphaned Files Check options:

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

Figure 7-7. Orphaned Files Configuration Options



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

Figure 7-8. Initiate Orphaned Files Check

Part III. Reference

Chapter 8. Other Administrative Tasks

Related information

- [Post Installation Configuration](#)
- [Business Application Groups](#)
- [Configure Event Log Files](#)
- [Review Event Logs](#)
- [Recloning Secondary or Tertiary Server](#)
- [Deploying Neverfail Engine 8.5 Cluster in Amazon Web Services Cloud Environment](#)
- [Deploying a Passive Node in an Amazon Web Services Cloud Environment](#)

1. Post Installation Configuration

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

1.1. Configure the VmAdapter Plug-in

About this task

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

Procedure

1. Launch the Engine Management Service UI for the server pair and login.
2. Navigate to **Settings > Application Protection > Plug-ins**.
3. Select the `VmAdapterNFPlugin.dll`
4. Click the **Edit** button.

The *Edit Plug-in* dialog is displayed.

5. 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
6. 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
7. If integration with vSphere HA monitoring is desired, select the **Integrate with vSphere HA monitoring** check box.

This option requires vSphere HA Application monitoring for the cluster and VM.
 8. Click **OK**.

1.2. Adding an Additional Network Interface Card

About this task

Neverfail Continuity Engine allows for installation using a single NIC on each Neverfail 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 Neverfail 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 Neverfail Engine is installed as a V2V Pair with the Primary server active and the Secondary server passive.

Procedure

1. Shutdown Neverfail 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 Neverfail Engine.
5. Allow the server to synchronize.

2. Business Application Groups

Neverfail Continuity 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 Continuity 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.

2.1. Installing the Business Application Plug-in

Before you begin

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.

About this task

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

Procedure

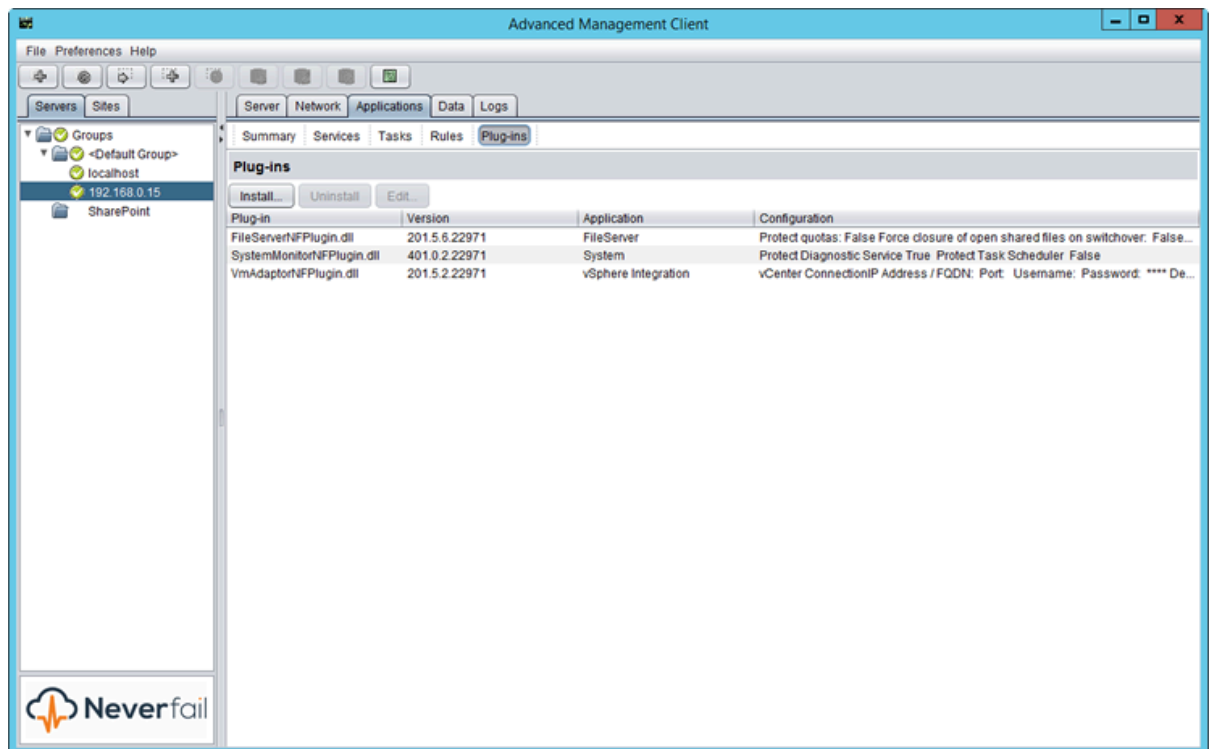
1. Download the `Z-SW-BusinessApplicationPlugin.201.5.[n].zip` file to a temporary location on the active server in the cluster.

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.

Figure 8-1. Applications: Plug-ins Page



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

Figure 8-2. Install Plug-in Dialog



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

Important

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

2.2. Creating a Business Application Group

Before you begin

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

About this task

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.

Procedure

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

The *Business Application Group Wizard* is displayed.

Figure 8-3. Business Application Group Wizard



3. Review the information in the *Create Business Application Group Wizard* and click **Next**.

The *Enter Basic Group Information* page is displayed.

Figure 8-4. Enter Basic Group Information Page

The screenshot shows a Windows-style dialog box titled "Create Business Application Group Wizard". The main area is titled "Enter Basic Group Information". It contains three text input fields: "Business Application Group Name" with the text "SharePoint Gp", "Primary Site Name" with the text "Austin", and "Secondary Site Name" with the text "Las Vegas". At the bottom right, there are three buttons: "<< Back", "Next >>", and "Cancel".

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

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

5. Add the name of the Primary Site.
6. Add the name of the Secondary (DR) site and click **Next**.

The *Add Servers to Business Application Group* page is displayed. A list of available servers is displayed in the left pane of the dialog.

Figure 8-5. Add Servers to Business Application Group Page

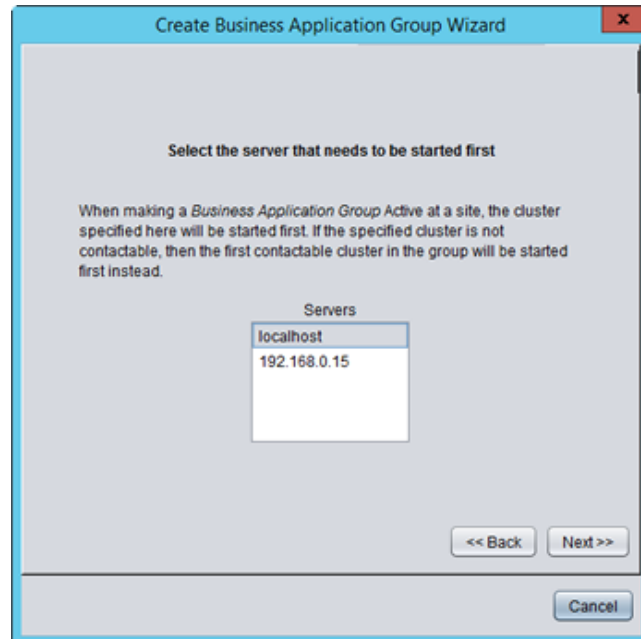
The screenshot shows the same "Create Business Application Group Wizard" dialog box, but on the "Add servers to the Business Application Group" page. It features two list boxes: "Available Servers" on the left containing "localhost" and "192.168.0.15", and "Added Servers" on the right which is currently empty. Between the list boxes are four buttons: "<<", "<", ">", and ">>". At the bottom right, there are three buttons: "<< Back", "Next >>", and "Cancel".

7. Select the servers to join the Business Application Group and click the > button to add the servers to the Business Application Group. Click **Next**.

The *Select First Server to Switch* page is displayed.

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

Figure 8-6. *Select First Server to Switch Page*



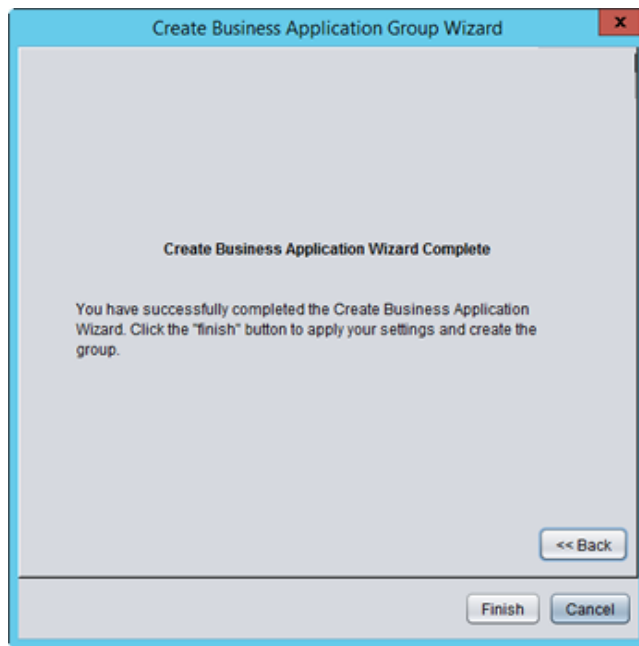
Note

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**.

Figure 8-7. Create Business Application Wizard Complete Page

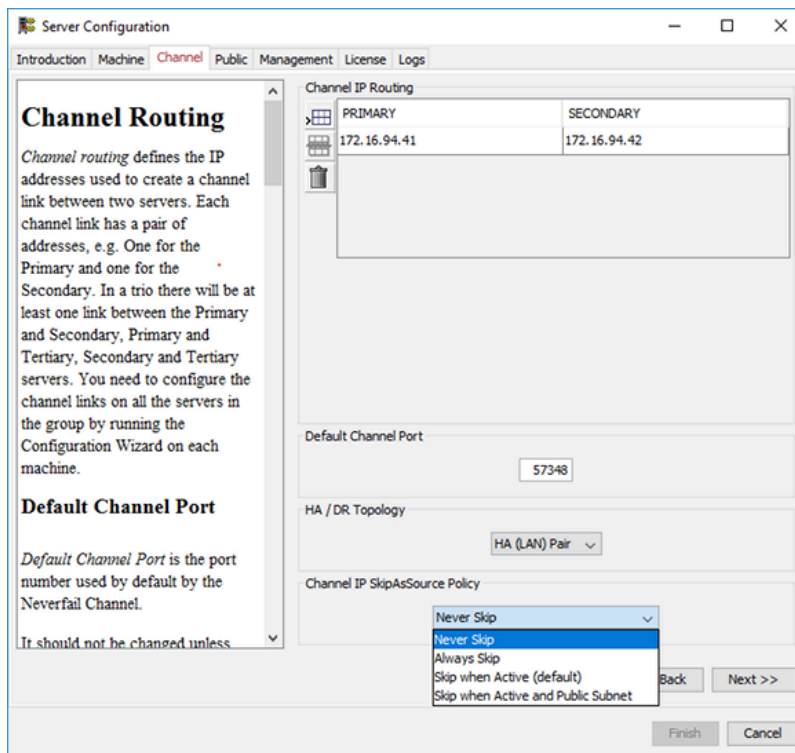


2.2.1. Configuring Neverfail Engine for Business Application Group

Procedure

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:

Figure 8-8. Channel IP SkipAsSource options



- 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).

2.3. Editing a Business Application Group

About this task

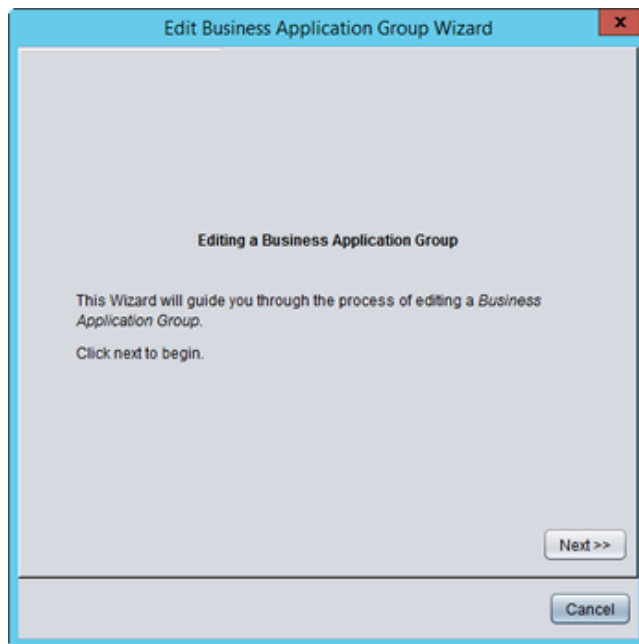
Engine Management Service allows you to edit the configuration of an existing Business Application Group.

Procedure

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

The *Edit Business Application Group Wizard* is displayed.

Figure 8-9. Edit Business Application Group Wizard



2. Click **Next**.

The *Enter Basic Group Information* page is displayed.

Figure 8-10. Enter Basic Group Information

The screenshot shows a Windows-style dialog box titled "Edit Business Application Group Wizard". The main area is titled "Enter Basic Group Information". It contains three text input fields: "Business Application Group Name" with the text "SharePoint Gp", "Primary Site Name" with the text "Austin", and "Secondary Site Name" with the text "Las Vegas". At the bottom right, there are two buttons: "<< Back" and "Next >>". At the bottom center, there is a "Cancel" button.

3. Edit the name of the Business Application Group, Primary Site, and/or the Secondary Site and click **Next**.

The *Select First Server to Switch* page is displayed.

Figure 8-11. Select First Server to Switch Page

The screenshot shows the same "Edit Business Application Group Wizard" dialog box, but at the "Select the server that needs to be started first" step. It includes explanatory 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 this is a list box titled "Servers" containing the entries "localhost" and "192.168.0.15", with "192.168.0.15" selected. The "<< Back", "Next >>", and "Cancel" buttons are also present at the bottom.

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**.

2.4. Dissolve a Business Application Group

About this task

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.

Procedure

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**.

Figure 8-12. Dissolve Business Application Group - Tool bar Button

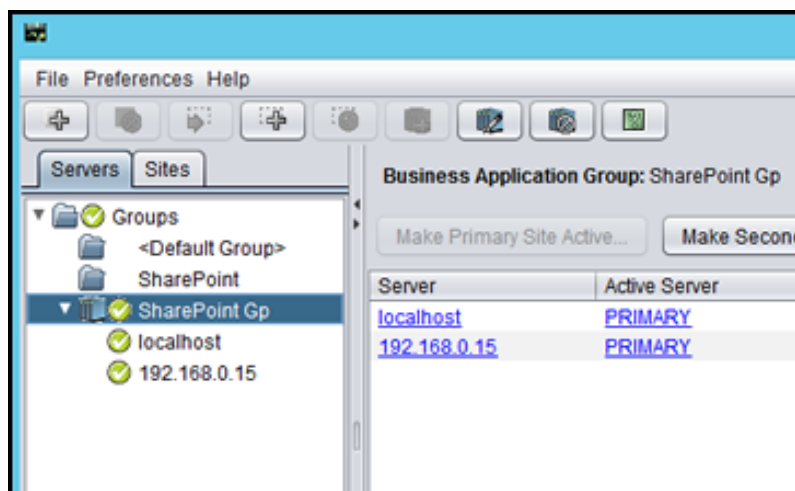
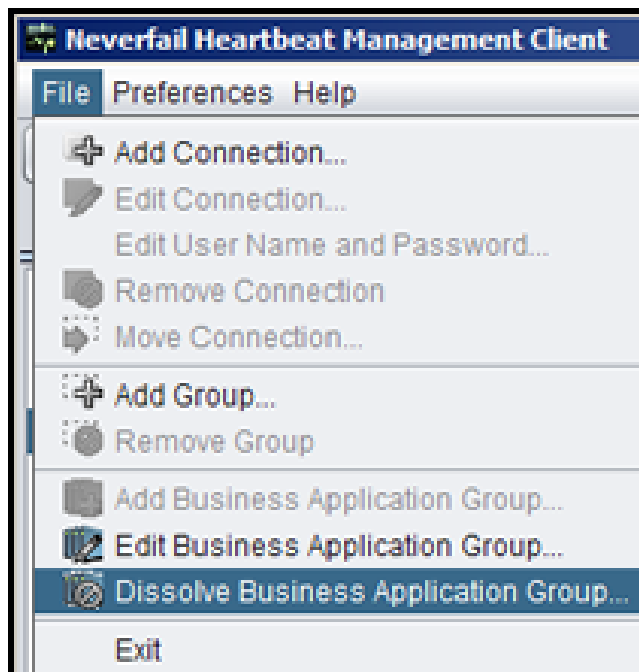


Figure 8-13. Dissolve Business Application Group - File Menu



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

Figure 8-14. Dissolve Business Application Group - Confirmation Dialog



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

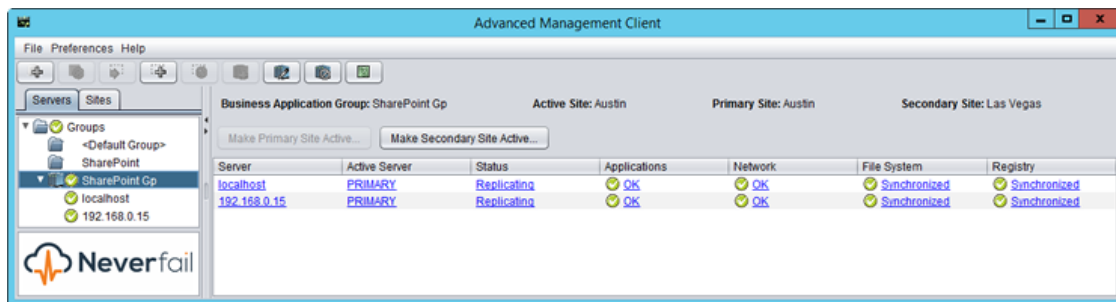
2.5. Business Application Switchover

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.

Figure 8-15. Business Application Group Summary Page



2.5.1. Performing a Business Application Switchover

Procedure

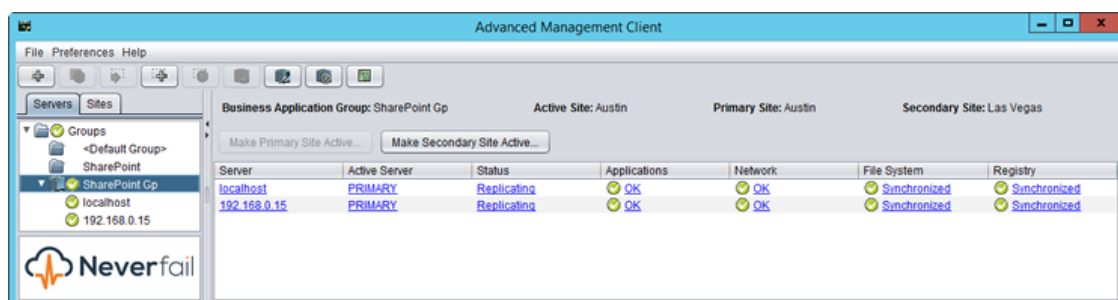
1. Launch the Neverfail Advanced Management Client.
2. Select the *Servers* tab in the left pane.
3. In the *Servers* pane, select the Business Application Group to switch.

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

Figure 8-16. Business Application Group Summary Page



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

2.6. 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.

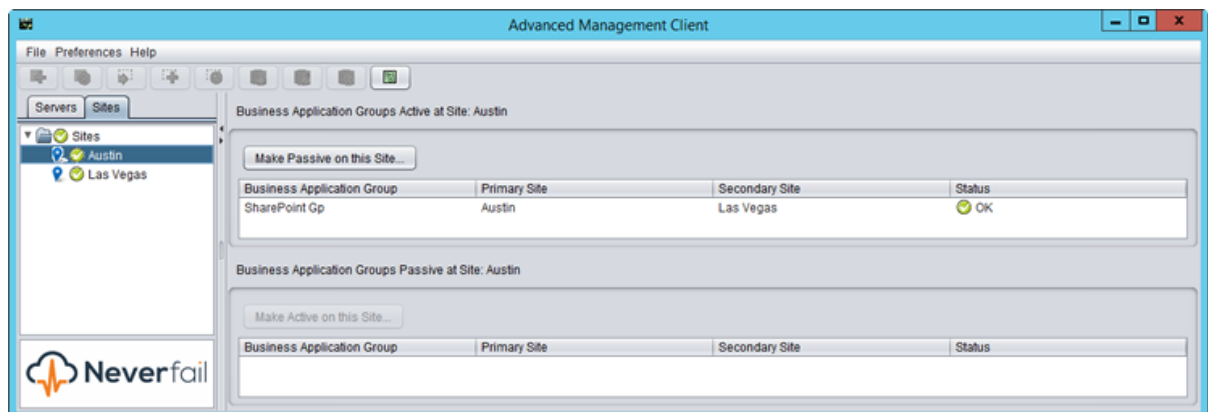
In the event of a WAN outage, the administrator needs to ensure that if the standby site is made active, then the administrator must shut down the previously active site to prevent both sites from being simultaneously active. To prevent both sites from being active at the same time, the administrator 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.

2.6.1. Performing a Site Switchover

Procedure

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

Figure 8-17. Neverfail Engine Sites Overview Page



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.

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

About this task

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

Procedure

1. Launch the Neverfail 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.

2.7. Uninstall the Business Application Plug-in

Before you begin

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.

About this task

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

Procedure

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.

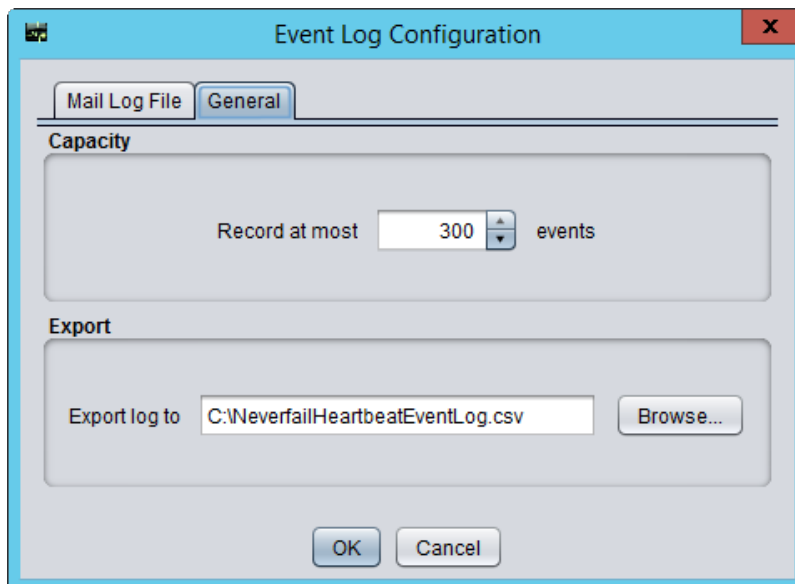
Note

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.

3. Configure Event Log Files

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

Figure 8-18. Event Log Configurations: General

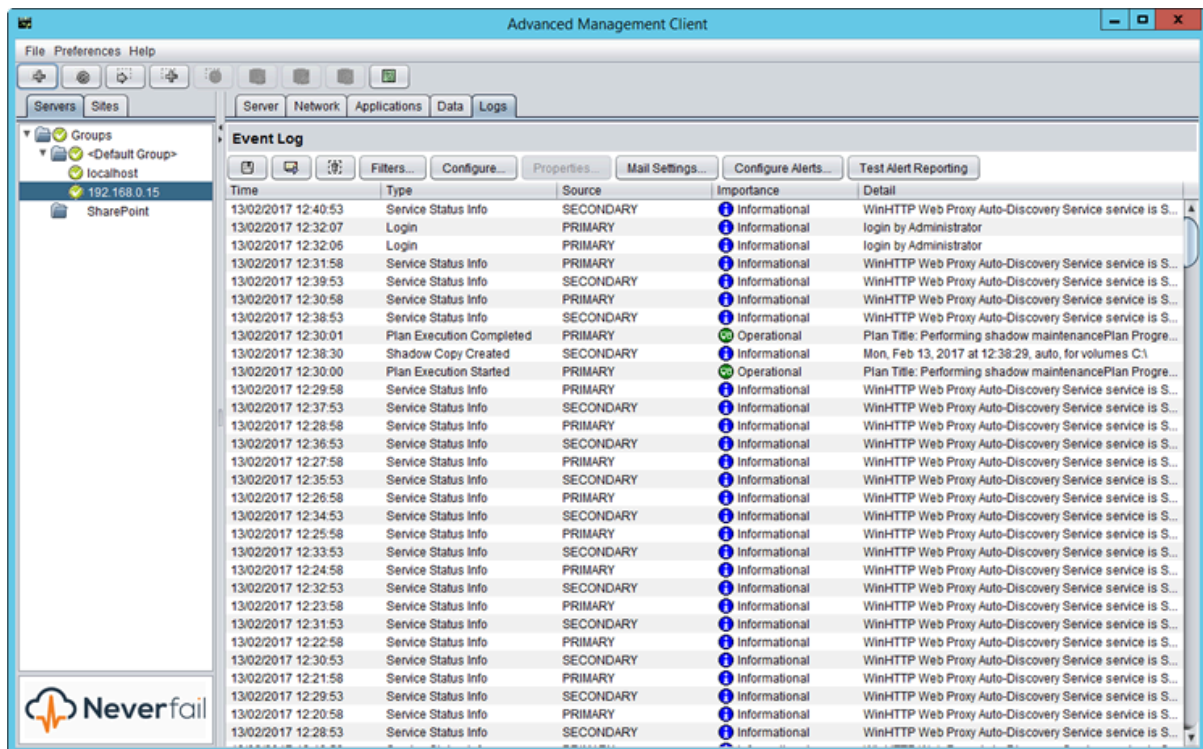


The length of the event list can also be adjusted using the *Record At Most* option. The default is to record 300 events but changing the value increases or decreases the length of the log list accordingly. After the logs are configured, click **OK** to commit the changes.

4. 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.

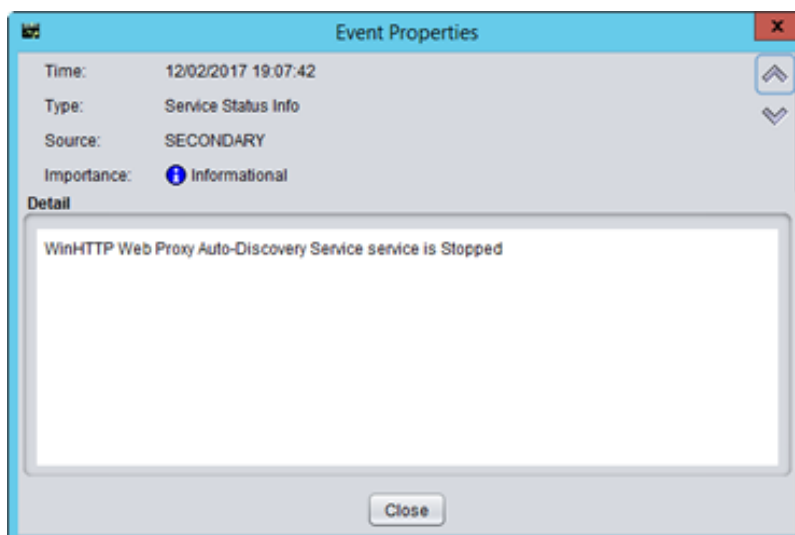
Figure 8-19. Event Log page



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





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

Figure 8-20. Event Log Properties



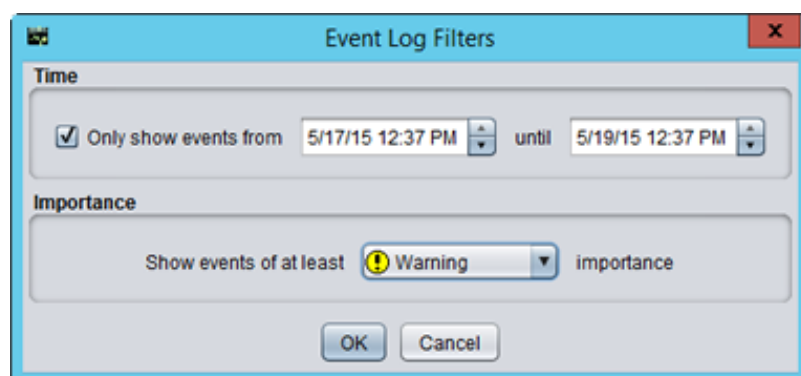
The *Event Properties* dialog gives the full detail and trace of the log that caused the event along with the source of the error aiding in troubleshooting. Further logs can be reviewed without having to close this window by using the **Up** and **Down** arrows of the dialog box to scroll through the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

There are four categories of importance of events that Neverfail 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.

Figure 8-21. Event Log Filters



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

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.

5. Recloning Secondary or Tertiary Server

The *Reclone Secondary or Tertiary Server* feature allows the administrator to perform a server reclone, either in place or scheduled. The **Secondary node, Tertiary node or both Secondary and Tertiary nodes** can be redeployed using this feature, while the **Primary node needs to be active** to serve as source for the recloning operation. This feature is available in the **Management menu**, at the bottom of the **Deploy** section.

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 [Configure Connection to VMware vCenter Server](#) must be configured 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, Manual and Scheduled Automated recloning options, considering the above conditions are met.

The *Reclone Secondary or Tertiary Server* feature provides three cloning options in the *Select clone type* section:

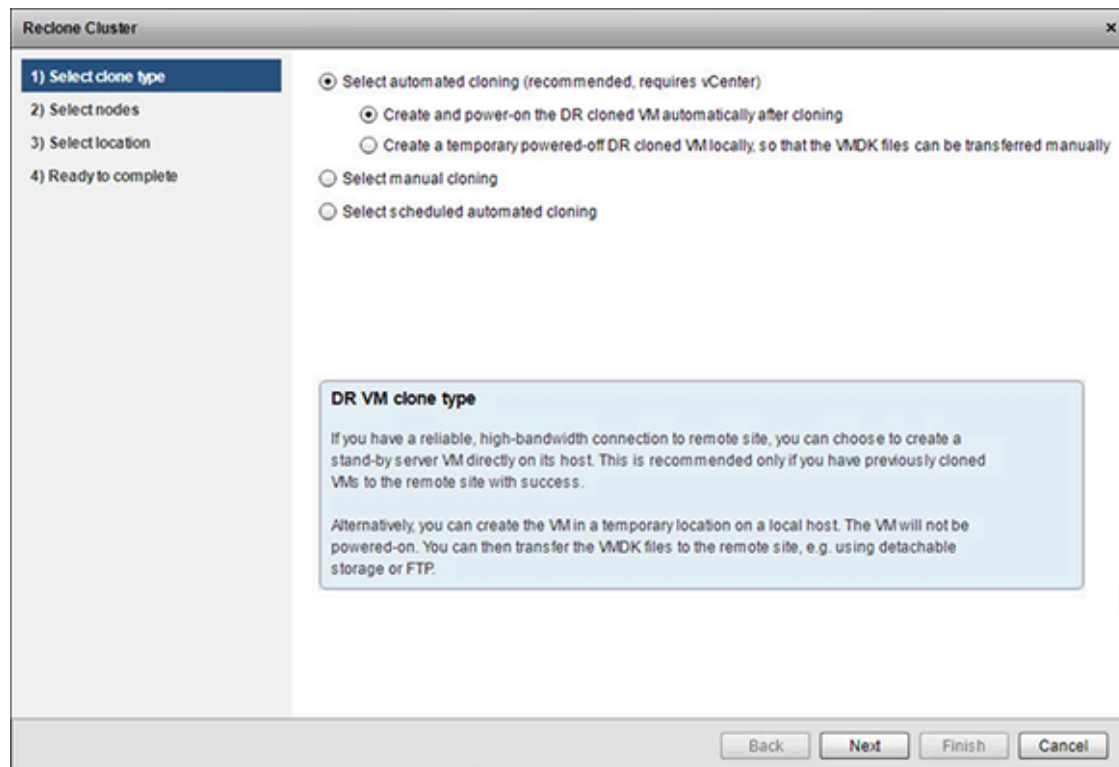
- Automated Recloning
- Manual Recloning
- Scheduled Recloning

5.1. Automated Recloning

The *Automated Recloning* task is completely handled by either vCenter Server or vCenter Converter.

This option is available when choosing the **Select automated cloning (recommended, requires vCenter)** option when configuring the clone type.

Figure 8-22. Select automated cloning



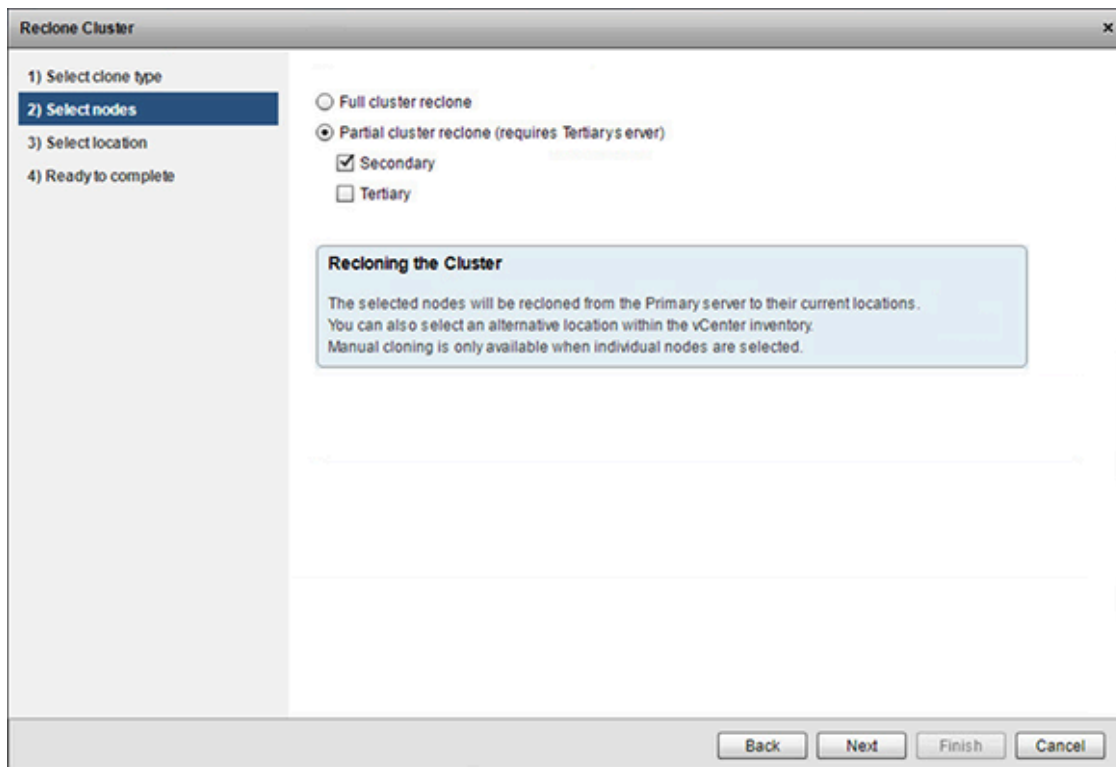
The cloning task is completely handled in an automated manner by either vCenter Server or vCenter Converter, providing that the vCenter Server connection, including the Default Host and vCenter Converter connection, if needed, have been properly configured.

The automated recloning task can handle the clones using two different approaches:

- **Create and power-on the DR cloned VM automatically after cloning.** This option is suitable when a high-bandwidth connection is available to the target clone host. It is recommended to make sure that you can successfully clone a VM to the remote host before using this option.
- **Create a temporary powered-off DR cloned VM locally, so that the VMDK files can be transferred manually.** This option is suitable when the cloned VM will be moved manually to the new host (for example using detachable storage or FTP to transfer the VMDK files).

Once the automated cloning option is chosen, the nodes to be cloned are available for selection in the *Select nodes* page.

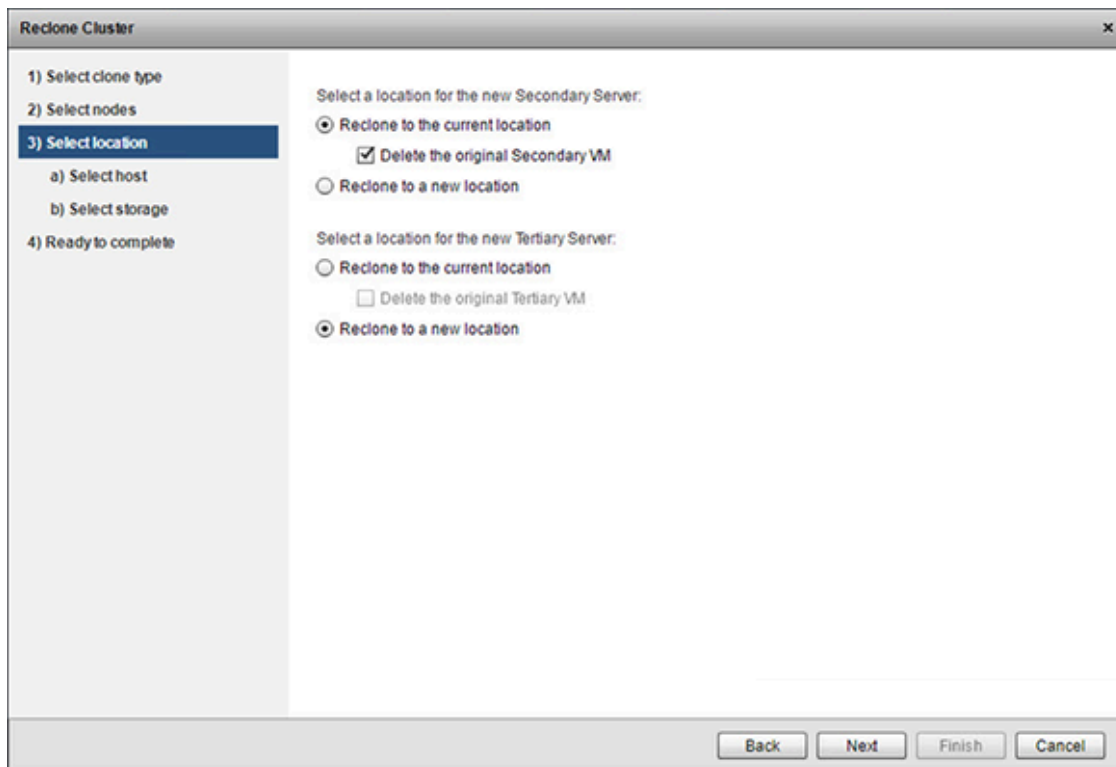
Figure 8-23. Select nodes



A **full cluster reclone** will clone the available passive nodes. If the cluster has a Tertiary node besides its Secondary, the **partial cluster reclone** option is available. This allows the recloning of only the Secondary or Tertiary nodes, or both, if selected so.

The selected nodes can be cloned to the same location as the original nodes or to a new host configured in the *Select location* page.

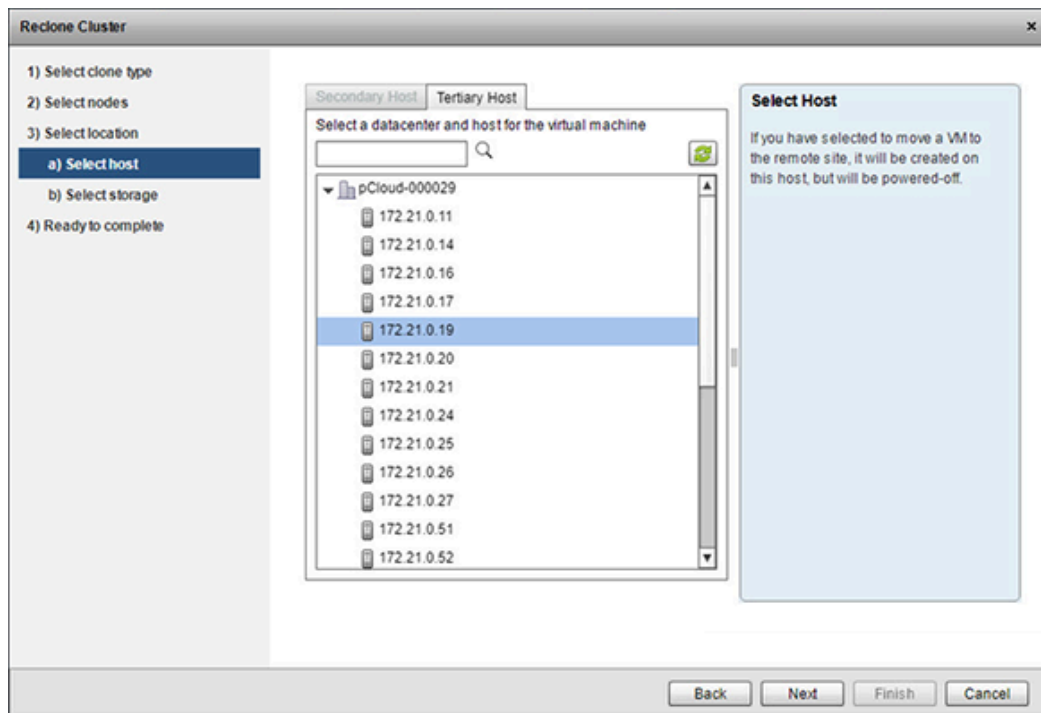
Figure 8-24. Select location



The location for each selected node can be configured as follows:

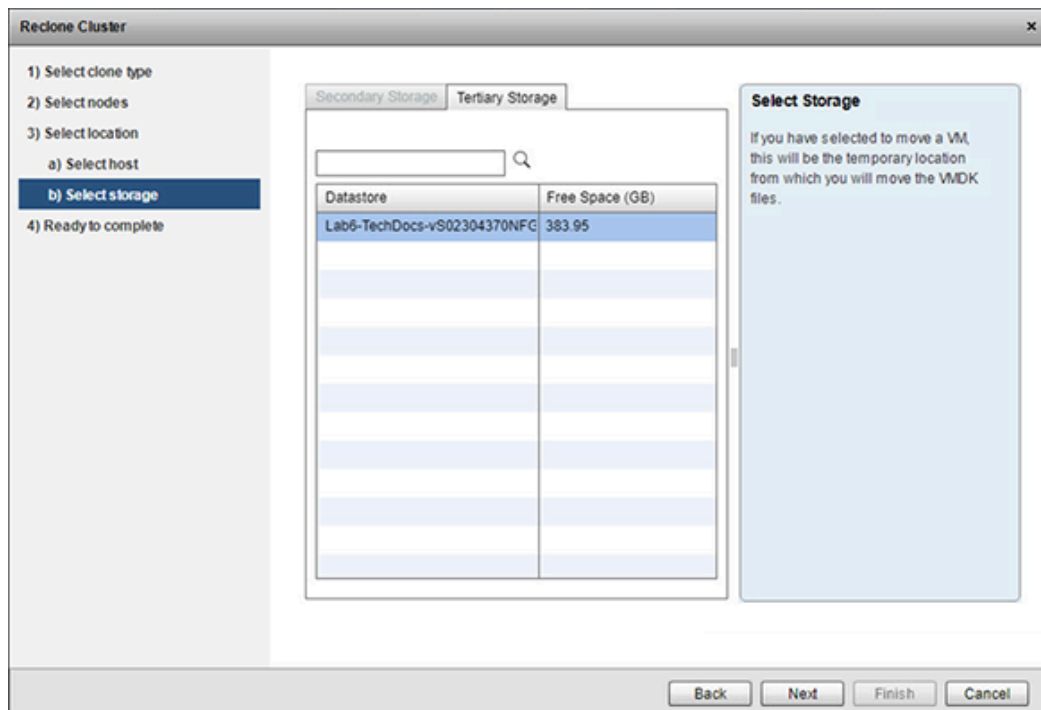
- **Reclone to the current location.** Uses the current location of the selected node as clone host. The Delete the original VM option allows the replacement of the selected node with its new clone.
- **Reclone to a new location.** Allows the selection of a new host and storage for the node's clone. When choosing this option, the **Select host** and **Select storage** sub-sections will be available:
 - The *Select host* page allows you to choose a target host through VMware vCenter Server, where the new clone will be created. The clone will be powered-off after creation. The host selection is available for each node marked for recloning.

Figure 8-25. Select host



- The *Select storage* page allows you to select the storage location for each node. When creating powered-off clones for manual transfer, the VMDK files of the cloned VMs will be stored here.

Figure 8-26. Select storage



Note

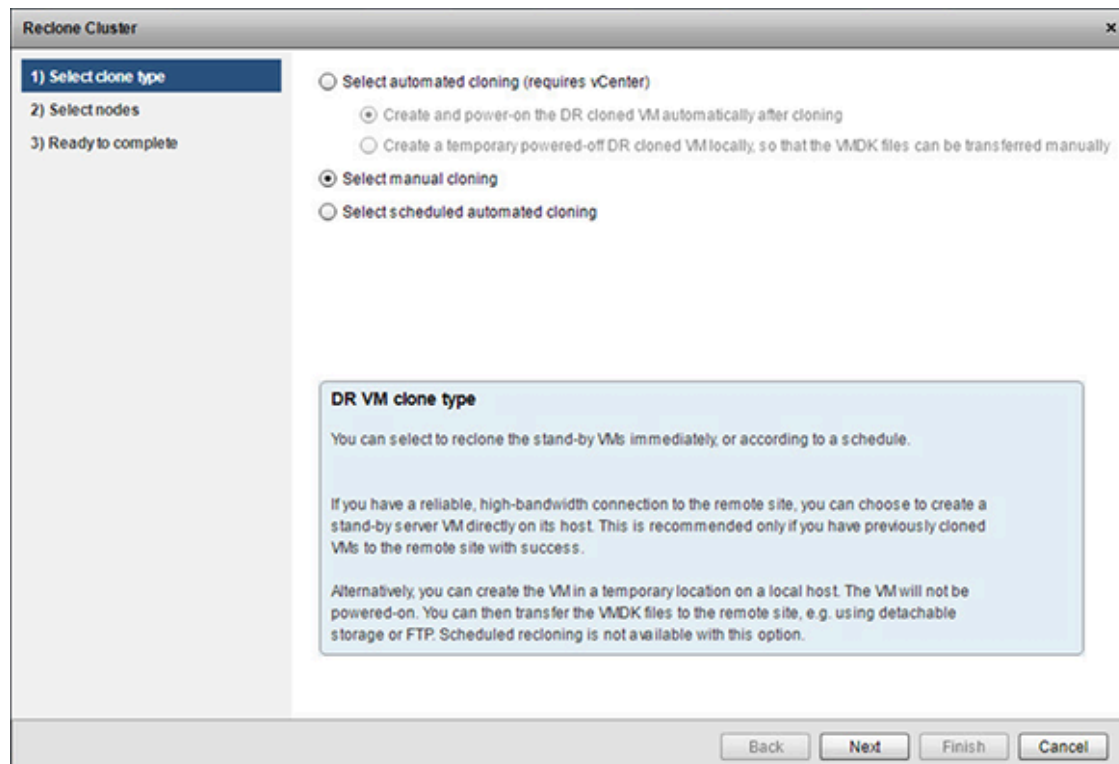
If the new location has not been configured for the recloning job or the EMS is not able to retrieve the original reclone target location from the Secondary or Tertiary nodes, the VMware vCenter Server Default Host will be used for the target reclone location. The Default Host needs to be configured in the [Configure Connection to VMware vCenter Server](#).

The **Ready to complete** page shows the summary of the recloning task. The automated recloning task will be executed immediately after its configuration, as soon as the cluster is in a ready state.

5.2. Manual Recloning

This cloning task is manually executed by the user. The manual recloning task can be executed at any time after its configuration, when initiated by the user.

Figure 8-27. Reclone manual cloning



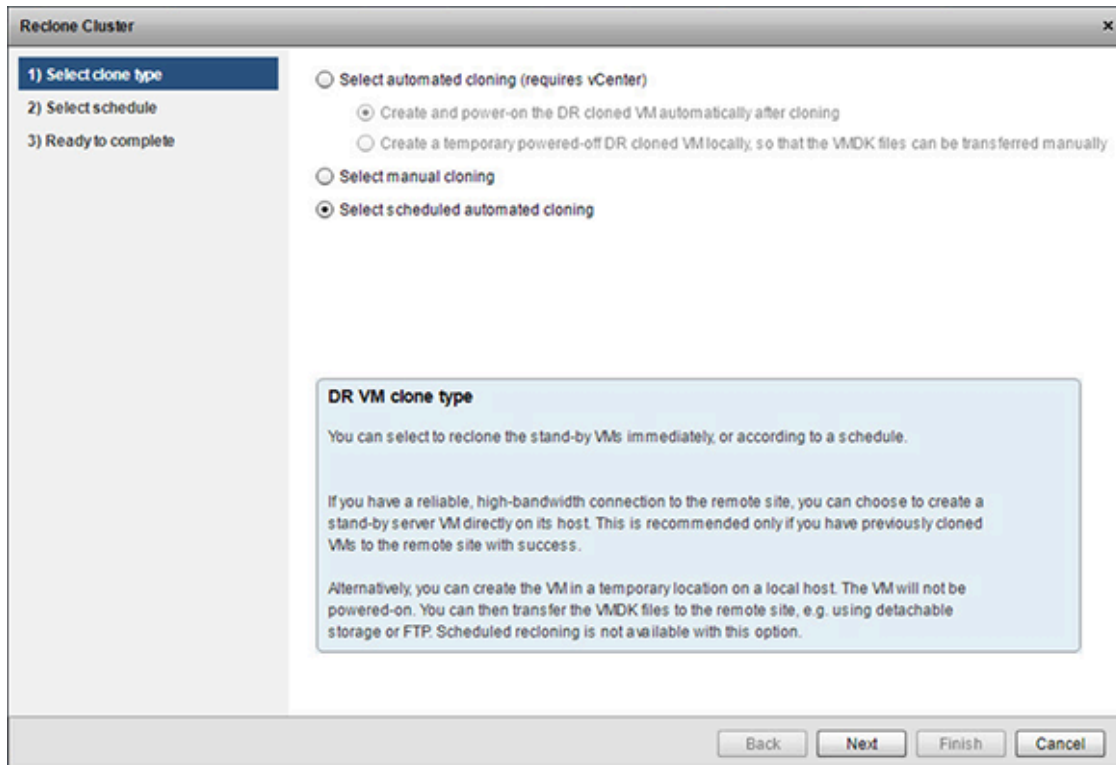
Once the manual cloning option is chosen, the nodes to be cloned are available for selection in the *Select nodes* page. Just like in Automated reclone, a **full cluster reclone** means that all available passive nodes will be cloned (manually, by the user). If the cluster has a Tertiary node besides its Secondary, the partial cluster reclone option is available. This allows the recloning of only the Secondary or Tertiary nodes, or both, if selected so.

The *Ready to complete* page shows the summary of the recloning task. The task will not be executed unless initiated by the user.

5.3. Scheduled Recloning

Allows the scheduling of a recloning task that executes in the same manner as the Automated recloning option, but at a specified time and using the specified repetition pattern.

Figure 8-28. Select scheduled automated cloning



Note

Scheduled recloning may not be available for Engine clusters upgraded from older versions until at least one automated reclone operation is performed before using the Scheduled recloning feature.

After selecting the Scheduled automated cloning option, the *Select schedule* page offers the possibility to enable the schedule and configure it.

Figure 8-29. Select schedule

The screenshot shows the 'Reclone Cluster' dialog box with the '2) Select schedule' step selected. The left sidebar shows three steps: '1) Select clone type', '2) Select schedule' (highlighted), and '3) Ready to complete'. The main area is titled 'Schedule automatic recloning:' and contains the following options:

- ☒ **Enable schedule** (with a 'Clear Schedule' button next to it)
- ☒ **Once every month on the** 15th (dropdown)
- ☐ **Twice every month on the** 1st and 15th (dropdown)
- ☐ **Every second week on** Monday (dropdown)

Below these are the 'Begin recloning:' options:

- ☒ **Starting at** 2:00 (dropdown)
- ☐ **Some time between the hours of** 2:00 (dropdown) **and** 3:00 (dropdown)

At the bottom, there are two checkboxes:

- ☐ **Delete original Secondary VM**
- ☐ **Delete original Tertiary VM**

The bottom of the dialog has four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

The automated recloning execution date can be programmed as follows:

- **Once every month.** The reclone task is executed in the selected day of every month.
- **Twice every month.** The reclone task is executed in the two selected days of every 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.** The reclone task can be scheduled to run once in two weeks, on the selected day. This allows a greater flexibility for scheduling two weeks reclone intervals without taking in consideration the length of the months.

The time of the execution set using the following options:

- **Starting at.** The reclone task is executed at the selected hour during the planned execution day(s). The Engine will wait for the cluster to be in the ready state within the span of the selected hour. If the cluster becomes ready for recloning within the time frame, the reclone task will be executed. If the cluster does not enter the ready state within the 60 minutes time frame, the reclone task will not be executed.
- **Some time between the hours.** The reclone task is executed in the time frame between the selected hours, as soon as the cluster is in the ready state. If the cluster does not enter the ready state within the specified time frame, the reclone task will not be executed. This option allows the configuration of a time frame bigger than 60 minutes in which the cluster can be ready for recloning.

The original passive nodes can be deleted after cloning if the **Delete original VM** option is selected.

A schedule can be reset (or cleared) using the **Clear Schedule** button available in the top side of the page.

The *Ready to complete* page shows the summary of the recloning schedule and the reclone task that will be performed.

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.

6. Deploying Neverfail Engine 8.5 Cluster in Amazon Web Services Cloud Environment

This topic covers the scenario in which the Neverfail Engine 8.5 Cluster is deployed in an Amazon Web Services (AWS) Cloud environment.

Networking and Topology Prerequisites

The deployment of Engine 8.5 Cluster in an Amazon Web Services environment depends on the following topology variables:

- **Amazon Web Services Account** - the same AWS account is recommended for deploying an Engine 8.5 Cluster. Cross-account deployment is possible but not recommended.
- **Virtual Private Cloud (VPC)** - the cluster deployment is possible on the same Virtual Private Cloud or on different Virtual Private Clouds.
- **Availability Zone (AZ)** - the cluster deployment is possible on the same Availability Zone or on different Availability Zones.
- **AWS Region** - the cluster deployment is possible on the same AWS Region or on different AWS Regions.

Considering the mentioned scenarios (see above), different network topologies are available for the Engine 8.5 Cluster deployment. However, as long as the **Engine nodes are able to communicate between themselves through a channel connection**, the Engine 8.5 Cluster will function correctly (considering the below mentioned Known Limitations).

The only challenge in deploying the Engine 8.5 Cluster in an AWS environment is meeting the required interconnectivity conditions. Fortunately, the procedures for interconnecting different instances, given the topology variables discussed above, are well documented by Amazon:

- Inter AWS Region connection: [Multiple Region Multi-VPC Connectivity](#)
- VPC-to-VPC peering:
 - [Creating and Accepting a VPC Peering Connection](#)
 - [Updating Your Route Tables for a VPC Peering Connection](#)
 - [Configurations with Routes to an Entire CIDR Block](#)
- Cross region AMI copying (required for manual cloning of DR node to a different VPC): [Cross Region EC2 AMI Copy](#)

Known limitations

1. **Engine 8.5 nodes should be deployed using the same AWS account.**

2. Engine supported Public IP scheme

Depending on AWS EC2 cloud instances abstraction layer, **you cannot have the same Engine Public IP shared between two different AWS instances. This means that the Engine traditional HA Pair or HA+DR Trio installs won't work in an all cluster AWS cloud deployment.** Hence, these are the supported scenarios for a AWS-to-AWS-(to-AWS) topology. Please note that some of the scenarios require manual reconfiguration so that the *different public IPs condition* is met.

- HA Pair with **different** Public IP addresses (requires post-install manual reconfiguration of different Public IPs - done from Configure Server Wizard on each node).
- DR Pair with different Public IPs on each node

Note

DR Pair with same Public IPs on each node installation requires post-install manual reconfiguration of different Public IPs - done from Configure Server Wizard on each node. The resulting configuration is equivalent with the one above.

-
- HA+DR trio with **different** public IPs for the HA nodes AND DR node: (requires post-install manual reconfiguration of different Public IPs - done from Configure Server Wizard on each HA node).

3. Engine nodes cloning type

Supported cloning type: manual using the AWS cloning approach, i.e.

- a. Go to the AWS instance
- b. Select the instance and click on instance > action
- c. Create image
- d. Launch another instance from the image created above

6.1. Installing Neverfail Engine 8.5 DR Pair in different Amazon Web Services VPCs

This use case example describes the procedure of deploying an Engine 8.5 Cluster in a AWS-to-AWS DR topology with Primary and Secondary nodes sitting in different VPCs, having the same AWS Region and Availability Zone.

Note

- This procedure can be applied also to the other supported Engine Pair scenarios - the only differences are the way the networking prerequisites between Engine nodes are implemented and the eventual manual public IP post-install configurations.
- Trio HA+DR deployments require to manually clone the Secondary node out of Primary then to add a DR Tertiary node by cloning it manually out of Secondary. The cloning procedure is the same for both Secondary and Tertiary nodes.

- Static routes definition may be required on the Engine nodes defined in different subnets in the same VPC. This depends on how the routing is defined between the subnets at VPC level.

Step 1. Networking Prerequisites

- Existing AWS account with 2 VPCs defined in the same Region and Availability Zone as described below.
- Primary-to-be server located in the Production VPC-A site.

VPC Configuration

- VPC-A (Production site)
 - IPv4 CIDR: 172.31.0.0/16
 - Subnet used by Production site: **172.31.71.0/24**
- VPC-B (DR site)
 - IPv4 CIDR: 172.32.0.0/16
 - Subnet used by Production site: **172.32.73.0/24**
- VPC-A to VPC-B Peering connection defined: [Configurations with Routes to an Entire CIDR Block](#)

Figure 8-30. Peer Connection

Create Peering Connection Actions ▾

Filter by tags and attributes or search by keyword

Name	Peering Connection ID	Status	Requester VPC	Accepter VPC	Requester CIDRs	Accepter CIDRs
A-2-B-peering	pcx-a5e1a6cc	Active	vpc-5d206934 V...	vpc-5674e33e V...	172.31.0.0/16	172.32.0.0/16

Peering Connection: pcx-a5e1a6cc

Description DNS Route Tables Tags

Requester VPC owner	507357548496	Accepter VPC owner	507357548496
Requester VPC ID	vpc-5d206934	Accepter VPC ID	vpc-5674e33e
Requester VPC Region	Ohio (us-east-2)	Accepter VPC Region	Ohio (us-east-2)
Requester VPC CIDRs	172.31.0.0/16	Accepter VPC CIDRs	172.32.0.0/16
VPC Peering Connection	pcx-a5e1a6cc	Peering connection status	Active
Expiration time	-		

- Routing table updated with correct peering and subnet association

Figure 8-31. VPC-A Routes

<input type="checkbox"/>	Name	Route Table ID	Explicitly Associat	Main	VPC
<input checked="" type="checkbox"/>	VPC-A-31	rtb-bcdafed5	2 Subnets	Yes	vpc-5d206934 VPC-A (prod)

rtb-bcdafed5 | VPC-A-31

Summary Routes Subnet Associations Route Propagation Tags

Edit

View: All rules

Destination	Target	Status	Propagated
172.31.0.0/16	local	Active	No
0.0.0.0/0	igw-9c841cf5	Active	No
172.32.0.0/16	pcx-a5e1a6cc	Active	No

Figure 8-32. VPC-B Routes

<input checked="" type="checkbox"/>	VPC-B-32	rtb-2447094c	1 Subnet	Yes	vpc-5674e33e VPC-B (DR)
-------------------------------------	----------	--------------	----------	-----	---------------------------

rtb-2447094c | VPC-B-32

Summary Routes Subnet Associations Route Propagation Tags

Edit

View: All rules

Destination	Target	Status	Propagated
172.32.0.0/16	local	Active	No
0.0.0.0/0	igw-38372f51	Active	No
172.31.0.0/16	pcx-a5e1a6cc	Active	No

- Each VPC Security Group updated with a inbound rule for allowing the traffic from the **remote/peered** subnet source, i.e.
 - Allow traffic from 172.32.73.0/24 source for VPC-A SG
 - Allow traffic from 172.31.71.0/24 source for VPC-B SG

Important

In all the next steps, you must assign to each EC2 instance (EMS, Primary, Secondary) ALL the private IP addresses that will be used on that given box, i.e. public IP-to-be, Channel IPs, Management IP (if exists).

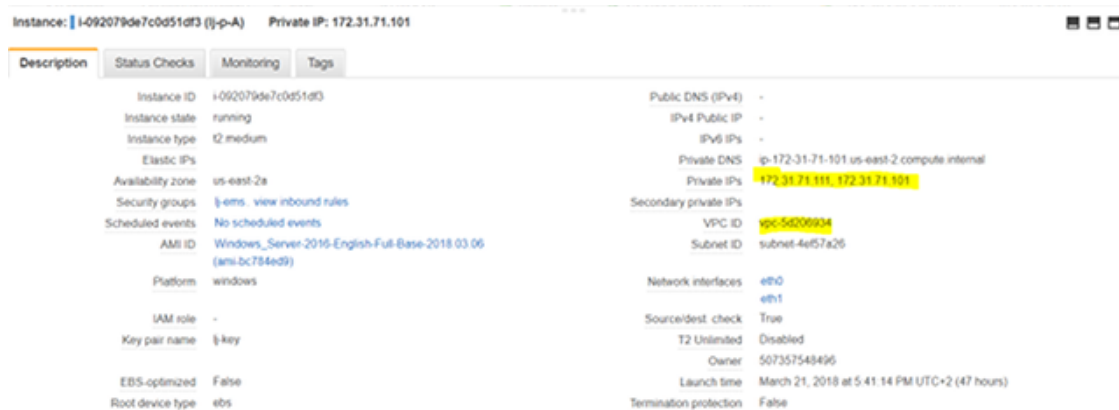
Step 2. Installing the Engine Management Service (EMS)

The EMS can be installed on any supported workstation or server which can access VPC-A and VPC-B subnets. The EMS installation procedure is described in the Engine 8.5 Installation Guide.

Step 3. Installing Engine on Primary Server located in VPC-A (production site)

- Using the EMS, install Engine on the Primary server as indicated in the Engine 8.5 Installation Guide.
- Make sure all the Public, Channel and Management IPs defined and managed by Engine are registered/configured on the AWS EC2 Primary instance. For example:
 - Public IP: 172.31.71.101/24
 - Channel IP: 172.31.71.111/24

Figure 8-33. Instance IPs



Step 4. Adding Secondary/DR Node located in VPC-B (DR site)

- Using the EMS, add a Secondary DR node as indicated in the Engine 8.5 Installation Guide.
 - Configure Secondary with the correct Public and Channel IP addresses:
 - Public IP: 172.32.73.102/24
 - Channel IP: 172.32.73.112/24

Configure also the correct GW and DNS server corresponding to the DR VPC-B site.

Figure 8-34. Public Address

Add a Stand-by Server For Disaster Recovery

- 1) Select public IP address
- 2) Select channel IP addresses
- 3) Select clone type
- 4) Select host (optional)
- 5) Select storage (optional)
- 6) Configure helper VM (optional)
- 7) Ready to complete

☐ The public (principal) IP address will be identical to the Primary server
☒ The public (principal) IP address will be different than on the Primary server

Network adapter: **P**

IP address	Subnet mask
172.32.73.102	255.255.255.0

Buttons: Add... Remove

Enter the gateway: 172.32.73.1

Enter the preferred DNS server: 172.32.73.1

Enter the alternate DNS server (optional):

Enter the user name for updating DNS servers: administrator

Enter the password: *****

Public IP Addresses

If the Primary and DR site use different subnets, the DR server requires a separate public IP address.

In this case, an account capable of updating the DNS servers must be specified.

On switchover or failover, DNS servers will then be updated with the IP address of the active server.

Buttons: Back Next Finish Cancel

Figure 8-35. Channel Address

Add a Stand-by Server For Disaster Recovery

- 1) Select public IP address
- 2) Select channel IP addresses
- 3) Select clone type
- 4) Select host (optional)
- 5) Select storage (optional)
- 6) Configure helper VM (optional)
- 7) Ready to complete

Primary server to Secondary server | Secondary server to Tertiary server | Tertiary server to Primary server

Select a network adapter for the channel: **Ch**

Enter an IPv4 address for the Primary: 172.31.73.111

Primary Subnet Mask (blank for default): 255.255.255.0

Enter an IPv4 address for the Secondary: 172.32.73.112

Secondary Subnet Mask (blank for default): 255.255.255.0

Channel IP Addresses

The addresses will be automatically added to each server to allow Engine to communicate and replicate data.

A persistent static route should be configured for the channel connection where routing is required

Buttons: Back Next Finish Cancel

- Choose the **manual cloning** type then when requested
- Wait until EMS informs you that the Secondary Server is ready to be manually cloned, then proceed with the AWS cloning as indicated above.
 - Go to the AWS Primary instance.
 - Select the instance and click on *Instance > Action > Create image*.
- Launch AWS Secondary instance from the Primary AMI image created above.
 - Make sure the Secondary instance type matches (at least) the CPU, memory, storage, and networking capacity configured for Primary instance.

- Configure the instance with ALL the IPs that will be used on the Secondary server (Public, Channel, Management IPs).

Figure 8-36. Network Interfaces

▼ Network interfaces ⓘ

Device	Network Interface	Subnet	Primary IP	Secondary IP addresses
eth0	New network interface ▼	subnet-29c74f41 ▼	172.32.73.102	Add IP
eth1	New network interface ▼	subnet-29c74f41 ▼	172.32.73.112	Add IP

- Configure the VPC-B corresponding Security Group.
- When Secondary instance is created it will connect automatically to the Primary.

Step 5. Post-installation and other deployment considerations

- **Trio:** for HA+DR trio deployment, repeat the Step 4 above, having as source the Secondary Server. Make sure the Tertiary instance is configured accordingly with the type, IP addresses, security group, etc.
- **Configuring Static Routes for Channel traffic:** If these are required then configure them as indicated in the [How to Stretch LAN to WAN in Neverfail IT Continuity Engine in a Primary - Secondary Configuration](#) Knowledge Base article (also in the following chapter).
- **Switchover considerations:** EMS Server should access both Primary and DR subnets. Otherwise, when Secondary is active it cannot connect to the newly active.

Figure 8-37. Engine Management Service

The screenshot displays the 'Neverfail CE Management Service' interface. The main window shows the 'Status' tab for a server pair. On the left, a diagram shows a 'Primary' server (172.31.71.101) and a 'Secondary' server (172.32.73.102) connected by a green arrow. Below this, there are checkboxes for 'Show IP Addresses', 'Public address', and 'Channel address'. The 'Plan Execution' section is empty. On the right, the 'Summary Status' section lists various metrics: Name (i-p-aws), Install Status, Product Version (VB 5 Update 1 (30192)), License Status (Expires in 3534 days), Active Server (Primary), Application State (Started - OK), Client Network (OK), Primary Status (Replicating), Data on Primary (Active), Secondary Status (Replicating), and Data on Secondary (Synchronized - Recovery Point (seconds): 0.0). Below this, the 'Applications and Platforms' section shows a table with columns 'Name' and 'Health', listing FileServer, System, and User Defined, all with 'OK' health. At the bottom, the 'Neverfail CE Management Service Tasks' section shows a table with columns 'Name', 'Target', 'Date Time', 'Status', and 'Detail'. It lists two tasks: 'Reconfigure Primary' and 'Install Primary', both completed successfully on 23 Mar 2018.

7. Deploying a Passive Node in an Amazon Web Services Cloud Environment

This topic covers the scenario in which the Secondary (or Tertiary, if applicable) cluster node, part of an on-premise Neverfail Engine Cluster, is moved into an Amazon Web Services (AWS) Cloud environment.

The procedure required to accomplish the described task is called stretching (LAN to WAN stretching) can be summed up as follows:

1. The Secondary Engine node must be prepared for movement.
2. The AWS Cloud environment must be configured to run the Secondary node.
3. The Secondary node must be imported in the AWS Cloud environment.
4. The Neverfail Engine must be configured on the moved Secondary and on the on-premise Primary.

Detailed Scenario

Initial state: the Engine cluster is hosted on-premise and has a Primary and a Secondary HA node. The DNS server is also hosted on-premise, in the same network environment. The two cluster nodes and the DNS server communicate through a local area network (LAN).

Example IP initial configuration:

IP	Value
Public and Channel subnet	192.168.69.x/24
Public IP	192.168.69.10
Primary Channel IP	192.168.69.211
Secondary Channel IP	192.168.69.212
DNS server IP	192.168.69.1

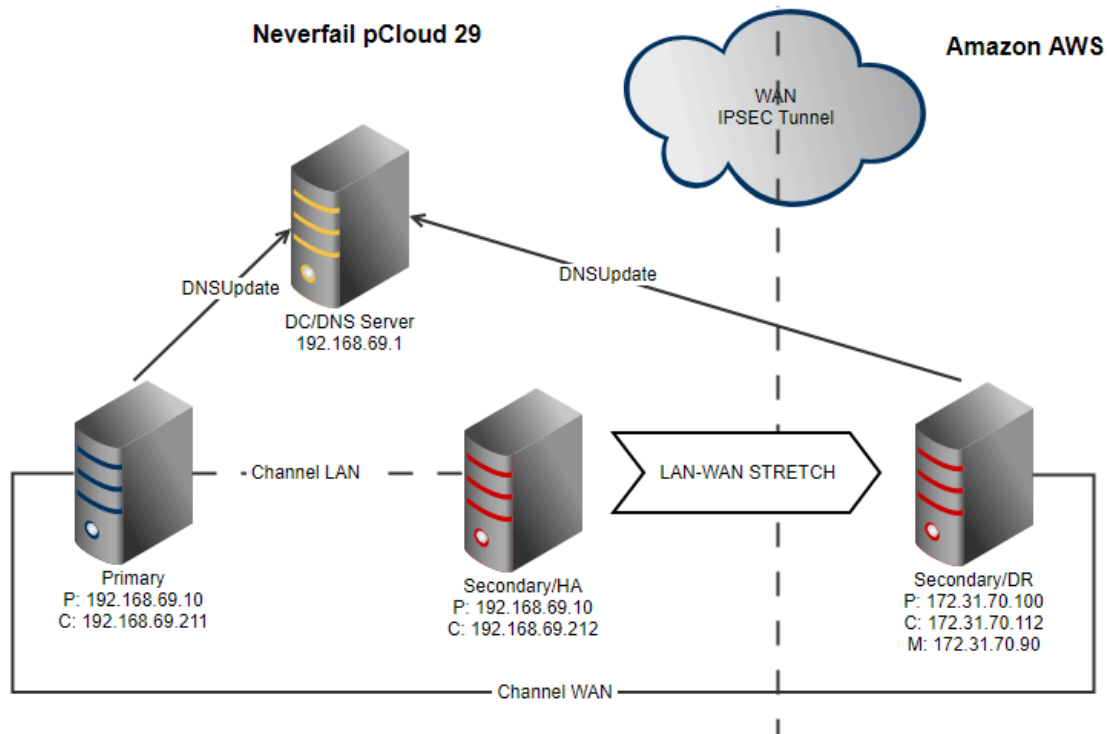
Target state: the Engine cluster is stretched (LAN - WAN stretching) so the Primary node and the DNS server reside in the same on-premise, LAN environment, while the Secondary node resides in the AWS Cloud environment. The AWS hosted Secondary is now serving as a DR node in the Engine cluster and is communicating with the Primary node and DNS server via WAN.

Example IP target configuration:

IP	Value
Primary Public and Channel subnet	192.168.69.x/24
Primary Public IP	192.168.69.10
Primary Channel IP	192.168.69.211
DNS server IP	172.31.70.x/24
Secondary Public and Channel subnet	172.31.70.x/24
Secondary Public IP	172.31.70.100

The following diagram illustrates this scenario:

Figure 8-38. LAN - WAN Stretching



Performing the Stretching Procedure

The stretching procedure implies the steps described below.

1. Creating the IPsec tunnel.

The tunnel is required for establishing a connection and allowing traffic between the on-premise Engine host and the AWS Cloud.

2. Preparing the Secondary node for stretching.

Note

The scenario implies that the Primary node is active and replicating on the Secondary passive node.

The Secondary cluster node must be prepared before moving it to AWS Cloud. Make sure to follow the steps below:

- Add the Domain Admin account to Neverfail Engine:
 - a. Login to the Neverfail Advanced Client.
 - b. Click **Application**.
 - c. Open the Tasks tab and highlight Neverfail Engine.
 - d. Click **User Accounts...**, click **Add**, and enter the Domain Admin account details.

e. Click **OK** and click **Close**.

- Shutdown Engine on all nodes.
- On the Secondary node, set the Engine service startup to **manual**.
- On the Secondary node, configure the UTC time settings as [required by AWS](#).
- On the Secondary node, set the IPv4 to **DHCP**. The IPv6 protocol should be disabled.
- On the Secondary node, remove all removable drives (including CD and Floppy) and network drives.
- Shut down the Secondary node.
- Export the Secondary node as OVA.

3. Setting up AWS Cloud.

Follow the AWS setup procedures detailed here to set up your AWS Cloud: [Setting Up with Amazon EC2](#).

4. Importing the Secondary OVA to AWS.

The detailed procedure is described here: [Importing a VM as an Image Using VM Import/Export](#). The essential steps are listed below:

- Create an Amazon S3 Bucket, required for uploading OVA images to AWS (prior to importing). Check out the procedure here: [How Do I Create an S3 Bucket?](#).
- Upload the Secondary OVA to the AWS S3 bucket.
- Create a new IAM user. This user will be used to import the virtual machine to AWS using the AWS CLI. While being logged with the root account in AWS console, go to **Services > Users** and select **Add User**. Note down the Access key ID and Secret key (they are required for CLI connections to AWS).

Note

Make sure to add sufficient permissions to the new AWS user.

-
- Install the AWS Command Line Interface, required by the VM Import Service Role. Check out this link for more information: [Install the AWS Command Line Interface on Microsoft Windows](#).
 - Create the **VM import Service role**. Check out this document for more information: [Importing a VM as an Image Using VM Import/Export](#).

After being created, the vmimport role should be available under IAM > Roles.

- Edit the `role-policy.json` file. Use the file bucket ARN where indicated as disk-image-file-bucket. The file bucket ARN is displayed when hovering and clicking the S3 bucket row.
- Attach the policy to the role created above as indicated in AWS import role procedure.
- Import the Secondary OVA to AWS AMI (considering that the upload is complete): Edit the `containers.json` file inserting the appropriate S3 bucket name and OVA file name (as displayed in the bucket) then import the OVA to AMI using the AWS CLI (as instructed in the AWS procedure). Check periodically for import task status, as instructed.

Note

The **ImportTaskId** will be used to check the status of the import operation.

5. Configuring and launching the EC2 Instance.

- Go to **AWS > EC2 > AMIs** and select the image imported at previous step.
 - Click **Launch** and continue with instance configuration prior of launching.
 - Choose the instance type which should match your Primary's resources.
 - Configure the instance details.
-

Important

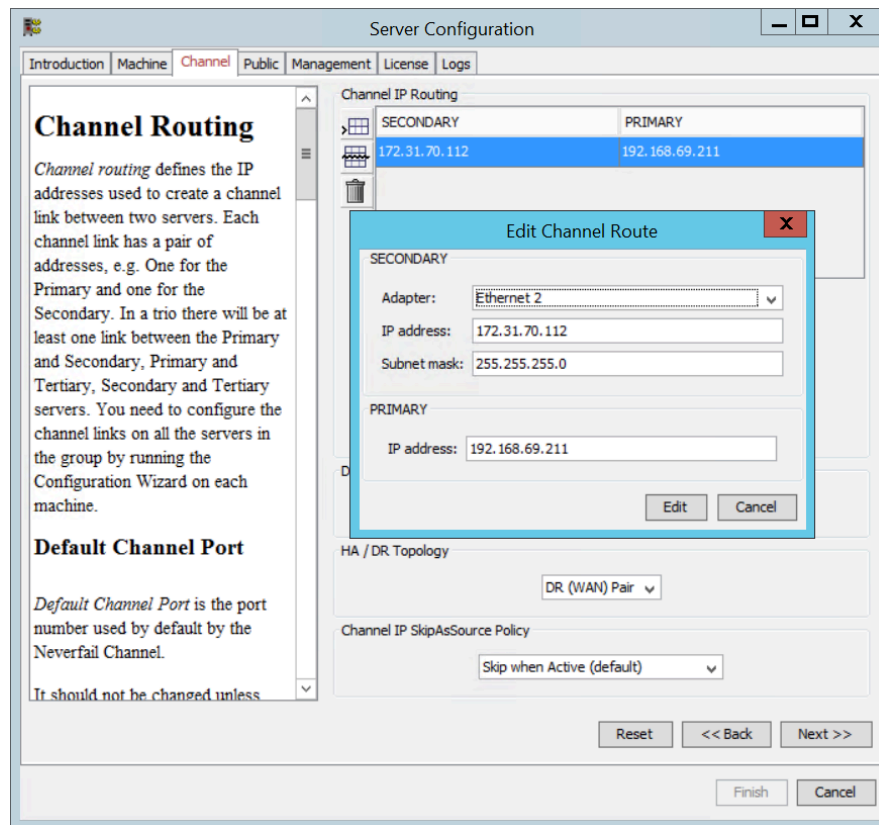
Choose the subnet configured for IPSec tunneling between the on-premise host and AWS. Also, enable the public IP auto assignation in order to be able to access (RDP) the AWS-hosted Secondary instance from a remote host (from a different subnet).

- Configure the desired storage size.
- Configure the Security group to allow traffic between the on-premise and the AWS subnets. Here you can also configure the access for RDP connections from other management computers.
- Create a new key or use an existing one then connect to the new instance.
- When the instance status is **running**, you can access (RDP) the AWS-hosted Secondary.

6. Reconfigure the Engine cluster as an on-premise-to-AWS DR pair.

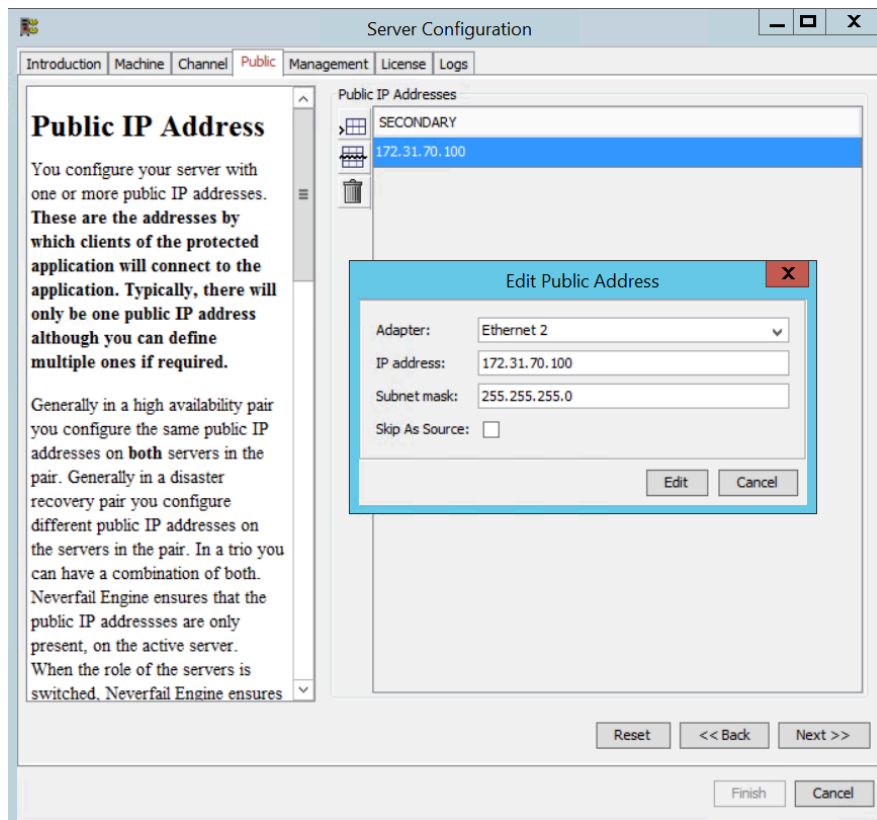
- **Configuring the migrated Secondary node.**
 - Launch *Configure Server Wizard* and Configure the new **Secondary Channel IP**. Make sure that an existing NIC adapter is selected.
 - Set the HA/DR Topology to **DR (WAN) Pair**.

Figure 8-39. Edit Secondary channel route



- Configure the new **Secondary Public IP**. Make sure that an existing NIC adapter is selected.

Figure 8-40. Edit public address



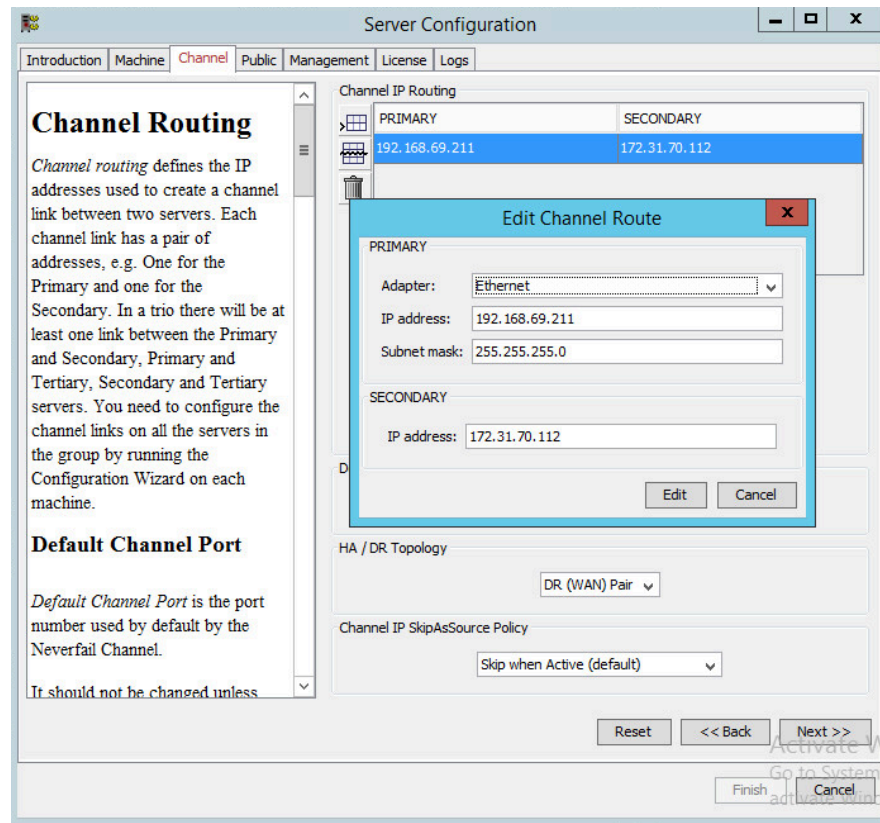
- Set the Engine service startup to **Automatic**.
- **Important:** On AWS console, select **Instance > Actions > Networking > Manage IP addresses** and assign to your EC2 instance ALL the private IP addresses that will be used on Secondary, i.e. public IP address Channel IP address, Management IP address (if

exists). This will allow traffic between the mentioned IP addresses and remote on-premises site.

- **Configuring the Primary node.**

- Launch the Configure Server Wizard and set the new **Secondary Channel IP**.
- Set the HA/DR Topology to **DR (WAN) Pair**.

Figure 8-41. Edit Primary channel route



- Start the Engine on both Primary and Secondary nodes.

- **Configuring DNSUpdate task.**

On the Primary Server, connect to **Advanced Management Client** and add two Network Configuration tasks, as follows:

- For the Primary server, select **Primary** radio button.
- Edit command `DNSUpdate -auto`.
- Click **Run As** and select from the menu the Domain Account previously configured in the *User Accounts* dialog.
- For the Secondary server, select **Secondary** radio button.
- Edit command `DNSUpdate -auto`.
- Click on **Run As** and select the Domain Account previously configured in the *User Accounts* dialog.

- **Configuring static routes for channel traffic.**

If static routes are required (for multi-NIC install), configure them as indicated here:

- Open **Routing and Remote Access** from Administrative Tools.
- Select the server name, then from the Action menu select **Configure and Enable Routing and Remote Access** to launch the configuration wizard.
- Select **Custom Configuration > LAN routing** and verify that the RRAS service is started.
- Select the server again, navigate to IP Routing and select **Static Routes**.
- From the Action menu select **New Static Route**.
- From the drop-down, select the channel interface and enter the destination channel IP followed by the mask 255.255.255.255 and the source machine gateway.
- Test the channel routing using the following command to ensure that all the packets will be sent using the channel IP and not the principal (public) IP.

```
pathping -n Channel_IP
```

Note

For a trouble free WAN implementation, it is recommended that you use RRAS for implementing static routes. Avoid using the interface ID when creating static routes using the "route" command because the interface ID is dynamic and increments each time a server is restarted or a NIC is disabled/enabled, and this change will make the route invalid.

Note

The persistent routes ensure that any communication with the channel network is in fact established via the physical channel NICs.

• **Configuring DR ping targets.**

On the Primary server update the Server Monitoring ping routing configuration:

- In the Neverfail Advanced Client, select Server Monitoring.
- On the Server Monitoring screen, in the Configure Pings section, click **Configure...**
- Browse to the Ping Routing tab of the new window.
- Update the Primary and Secondary IP addresses to match the new IP scheme implemented in the steps above. Update both the Ping From and Ping To fields.

• **Configuring VMware Tools service protection.**

Issue: Considering initial HA pair was a VMware V2V, the VMware Tools service is installed and protected by the vSphere Integration plugin. The service is also present on the stretched AWS instance, though it cannot be started in this new context. Thus. if a switchover is

attempted, the service will be attempted to be started on the S/A but, as expected, will fail (application errors will be logged).

WORKAROUND (use one of the following):

- uninstall vSphere Integration plugin if this is not used anymore when Primary is active
- modify target states for the VMware Tools service to be Any/Any: This can be done by Adding the service to User Defined protection while Primary is active

- **Switchover considerations.**

The Engine Management Server should access both Primary and DR subnets. Otherwise, when Secondary is active it cannot connect to the newly active.

Chapter 9. Troubleshooting

Related information

- [Two Active Servers](#)
- [Two Passive Servers](#)
- [Invalid Neverfail Continuity Engine License](#)
- [Synchronization Failures](#)
- [Channel Drops](#)
- [MaxDiskUsage Errors](#)
- [Application Slowdown](#)

1. Two Active Servers

Condition

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.

Solution

Cause

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 `nbstat -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.

Remedy

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 **P/-** (Primary and passive) or **S/-** (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 **P/A** (Primary and active) or **S/A** (Secondary and active).
11. Log into the Neverfail Advanced Management Client.
12. Verify that the servers have connected and replication has started.

2. Two Passive Servers

Condition

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.

Solution

Cause

- 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.

Remedy

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.

3. Invalid Neverfail Continuity Engine License

Condition

The Neverfail Continuity 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.

Solution**Cause**

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.

Remedy

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.

4. 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.

4.1. Services Running on the Passive Server

Condition

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

Solution

Cause

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 Continuity 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.

Remedy

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.

4.2. Neverfail Channel Incorrectly Configured

Condition

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*.

Solution

Cause

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

Remedy

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.

4.3. Incorrect or Mismatched Disk Configuration

Condition

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.

Solution

Cause

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**`.

Remedy

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?

4.4. The Passive Server has Less Available Space than the Active Server

Condition

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.`

Solution

Cause

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.

Remedy

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.

4.5. Unprotected File System Features

Condition

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*.

Solution

Cause

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](#).

Remedy

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.

4.6. 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.

4.6.1. Resource Issues

Condition

Neverfail Engine logs the following error message:

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

Solution

Cause

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

Remedy

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

4.6.2. Registry Security Issues

Condition

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

Solution

Cause

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.

Remedy

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

5. Channel Drops

5.1. Performance Issues

Condition

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.

Solution

Cause

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 Continuity Engine R2 service itself experiences problems and may hang or terminate unexpectedly.

Remedy

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.

5.2. Passive Server Does Not Meet Minimum Hardware Requirements

Condition

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

Solution

Cause

A 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.

Remedy

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.

5.3. Hardware or Driver Issues on Channel NICs

Condition

The Neverfail Channel drops or disconnects and reconnects intermittently.

Solution

Cause

- 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

Remedy

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.

5.4. Firewall Connection

Condition

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.

Solution

Cause

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.

Remedy

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.

5.5. Incorrect Neverfail Channel Configuration

Condition

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

Solution

Cause

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.

Remedy

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.

5.6. Subnet/Routing Issues — In a LAN

Condition

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

Solution

Cause

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

Remedy

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.

5.7. Subnet/Routing Issues — In a WAN

Condition

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

Solution

Cause

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.

Remedy

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.

6. 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 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 (NFChannelExceededMaxDiskUsageException)`, the following conditions exist:

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

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

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

Condition

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.

Solution

Cause

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.

Remedy

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.

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

Condition

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.

Solution

Cause

- 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:

Remedy

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.

6.3. [L20]Out of disk space (NFChannelOutOfDiskSpaceException)

Condition

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

Solution

Cause

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.

Remedy

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.

7. 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.

7.1. Poor Application Performance

Condition

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

Solution

Cause

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.

Remedy

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

7.2. Servers Could Accommodate the Initial Load but the Load has Increased

Condition

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

Solution

Cause

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.

Remedy

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.

7.3. One Server is Able to Cope, but the Other Cannot

Condition

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

Solution**Cause**

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.

Remedy

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.

7.4. Scheduled Resource Intensive Tasks

Condition

Resource-intense scheduled tasks impact performance at certain times.

Solution**Cause**

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.

Remedy

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.

Chapter 10. Neverfail SCOPE Data Collector Service Overview

Related information

- [Using Neverfail SCOPE Data Collector Service](#)
- [Neverfail SCOPE Analysis Reports](#)

1. Using Neverfail SCOPE Data Collector Service

Daily Usage

The Neverfail SCOPE Data Collector Service collects configuration and performance data for pre-implementation analysis, license key generation, and assisting in support of Neverfail Continuity Engine.

The Neverfail SCOPE Data Collector Service runs as a service that requires no user intervention to log daily configuration and performance data. There is no need for any day-to-day user interaction with Neverfail SCOPE Data Collector Service. Log files can be collected and sent to Neverfail Support for analysis if desired.

Collecting Log Files

The Neverfail SCOPE Data Collector Service can be used both pre and post implementation of Neverfail.

- **Pre-Implementation**

Neverfail SCOPE Data Collector Service maintains a single file which is needed to obtain a pre-implementation report and to generate a license key. The data file created by Neverfail SCOPE Data Collector Service may be available as soon as 15 minutes after installing the collector service, but on systems with many shared files and folders the collection process can take an hour or more. If you require a full performance report you should wait at least 24 hours before collecting the file and sending it to Neverfail Support. The file contains the latest configuration data and the most recent 24 hours worth of performance data.

- **Post-Implementation**

To receive configuration or performance analysis you must collect the *Candidate for Upload files* and manually forward to Neverfail Support for analysis and report creation.

1.1. Configuring Neverfail SCOPE Data Collector Service

The SCOPE Configuration Tool

Neverfail strongly recommends contacting Neverfail Support staff to change these settings.

To use the *SCOPE Configuration Tool*, select **Start > All Programs > Neverfail > SCOPE > SCOPE Configuration Tool**. The *SCOPE Configuration Tool* opens in a new window.

The *SCOPE Configuration Tool* consists of four tabs: **General**, **Connectivity**, **Data Files** and **Support**. The features of each tab are described in the associated sections of this document.

Additionally, a link to *Neverfail SCOPE Data Collector ServiceOnline Help* can be found in the lower left corner of the window.

1.1.1. Configure the General tab

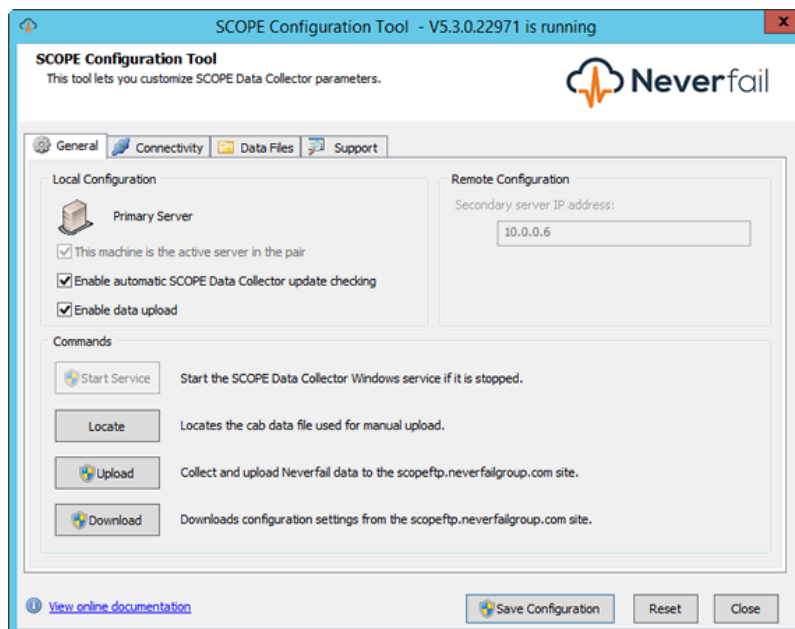
About this task

The *General* tab features controls for manually configuring IP addressing of the Secondary and Tertiary (if installed) servers, specifying the active server in the cluster, and enabling automatic update checking. The *General* tab also allows you to start the Neverfail SCOPE Data Collector Service Windows service, to upload collected Neverfail SCOPE Data Collector Service data, to download configuration settings from the and to locate the .CAB file for manual uploading.

Procedure

1. Select the **General** tab.

Figure 10-1. *SCOPE Configuration Tool - General tab*



Option	Description
Start Service	Starts the Neverfail SCOPE Data Collector Service Windows service if it is stopped.
Upload	Uploads the current .cab file typically located in the default location: <ul style="list-style-type: none"> On Windows 2008 installations: C:\ProgramData\ Neverfail-SCOPE \Data

Option	Description
	<ul style="list-style-type: none"> On Windows 2012 installations: C:\ProgramData\ Neverfail-SCOPE \Data <p>See expanded description below for more information about this feature.</p>
Download	Downloads configuration and Neverfail SCOPE Data Collector Service updates, if available, from the Neverfail Extranet
Locate	Locates the .cab files for manual upload

When you click **Upload**, Neverfail SCOPE Data Collector Service gathers all data. Do not close the application until it has finished gathering the data. After all data is gathered, Neverfail SCOPE Data Collector Service uploads it.

- After making configuration changes, click **Save Configuration** to save your changes, or click **Reset** to restore the default configuration.

1.1.2. Configure the Data Files tab

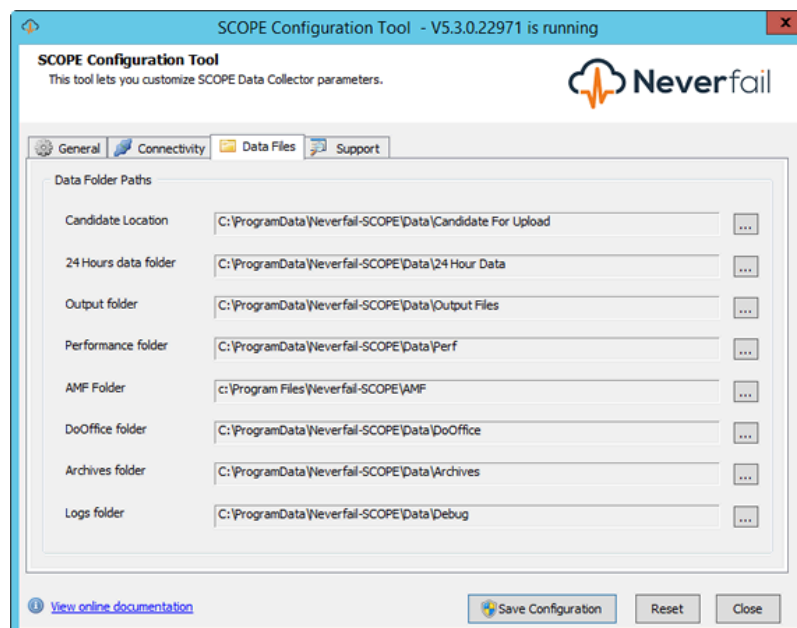
About this task

The *Data Files* section allows you to configure the file locations for Neverfail SCOPE Data Collector Service.

Procedure

- Select the Data Files tab.

Figure 10-2. SCOPE Configuration Tool - Data Files tab



Use the *Data Files* page to change the location where data files are stored.

2. After making configuration changes, click **Save Configuration** to save your changes, or click **Reset** to restore the default configuration.

1.1.3. Configure the Connectivity tab

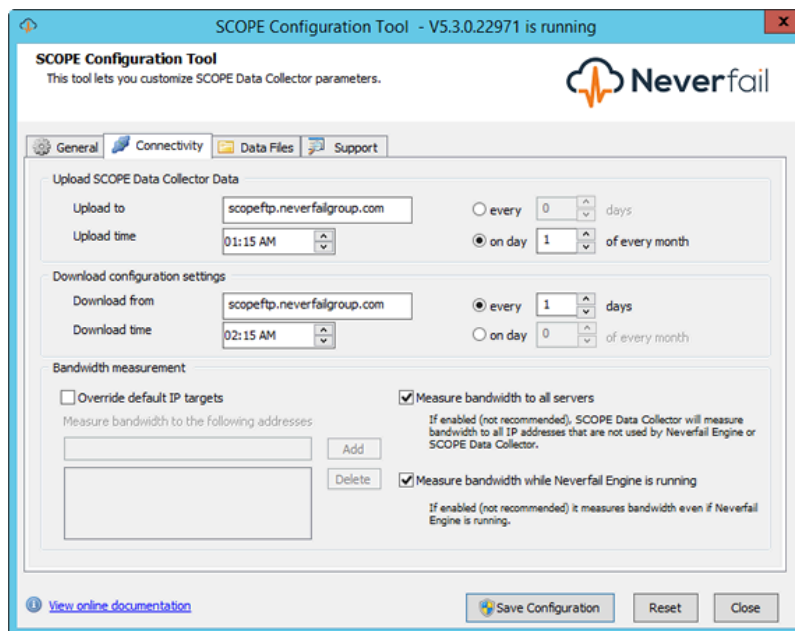
About this task

The *Connectivity* tab features controls for scheduling automated uploads of Neverfail SCOPE Data Collector Service data, downloads of Neverfail SCOPE Data Collector Service configuration data, and to configure bandwidth measurements.

Procedure

1. Select the **Connectivity** tab.

Figure 10-3. SCOPE Configuration Tool- Connectivity tab



The *Upload SCOPE Data Collector Data* pane in the **Connectivity** page provides ways to manually configure the upload destination address and select a schedule for automated uploads of Neverfail SCOPE Data Collector Service data. Scheduled uploads can follow a regular schedule of a set number of days (for example, every 7 days), or on a specified day (for example, on the 15th of the month). You also specify the time to perform the upload.

The *Download configuration settings* pane in the **Connectivity** page provides similar configuration settings to schedule automated downloads of Neverfail SCOPE Data Collector Service configuration data.

The *Bandwidth measurement* pane in the **Connectivity** page is used to configure bandwidth measurements. If you need to measure bandwidth using IP addresses other than the ones used for the Neverfail Channel, select the *Override default IP targets* check box and add new IP addresses by typing them into the text box and clicking **Add**. Remove IP addresses by selecting them from the list and clicking **Delete**.

Use the two check boxes on the right side of the *Bandwidth measurement* pane to measure the bandwidth between the local server and any other servers on the network running Neverfail SCOPE Data Collector Service but not running Neverfail Engine, or to measure bandwidth while Neverfail Engine is running.

By default, Neverfail SCOPE Data Collector Service does not measure bandwidth when Neverfail Engine is running to avoid overloading the busy Neverfail Channel. You can run Neverfail SCOPE Data Collector Service while Neverfail Engine is running if you use network connections for Neverfail SCOPE Data Collector Service that are separate from those used by Neverfail Engine. After configuring separate network connections for use by Neverfail SCOPE Data Collector Service, select the *Measure bandwidth while Heartbeat is running* checkbox.

To measure bandwidth to all servers in the Cluster using the Neverfail Channel, add their IP addresses and select the *Measure bandwidth to all servers* checkbox to prevent those IP addresses from being filtered out by default.

2. After making configuration changes, click **Save Configuration** to save your changes, or click **Reset** to restore the default configuration.

1.1.4. Configure the Support tab

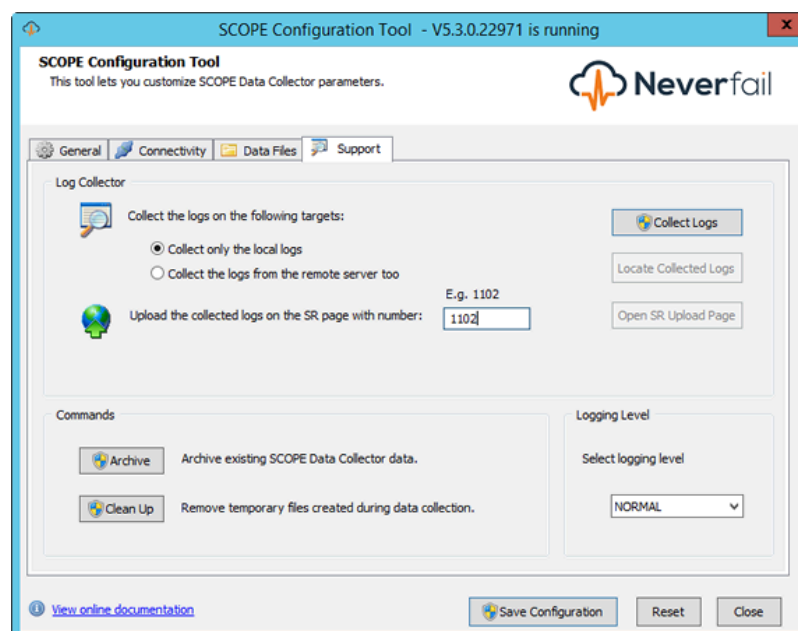
About this task

Use the controls on the *Support* tab to associate a Support Request number (S.R. number) with a specific set of Neverfail SCOPE Data Collector Service data, to control how this data is stored, and to select the logging level.

Procedure

1. Select the **Support** tab.

Figure 10-4. Neverfail SCOPE ConfigurationTool - Support tab



In the *Log Collector* pane of the **Support** page, type the SR (Support Request) number into the *Upload the collected logs on the SR page with number:* text box, then click **Open SR Upload Page**. The collected Neverfail SCOPE Data Collector Service data is uploaded to the SR.

Note

This action requires the server to have internet access, as the <https://neverfail.com/> page is opened to facilitate the upload.

Click **Collect Logs** to re-gather the Neverfail SCOPE Data Collector Service logs manually upon command. After re-gathering the logs, the **Locate Collected Logs** button becomes active and when clicked, automatically navigates to the location of the .CAB file.

In the *Commands* pane, click **Archive** to archive the existing Neverfail SCOPE Data Collector Service data, and click **Clean Up** to remove temporary files created during data collection. In the *Logging Level* pane, select a logging level (DEBUG or NORMAL) from the drop-down list

2. After making configuration changes, click **Save Configuration** to save your changes, or click **Reset** to restore the default configuration

1.1.5. Automatic Configuration

If *Enable automatic SCOPE Data Collector update checking* is selected on the **General** page, the service connects at the intervals specified and to the address specified in the *Download configuration settings* pane of the **Connectivity** page and downloads a `global.cfg` file (if available), which contains overrides for the default parameters. The values in the `global.cfg` file are stored in the registry and will override the existing local values.

The service then looks for the `<machineID>.cfg` file, where `<machineID>` is the globally unique ID (GUID) of the server, which was set when the machine first ran Neverfail SCOPE Data Collector Service. If the file is found, then the values in `<machineID>.cfg` file are stored in the registry and will override any existing values.

When *Enable automatic SCOPE Data Collector update checking* is selected, any manual configuration changes made using the local SCOPE Configuration Tool will be overridden by the `global.cfg` and/or `<machineID>.cfg` files. If you prefer to use a customized manual configuration instead of accepting automatic configuration, use the methods described below in *Manual Configuration*.

1.1.6. Manual Configuration

Neverfail SCOPE Data Collector Service can be configured manually to adjust the Neverfail SCOPE Data Collector Service parameters using the procedures below.

- If the server has no Internet access, use the SCOPE Configuration Tool to set the required parameters.
- If the server has Internet access and you do not wish to use the global settings, create a machine-specific .CFG file using the *SCOPE Configuration* page on the Neverfail Extranet, or clear the *Enable automatic SCOPE Data Collector update checking* check box on the **General** page of the SCOPE Configuration Tool.

1.1.7. Neverfail SCOPE Data Collector Service Parameters

The Neverfail SCOPE Data Collector Service uses values stored in the registry to control its operational parameters.

These values can be adjusted by using the SCOPE Configuration Tool and/or through automatic configuration. It is important to understand these parameters and the interaction between the SCOPE Configuration Tool and the automatic configuration feature.

Note

The parameters on the following pages are designed to work with the online analyzer. Always consult Neverfail Support before adjusting.

Parameter Name	Default Value	Description
AdditionalFilesForUpload		Additional files to be added to the auto-uploaded .CAB file
Current Version		Neverfail SCOPE Data Collector Service version
ForcedTimeStampStart		Timestamp data gathering started
ForcedTimeStampStop		Timestamp data gathering stopped
Gathering Percent	0x00000064(100)	Percent of data gathered
Last File ID		The ID of the last generated .cab file
Last Job Status	StoreData	Last job done
Last Upgrade	never	Last Neverfail SCOPE Data Collector Service autoupgrade timestamp
LastForcedFilename		Last Neverfail SCOPE Data Collector Service data file - used by log collector
24 Hour Data Location	C:\ProgramData\Neverfail-SCOPE\Data\24 Hour Data	The 24 Hour Data file location
Amf Folder	C:\Program Files\Neverfail\SCOPE\AMF	AMF plug-ins location
Archives Folder	C:\ProgramData\Neverfail-SCOPE\Data\Archives	The archives location
Bin Folder	C:\Program Files\Neverfail\SCOPE	Location of Neverfail SCOPE Data Collector Service binaries
Candidate Location	C:\ProgramData\Neverfail-SCOPE\Data\Candidate For Upload	Location of the file to be uploaded
DB Root Path	C:\Document and Settings\All Users\ApplicationData\Neverfail-SCOPE\Data\WebServiceDB	The location where Neverfail SCOPE Data Collector Service stores data
DoOffice Location	C:\ProgramData\Neverfail-SCOPE\Data\DoOffice	The DoOffice measurement location
Log Dir	C:\ProgramData\Neverfail-SCOPE\Data\Debug	The logs directory
Output File Path	C:\ProgramData\Neverfail-SCOPE\Data\OutputFiles	The midnight files location

Parameter Name	Default Value	Description
Performance Output Path	C:\ProgramData\Neverfail-SCOPE\Data\Perf	The performance and history files location
Root Folder	C:\ProgramData\Neverfail-SCOPE\Data	The root folder of all the data subfolders
Auto Logs Cleanup (Days)	90	Timeout for log files
Bandwidth IPs		Selected bandwidth IP addresses
ManagedMemoryThreshold (MB)	1024	Management threshold in MB - if reached, Neverfail SCOPE Data Collector Service service is restarted
PrivateMemoryThreshold(MB)	1024	Management threshold in MB - if reached, Neverfail SCOPE Data Collector Service service is restarted
Proxy Encrypted Method	PlainText	How the proxy credentials should be encrypted
Proxy Password		Password used for the proxy server
Proxy Server		The proxy servers IP address
Proxy UserName		Username used for the proxy server
Reference GMT	0	Used to generate random upload time
Time window	5	Used by upload time randomization to randomly select a time in GMT+0 from 24:00 to 05:00 (Windows time value)
Active Server	True	Server Role
AMF Periodic Rules (Minutes)	15	Used to trigger AMF tasks
Config Frequency (Hours)	24	Static data gathering
Config Time		Set to an hour when the static data should be gathered
Download Frequency	01,00	The frequency at which updates will be downloaded (See note below)
Download Time	02:25	The time at which updates will be downloaded in a 24 Hour format
Download URL	scopeftp.neverfailgroup.com	The URL where program updates are located
Identity	PrimaryServer	Neverfail SCOPE Data Collector Service's Identity
Max Walk Time (Minutes)	60	Timeout for parsing the file system - Shares
ModelType	True	The type of Cluster (pair or trio)
Performance Interval (Minutes)	15	The frequency for collecting performance data (5, 10, 15, 30)
Primary's IP	<IP_address_of_primary_server>	The IP address of the Primary server (blank on the Primary server)
Secondary's IP	<IP_address_of_secondary_server>	The IP address of the Secondary server (blank on the Secondary server)
Socket Bandwidth Port	61000	The port on which the service will listen for connections from remote management utilities [this can be customized if needed]
Socket Forward Port	62000	The port used to send and receive remote data

Parameter Name	Default Value	Description
Tertiary's IP	<IP_address_of_tertiary_server>	The IP address of the Tertiary server (blank on the Tertiary server)
Upload Frequency	07,00	The frequency which data will be uploaded at (See note below)
Upload Time	01:20	Scheduled time of upload
Upload URL	scopeftp.neverfailgroup.com	The URL of the FTP server to upload data to
AMFDisableInfoLogging	True	Disable AMF message logging from Neverfail SCOPE Data Collector Service
AMFJobDisable	False	Disable the AMF from Neverfail SCOPE Data Collector Service
Auto Update Enabled	1	Enable updating of local configuration and binaries
Auto Upload Enabled	1	Enable automatic upload of data
CheckForSplitBrain	True	Checks if two Neverfail servers in a Cluster are both active
CheckMemoryConsumption	True	Checks memory consumption and if they exceed the threshold, Neverfail SCOPE Data Collector Service service is restarted
EnableFtpSSL	False	Send ftp data using secured sockets (SSL)
EnableManagementServices	False	Operates as a management server for SLM
EnableOutputCompressing	True	Compress the midnight files that are older than 1 month to a <monthName> <year> .cab file
EnableServerDataHistory	False	Keep the remote servers data in case one goes down for 1 day and append it to the Neverfail SCOPE Data Collector Service data file
LimitEventsTo24h	False	Limit events to 24 hours
Randomized	1	Randomize upload time/download time
Remote Management Enabled	1	Enable remote management
Upgrade On Server Activation	1	Perform Neverfail SCOPE Data Collector Service configuration and binaries auto-update in case the server becomes active
Veto HB Settings Constraint	0	Disregard the 'Heartbeat must be stopped to measure bandwidth' constraint
Veto SCOPE Bandwidth All	False	Measure bandwidth to all given IP addresses
Veto SCOPE Bandwidth IPs	0	Measure bandwidth only to IPs from the same Cluster (set in Neverfail Engine and Neverfail SCOPE Data Collector Service)

Note

The performance frequency is currently locked to 900 seconds (15 minutes) in order to maintain compatibility with the analyzer. Changes to this value will be ignored.

The upload and download frequencies are specified as two pairs of digits separated by a comma. Such as 01,00 or 00,08. The first pair designates a period in number of days between uploads/downloads, the second pair specifies the day of the month on which uploads/downloads should take place. Only one of these pairs of digits should be specified and the other must be 00. For example, 07,00 means every 7 days, 00,07 would mean on the 7th of every month.

1.1.8. Configure Bandwidth Measurement

Before you begin

To calculate the bandwidth available between two servers, you must install Neverfail SCOPE Data Collector Service on both servers.

About this task

Neverfail SCOPE Data Collector Service can measure the bandwidth between servers in the cluster but must be configured prior to initiating the measurement.

Procedure

1. Configure one server as the Primary server and the other as the Secondary server.
2. Connect the two network cards to one another in the same way you propose to configure the dedicated channel link between your Neverfail server pair. This connection may be a dedicated crossover cable, or it may be set up over a LAN or WAN.
3. Configure the two network cards with appropriate static IP addresses to allow network traffic between them. You should test the link before running Neverfail SCOPE Data Collector Service.
4. On the Primary server, configure the correct IP address for the Secondary server, and on the Secondary server, configure the correct IP address for the Primary server.
5. On the Primary server: Start the SCOPE Configuration Tool application by navigating to **Start > All Programs > Neverfail > SCOPE > SCOPE Configuration Tool**.
6. Select the **General** tab.
7. Set the server role to active by selecting the *This machine is the active server in the pair* check box.
8. Enter the IP Address of the Secondary server in the *Remote Configuration* pane.
9. Save and exit the SCOPE Configuration Tool.
10. On the Remote (Secondary) server: Start the SCOPE Configuration Tool application by navigating to **Start > All Programs > Neverfail > SCOPE > SCOPE Configuration Tool**.
11. Select the **General** tab.
12. Select the Secondary server role by clearing (un-checking) the *This machine is the active server in the pair* check box.
13. Save and exit the SCOPE Configuration Tool.

1.2. Neverfail SCOPE Data Collector Service Network Ports

The Neverfail SCOPE Data Collector Service service uses the network ports listed in the following table. For full operation of Neverfail SCOPE Data Collector Service, these ports must be opened on any firewalls.

Ports	Default Use
62000	Inter-process communications between the Primary and Secondary servers and remote management. This port is customizable.
61000	Bandwidth calculations between the Primary and Secondary servers. This port is customizable.

1.3. Daylight Savings Time

Neverfail SCOPE Data Collector Service does not use an internal time but instead uses the server's clock to operate.

Since Neverfail SCOPE Data Collector Service uses the server's clock, manually adjusting the server time may result in longer or shorter periods between data capture.

For example, if Neverfail SCOPE Data Collector Service is configured to gather data at 17:15, but an administrator or automated process resets the server's clock at 17:02 to 16:02 (-1 hour), Neverfail SCOPE Data Collector Service still gathers the data at 17:15 by the server's clock. In the performance data, the timestamp will contain 17:15 resulting in 25 hours worth of data.

2. Neverfail SCOPE Analysis Reports

2.1. Neverfail SCOPE Reports

The Neverfail SCOPE Report provides the results of a detailed interrogation of the server environment.

Neverfail SCOPE Analysis Report

As stated previously, Neverfail SCOPE is a combination of both software and process, designed to ensure a stable server environment and a successful Neverfail Engine implementation. The information collected must be uploaded to the Neverfail Extranet for analysis. If Neverfail SCOPE is configured for automatic upload, this task is accomplished automatically and requires no user input. If Neverfail SCOPE is not configured for automatic upload, you must upload the collected information manually

Once uploaded, the raw data file is immediately analyzed and a Neverfail SCOPE Report is available for viewing using a standard web browser

The Neverfail SCOPE Report provides information about:

- Windows version, including Service Packs and Hotfixes
- System memory (RAM)

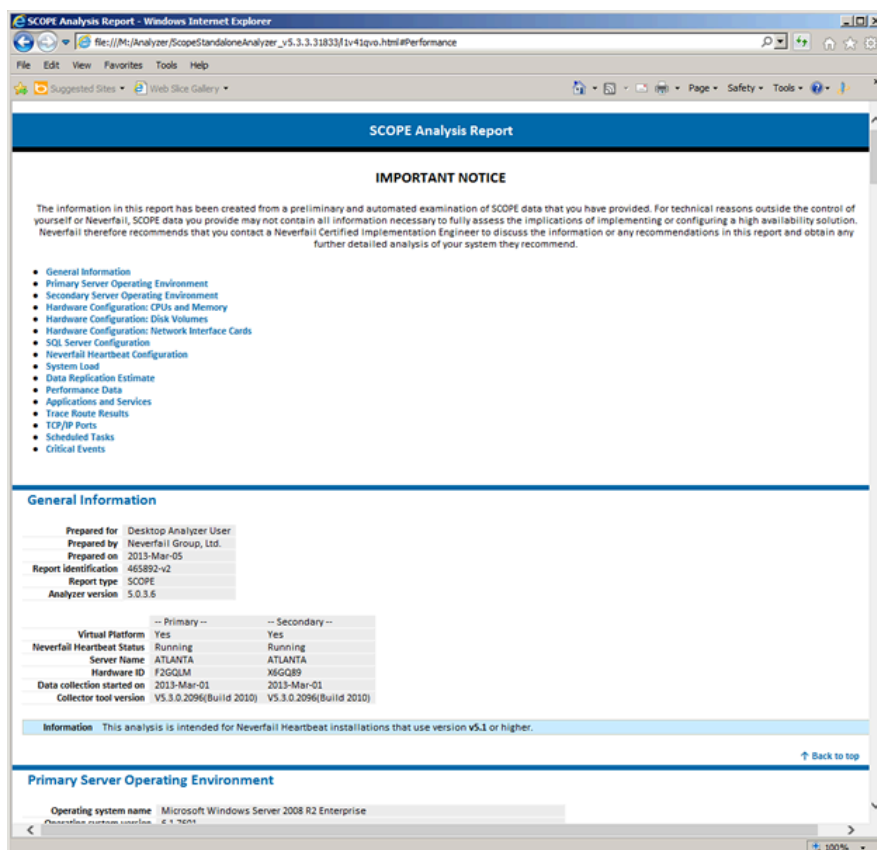
- Disk size, type, partition structure, and available space
- Shared folders
- Windows services
- Third-party application services
- Optional available bandwidth measurement and replication bandwidth estimate if Neverfail SCOPE runs for at least 24 hours
- A detailed performance report
- Recommended changes (in red)

Note

The required bandwidth estimate is based upon an actual network measurement using server disk activity. You can use the estimate as a guide to determine bandwidth requirements for the dedicated Neverfail Channel link between servers.

The Neverfail SCOPE Analysis Report provides an overview of the analyzed criteria and identifies any areas of the current environment that are likely to pose problems when implementing Neverfail Engine. Problems that must be resolved before installing Neverfail Engine are highlighted in red for easy identification. This report should be reviewed in its entirety to ensure that the current server environment is adequate for a successful Neverfail Engine installation.

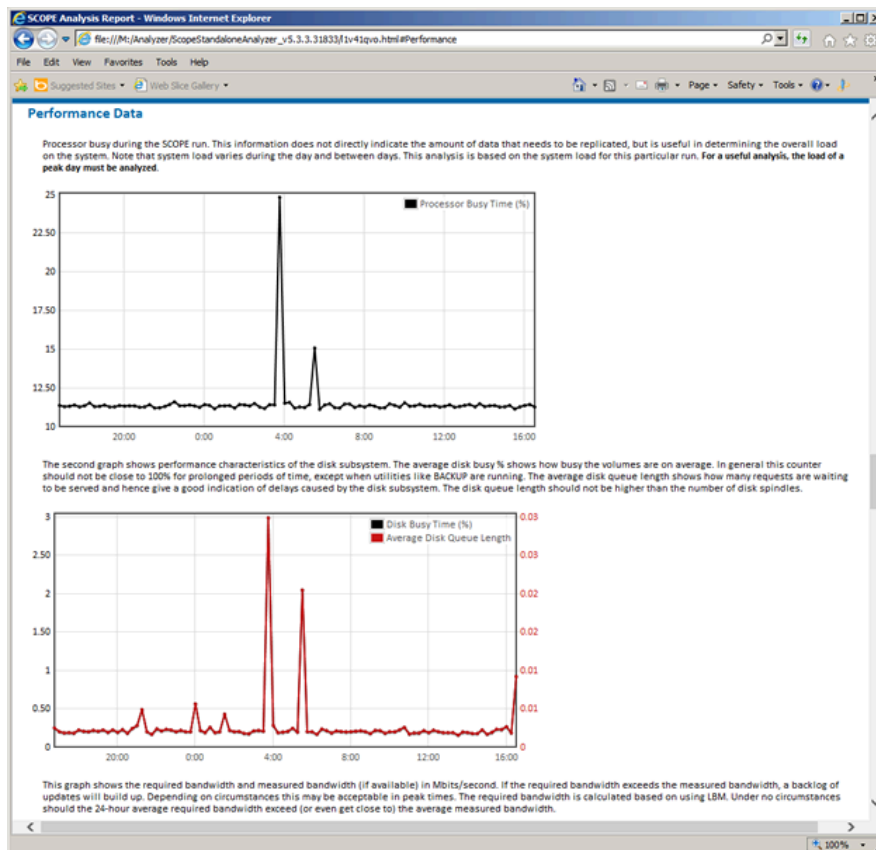
Figure 10-5. SCOPE Analysis Report



2.2. Neverfail SCOPE Graphs

The Neverfail SCOPE Analysis Report provides an overview of the analyzed criteria and identifies any areas of the current environment that are likely to pose problems when implementing Neverfail Engine. Problems that must be resolved before installing Neverfail Engine are highlighted in red for easy identification. This report should be reviewed in its entirety to ensure that the current server environment is adequate for a successful Neverfail Engine installation.

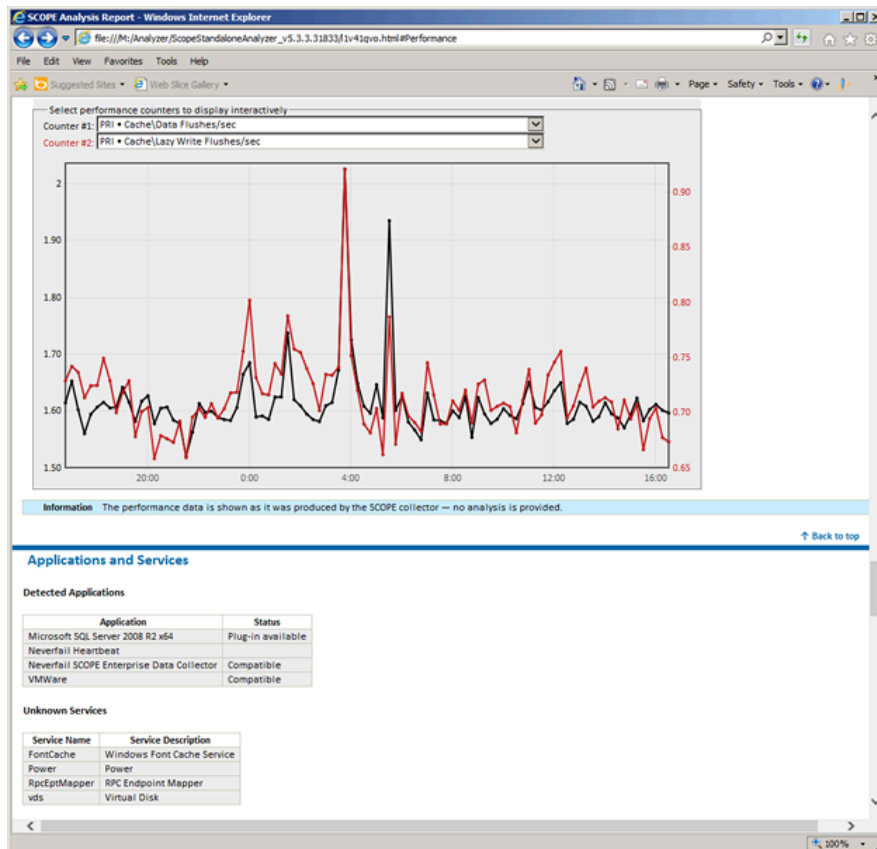
Figure 10-6. Neverfail SCOPE Graphs



2.3. Neverfail SCOPE Performance Counters

The Neverfail SCOPE Performance Counter graph provides for selection of a variety of counters and permits comparison between servers. Placing the cursor over a data point displays the exact value of the counter.

Figure 10-7. Neverfail SCOPE Performance Counters



Glossary

Glossary

Active

The functional state or role of a server when it is visible to clients through the network, running protected applications, and servicing client requests.

Alert

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

Active Directory (AD)

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

Active–Passive

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

Advanced Configuration and Power Interface (ACPI)

A specification that dictates how the operating system can interact with the hardware especially where power saving schemes are used. The Primary, Secondary, and Tertiary servers must have identical ACPI compliance.

Asynchronous

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

Basic Input/Output System (BIOS)

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

Cached Credentials

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

Channel Drop

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

Channel NIC (Network Interface Card)

A dedicated NIC used by the Neverfail Channel.

Checked

The status reported for user account credential (username/password) validation.

Cloned Servers

Servers that have identical configuration settings, names, applications, Security Identifiers (SIDs) and IP addresses, following the installation of Neverfail Engine.

Cloning Process

The Neverfail Continuity Engine process whereby all installed programs, configuration settings, and the machine name, Security Identifier (SID), and IP addresses are copied to another server.

Cluster

A generic term for a Neverfail Engine Pair or Trio and the set of machines (physical or virtual) involved in supporting a single protected server. A Neverfail Engine Cluster can include the machines used in a VMware or Microsoft cluster.

Connection

Also referred to as Cluster Connection. Allows the Engine Management Service to communicate with a Neverfail Engine Cluster, either on the same machine or remotely.

Crossover Cable

A network cable that crosses the transmit and receive lines.

Data Replication

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

Data Rollback Module

A Neverfail Continuity Engine module that allows administrators to rollback the entire state of a protected application, including files and registry settings, to an earlier point-in-time. Typically used after some form of data loss or corruption.

Degraded

The status reported for an application or service that has experienced an issue that triggered a Rule.

Device Driver

A program that controls a hardware device and links it to the operating system.

Disaster Recovery (DR)

A term indicating how you maintain and recover data with Neverfail Engine in event of a disaster such as a hurricane or fire. DR protection can be achieved by placing the Secondary server at an off-site facility, and replicating the data through a WAN link.

DNS (Domain Name System) Server

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

Domain

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

Domain Controller (DC)

The server responsible for maintaining privileges to domain resources; sometimes called AD controller in Windows 2003 and above domains.

Dualled

A way to provide higher reliability by dedicating more than one NIC for the Neverfail Channel on each server.

Failover

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

Full System Check (FSC)

The internal process automatically started at the initial connection or manually triggered through the Manage Server GUI to perform verification on the files and registry keys and then synchronize the differences.

Fully Qualified Domain Name (FQDN)

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

Global Catalog

A global catalog is a domain controller that stores a copy of all Active Directory objects in a forest. The global catalog stores a full copy of all objects in the directory for its host domain and a partial copy of all objects for all other domains in the forest.

Graceful (Clean) Shutdown

A shutdown of Neverfail Engine based upon completion of replication by use of the Engine Management Service, resulting in no data loss.

Group

An arbitrary collection of Neverfail Engine Clusters used for organization.

Hardware Agnostic

A key Neverfail Continuity Engine feature allowing for the use of servers with different manufacturers, models, and processing power in a single Neverfail Engine Cluster.

Heartbeat

The packet of information issued by the passive server across the channel, which the active server responds to indicating its presence.

High Availability (HA)

Keeping users seamlessly connected to their applications regardless of the nature of a failure. LAN environments are ideally suited for HA.

Hotfix

A single, cumulative package that includes one or more files that are used to address a problem in a product.

Identity

The position of a given server in the Neverfail Continuity Engine Cluster: Primary, Secondary, or Tertiary.

Install Clone

The installation technique used by Neverfail Continuity Engine to create a replica of the Primary server using NTBackup or Wbadmin and to restore the replica to the Secondary and/or Tertiary servers.

Low Bandwidth Module (LBM)

A Neverfail Continuity Engine module that compresses and optimizes data replicated between servers over a WAN connection. This delivers maximum data throughput and improves application response time on congested WAN links.

Machine Name

The Windows or NETBIOS name of a computer.

Management IP Address

An additionally assigned unfiltered IP address in a different subnet than the Public or Neverfail Channel IP addresses used for server management purposes only.

Many-to-One

The ability of one physical server (hosting more than one virtual server) to protect multiple physical servers.

Network Monitoring

Monitoring the ability of the active server to communicate with the rest of the network by polling defined nodes across the network at regular intervals.

Neverfail Channel

The IP communications link used by the Neverfail system for the heartbeat and replication traffic.

Neverfail Continuity Engine

The core replication and system monitoring component of the Neverfail solution.

Neverfail Extranet

The Neverfail web site dedicated to supporting partners and customers by providing technical information, software updates, and license key generation.

Neverfail Engine Packet Filter

The network component, installed on all servers, that controls network visibility.

Neverfail License Key

The key obtained from the Neverfail extranet that allows the use of components in the Neverfail suite; entered via the License wizard of the Engine Management Service User Interface, or through the Configure Server Wizard.

Neverfail Pair

Describes the coupling of the Primary and Secondary server in a Neverfail solution.

Neverfail Plug-ins

Optional modules installed into a Neverfail Continuity Engine server to provide additional protection for specific applications.

Neverfail SCOPE

The umbrella name for the Neverfail process and tools used to verify the production servers health and suitability for the implementation of a Neverfail solution.

Neverfail SCOPE Report

A report provided upon the completion of the Neverfail SCOPE process that provides information about the server, system environment, and bandwidth.

Neverfail Switchover/Failover Process

A process unique to Neverfail in which the passive server gracefully (switchover) or unexpectedly (failover) assumes the role of the active server providing application services to connected clients.

Pair

See Neverfail Continuity Engine Pair above.

Passive

The functional state or role of a server when it is not delivering service to clients and is hidden from the rest of the network.

Pathping

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

Plug-and-Play (PnP)

A standard for peripheral expansion on a PC. On starting the computer, PnP automatically configures the necessary IRQ, DMA and I/O address settings for the attached peripheral devices.

Plug-in

An application specific module that adds Neverfail Continuity Engine protection for the specific application.

Pre-Clone

An installation technique whereby the user creates an exact replica of the Primary server using VMware vCenter Converter or other 3rd party utility prior to the initiation of installation and uses the replica as a Secondary and or Tertiary server.

Pre-Installation Checks

A set of system and environmental checks performed as a prerequisite to the installation of Neverfail Engine.

Primary

An identity assigned to a server during the Neverfail Engine installation process that normally does not change during the life of the server and usually represents the production server prior to installation of Neverfail Engine.

Protected Application

An application protected by the Neverfail Continuity Engine solution.

Public IP Address

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

Public Network

The network used by clients to connect to server applications protected by Neverfail Continuity Engine.

Public NIC

The network card which hosts the Public IP address.

Quality of Service (QoS)

An effort to provide different prioritization levels for different types of traffic over a network. For example, Neverfail Engine data replication may have a higher priority than ICMP traffic, as the consequences of interrupting data replication are more obvious than slowing down ICMP traffic.

Receive Queue

The staging area on a passive server used to store changes received from another server in the replication chain before they are applied to the disk/registry on the passive server.

Remote Desktop Protocol (RDP)

A multi-channel protocol that allows a user to connect to a computer running Microsoft Terminal Services.

Replication

The generic term given to the process of intercepting changes to data files and registry keys on the active server, transporting the changed data across the channel, and applying them to the passive server(s) so the servers are maintained in a synchronized state.

Role

The functional state of a server in the Neverfail Continuity Engine Cluster: active or passive.

Rule

A set of actions performed by Neverfail Continuity Engine when defined conditions are met.

Secondary

An identity assigned to a server during the Neverfail Engine installation process that normally does not change during the life of the server and usually represents the standby server prior to installation of Neverfail Engine.

Security Identifier (SID)

A unique alphanumeric character string that identifies each operating system and each user in a network of Windows 2008/2012 systems.

Send Queue

The staging area of the active server used to store intercepted data changes before being transported across Neverfail Channel to a passive server in the replication chain.

Server Monitoring

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

Shared Nothing

A key feature of Neverfail Continuity Engine in which no hardware is shared between the Primary or Secondary servers. This prevents a single point of failure.

SMTP

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

SNMP

Simple Network Management Protocol (SNMP) is an Internet-standard protocol for managing devices on IP networks.

Split-Brain Avoidance

A unique feature of Neverfail Continuity Engine that prevents a scenario in which Primary and Secondary servers attempt to become active at the same time leading to an active-active rather than an active-passive model.

Split-Brain Syndrome

A situation in which more than one server in a Neverfail Engine Cluster are operating in the active mode and attempting to service clients, resulting in the independent application of different data updates to each server.

Subnet

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

Storage Area Network (SAN)

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

Switchover

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

Synchronize

The internal process of transporting 64KB blocks of changed files or registry key data, through the Neverfail Channel, from the active server to the passive server to ensure the data on the passive server is a mirror image of the protected data on the active server.

System Center Operations Manager (SCOM)

System Center Operations Manager is a cross-platform data center management server for operating systems and hypervisors.

System State

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

Task

An action performed by Neverfail Engine when defined conditions are met.

Tertiary

An identity assigned to a server during the Neverfail Continuity Engine installation process that normally does not change during the life of the server and usually represents the disaster recovery server prior to installation of Neverfail Continuity Engine.

Time-To-Live (TTL)

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

Traceroute

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

Trio

A Neverfail cluster comprising three servers, a Primary, Secondary and Tertiary, in order to provide High Availability and Disaster Recovery.

Ungraceful (Unclean) Shutdown

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

Unprotected Application

An application that is not monitored nor its data replicated by Neverfail Continuity Engine.

Virtual Private Network (VPN)

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

Windows Management Instrumentation (WMI)

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