

Installation Guide

For

Neverfail Continuity Engine v8.1

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 © 2017 Neverfail, LLC. All rights reserved.



Contents

Preface: About This Book.....	iv
Chapter 1: Introduction.....	6
Neverfail Continuity Engine Concepts.....	6
Communications.....	9
Neverfail Engine Switchover and Failover Processes.....	10
Chapter 2: Implementation.....	11
Neverfail Continuity Engine Implementation.....	11
Environmental Prerequisites.....	11
Supported Environments.....	12
Unsupported Environments.....	12
Minimal VMware Permissions Requirements:.....	12
Pre-Install Requirements.....	13
Server Deployment Architecture Options.....	16
Virtual-to-Virtual.....	16
Physical-to-Virtual.....	16
Physical-to-Physical.....	17
Cloning Technology Options.....	17
Application Component Options.....	18
Networking Configuration.....	18
Local Area Network (LAN).....	19
Wide Area Network (WAN).....	19
Network Interface Card (NIC) Configuration.....	20
Firewall Configuration Requirements.....	21
Anti-Malware Recommendations.....	23
Chapter 3: Installing Neverfail Continuity Engine	24
Installing Neverfail Continuity Engine	24
Deploying Neverfail Engine on the Primary Server.....	25
Automated Deployment of Stand-by Servers with Automatic Cloning.....	26
Semi-Automatic Deployment of Stand-by Servers Leveraging Manual Cloning.....	28
Using the Engine Management Service User Interface.....	30
Configure Connection to VMware vCenter Server.....	30
Configure VMware vCenter Converter.....	31
Protected Servers.....	33
Management.....	33
Summary.....	63
Status.....	63
Events.....	65
Services.....	66
Data.....	70
Shadows.....	73

Tasks.....	80
Rules.....	83
Settings.....	85
Actions.....	92
Post Installation Configuration.....	95
Configure the VmAdapter Plug-in.....	95
Adding an Additional Network Interface Card.....	96
Appendix A: Installation Verification Testing.....	98
Testing a Neverfail Engine Pair.....	98
Exercise 1 - Auto-switchover.....	98
Exercise 2 - Data Verification.....	100
Exercise 3 - Switchover.....	101
Testing a Neverfail Engine Trio.....	101
Exercise 1 - Auto-switchover.....	102
Exercise 2 - Managed Switchover.....	103
Exercise 3 - Data Verification.....	105
Glossary.....	107

About This Book

The Installation Guide provides information about installing Neverfail Continuity Engine, including implementation in a Local Area Network (LAN) and/or Wide Area Network (WAN). This book provides an overview of installation procedures and guidance for the configuration of Neverfail Continuity Engine when the Secondary and Tertiary servers are virtual.

Intended Audience

This guide assumes the reader has a working knowledge of networks including the configuration of TCP/IP protocols and domain administration, notably in Active Directory and DNS.

Overview of Content

This guide is designed to provide guidance on the installation and configuration of Neverfail Continuity 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 — *Introduction* presents an overview of Neverfail Continuity Engine concepts including the Switchover and Failover processes.
- Chapter 2 — *Implementation* discusses environmental prerequisites and pre-install requirements for installation, options for server architecture, application components, and network configurations. It also gives guidance on anti-malware solutions, and provides a convenient summary of supported configurations as you perform the installation.
- Chapter 3 — *Installing* describes the installation process, guides you through installation on the Primary, Secondary, and Tertiary (if deployed) servers, and through post-installation configuration.
- Appendix A — *Installation Verification* provides a quick, simple procedure to verify that Neverfail Continuity Engine is properly installed and initially configured.

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

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

Technical Support and Education Resources

The following sections describe technical support resources available to you. To access the current version of this book and other related books, go to <http://www.neverfail.com/services-and-support/>

Online and Telephone Support

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

Support Offerings

To find out how Neverfail Support offerings can help meet your business needs, go to <http://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 Continuity Engine servers. To access information about education classes, certification programs, and consulting services, go to <http://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 http://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.

Chapter 1

Introduction

Neverfail Continuity Engine is a Windows based service specifically designed to provide High Availability and/or Disaster Recovery for server configurations in one solution without any specialized hardware.

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.

Topics:

- [*Neverfail Continuity Engine Concepts*](#)
- [*Communications*](#)
- [*Neverfail Engine Switchover and Failover Processes*](#)

Neverfail Continuity Engine Concepts

Overview

Neverfail Continuity Engine consists of the Engine Management Service that is used to deploy and manage the Neverfail Engine nodes that provide for application-aware continuous availability used for protecting critical business systems. The Engine Management Service can be installed on vCenter Server or another Windows server with access to a remote instance of vCenter Server and is accessible via common web browsers.

Using the Engine Management Service User Interface (UI), 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.

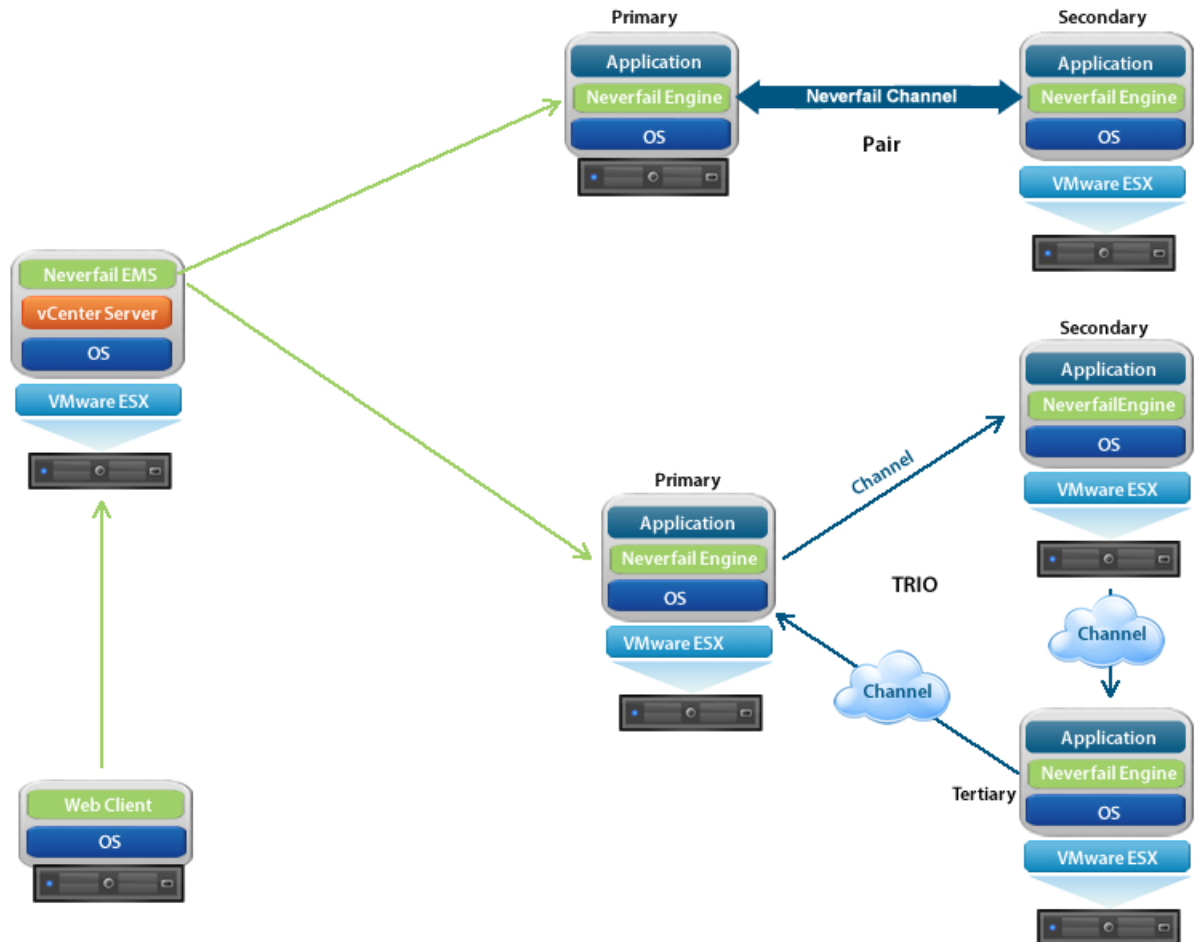


Figure 1: Deployment Architecture

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 a Cluster of 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 an **Identity** (*Primary/Secondary/Tertiary*) and a **Role** (*Active/Passive*). Identity is used to describe the physical instance of the server while the role is used to describe what the server is doing. When the identity is assigned to a server it normally will not change over the life of the server whereas the role of the server is subject to change as a result of the operations the server is performing. When Neverfail 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).

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

In its simplest form, Neverfail Engine operates as a Neverfail Engine Pair with one server performing an active role (normally the Primary server) while the other server performs a passive role (normally the Secondary server). The server in the active role provides application services to users and serves

as the source for replication while the server in the passive role serves as the standby server and target for replicated data. This configuration supports replication of data between the active and passive server over the Neverfail Channel.

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

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

Architecture

Neverfail Engine software is installed on a [Primary](#) (production) server, a [Secondary](#) (ready-standby) server, and optionally, a [Tertiary](#) (also a ready-standby) server. These names refer to the identity of the servers and never change throughout the life of the server (except in the special case described above).

Note: *In this document, the term “Cluster” refers to a Neverfail Engine Cluster. Refer to the [Glossary](#) for more information about Neverfail Engine Clusters.*

Depending on the network environment, Neverfail Continuity Engine can be deployed in a Local Area Network (LAN) for High Availability and/or Wide Area Network (WAN) for Disaster Recovery, providing the flexibility necessary to address most network environments.

When deployed, one of the servers performs the [Role](#) of the [Active](#) server that is visible on the Public network while the other is [Passive](#) and hidden from the Public network but remains as a ready-standby server. The Secondary server has the same domain name, uses the same file and data structure, same Public network address (in a LAN), and can run all the same applications and services as the Primary server. Only one server can display the Public IP address and be visible on the Public network at any given time. Neverfail Engine software is symmetrical in almost all respects, and either the Primary server, Secondary server, or Tertiary server (if applicable) can take the active role and provide protected applications to the user.

Protection Levels

Neverfail Continuity Engine provides the following protection levels:

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

Neverfail Continuity Engine provides all five protection levels continuously, ensuring all facets of the user environment are maintained at all times, and that the Public network continues to operate through as many failure scenarios as possible.

Communications

Neverfail Continuity Engine communications consist of two crucial components, the Neverfail Channel and the Public network.

To accommodate communications requirements, Neverfail Engine can be configured with either a single NIC configured with both the Public IP address and the Neverfail Channel IP address on the same NIC or multiple NICs. Separate NICs can be dedicated for the Public and Channel IP addresses, but this is not a requirement.

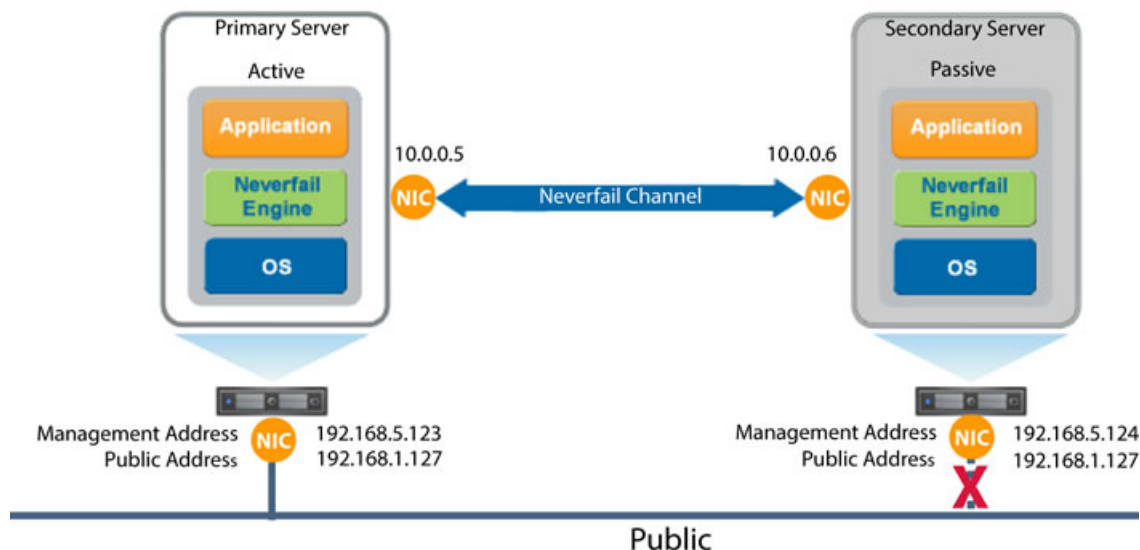


Figure 2: Communications Between Primary and Secondary Servers

Neverfail Channel

The first component is the Neverfail Channel which provides communications between the active and passive servers. The Neverfail Channel is used for control and data transfer from the active server to the passive server and for monitoring of the active server's status by the passive server.

The Channel IP addresses can be in the same or a different subnet as the Public IP address. NetBIOS will be filtered for the Neverfail Channel on the active and passive servers to prevent server name conflicts.

The NICs that support connectivity across the Neverfail Channel can be standard 10/100/1000 Base-T Ethernet cards providing a throughput of up to 1000 Mbits per second across standard Cat-5 cabling or virtual NICs configured on a virtual machine.

When configured for a WAN deployment, if the Channel IP addresses are in the same subnet as the Public IP Address, then they will be routed via the default gateway in a WAN deployment. Alternatively you can configure the Neverfail Channel to use static routes over switches and routers to maintain continuous communications independent from corporate or public traffic.

Public Network

The second component is the Public network used by clients to connect to the active server. The Public network provides access to the Public IP address used by clients to connect to the active server.

The Public IP address is a static IP address that is only available on the currently active server and is the IP address a client uses to connect to the active server. It must be configured as a static IP address, that is, not DHCP (Dynamic Host Configuration Protocol) enabled. In the figure above, the IP address is configured as 192.168.1.127. The Public IP address is common to the active and passive servers in a LAN and is always available on the currently active server in the cluster. In the event of a switchover or failover, the Public IP address is removed from the previously active server and is then available on the new active server. When configured, a Management IP address will provide access to a server regardless of the role of the server.

Management IP Address

After installation, all servers in the cluster can be configured with separate Management IP addresses that allow access to the server when the server is in the passive role. The Management IP address is a static IP address in a different subnet than the Public IP address or Neverfail Channel IP address and is always available for administrators to access the server.

Neverfail Engine Switchover and Failover Processes

Neverfail Engine uses four different procedures – managed switchover, automatic switchover, automatic failover, and managed failover – to change the role of the active and passive servers depending on the status of the active server.

- *Managed Switchover* – To perform a Managed Switchover, navigate to the *Actions* drop-down of the Engine Management Service UI and click to make one of the stand-by servers active to initiate a managed switchover or you can click **Make Active** on the Neverfail Advanced Management Client *Server: Summary* page. When a managed switchover is triggered, the running of protected applications is transferred from the active machine to the passive machine in the server pair. The server roles are reversed.
- *Automatic Switchover* – Automatic switchover (auto-switchover) is similar to failover (discussed in the next section) but is triggered automatically when system monitoring detects failure of a protected application.
- *Automatic Failover* – Automatic failover is similar to automatic switchover (discussed above) but is triggered when the passive server detects that the active server is no longer running properly and assumes the role of the active server.
- *Managed Failover* – Managed failover is similar to automatic failover in that the passive server automatically determines that the active server has failed and can warn the system administrator about the failure, but no failover actually occurs until the system administrator manually triggers this operation (the default configuration in a DR environment).

Chapter 2

Implementation

This chapter discusses the deployment options and prerequisites to successfully implement Neverfail Continuity Engine and provides a step-by-step process to assist in selecting options required for installation.

Topics:

- [*Neverfail Continuity Engine Implementation*](#)
- [*Environmental Prerequisites*](#)
- [*Minimal VMware Permissions Requirements:*](#)
- [*Pre-Install Requirements*](#)
- [*Server Deployment Architecture Options*](#)
- [*Cloning Technology Options*](#)
- [*Application Component Options*](#)
- [*Networking Configuration*](#)
- [*Firewall Configuration Requirements*](#)
- [*Anti-Malware Recommendations*](#)

Neverfail Continuity Engine Implementation

Neverfail Continuity Engine is a versatile solution that provides multiple configurations to suit user requirements. It can be deployed in a LAN for high availability and/or across a WAN to provide disaster recovery.

During the installation process, Engine Management Service performs a variety of checks to ensure the server meets the minimum requirements for a successful installation. A critical stop or warning message appears if the server fails a check. You must resolve critical stops before you can proceed with setup. Prior to installing Neverfail Continuity Engine, select the deployment options you intend to use. The installation process will prompt you to select options throughout the procedure to create the configuration you want.

Environmental Prerequisites

Neverfail Continuity Engine supports the following environments listed below.

Supported Environments

- Neverfail Continuity Engine is supported on the following versions of Windows Server
 - Windows Server 2008 R2 Standard/Enterprise/Datacenter
 - Windows Server 2012 Standard/Datacenter
 - Windows Server 2012 R2 Standard/Datacenter
 - Windows Server 2016 Standard/Datacenter

Unsupported Environments

- Neverfail Continuity Engine is not supported across the following:
 - A server where Engine Management Service is already running
 - On a server deployed as a [Domain Controller \(DC\)](#)
 - On a server deployed as a [Global Catalog](#)
 - On a server deployed as a [DNS \(Domain Name System\) Server](#)
 - On an IA-64 Itanium Platform

Minimal VMware Permissions Requirements:

Procedure

To create a Neverfail Continuity Engine install user:

1. Using the VMware vSphere Client, log into vCenter Server as an Administrator.
2. Navigate to **Home > Roles**.
3. Select the *Read-only* role.
4. Right-click the role and click **Clone**.
5. Rename the new role. For example, Neverfail Continuity Engine.
6. Right-click the newly cloned role and select *Edit Role*.
7. Add the following privileges:

Note: *The below listed permissions are the minimal required permissions to perform an installation.*

- **Datastore > Allocate Space**
- **Datastore > Browse Datastore**
- **Extension**
- **Global > Log Event**
- **Network > Assign Network**
- **Resource > Assign Virtual Machine to Resource Pool**
- **Resource > Migrate powered off virtual machine**
- **Resource > Migrate powered on virtual machine**
- **Tasks**
- **Virtual Machine > Configuration**

- **Virtual Machine > Interaction > Configure CD Media**
 - **Virtual Machine > Interaction > Device Connection**
 - **Virtual Machine > Interaction > Power On**
 - **Virtual Machine > Interaction > Power Off**
 - **Virtual Machine > Interaction > Reset**
 - **Virtual Machine > Inventory**
 - **Virtual Machine > Provisioning**
 - **Virtual Machine > Snapshot Management**
8. Map the vCenter Server user account configured in Engine Management Client (EMS) to the newly created Neverfail Continuity Engine role, at the vCenter Server level.
- a) Select the top level for vCenter Server, then click the **Permissions** tab.
 - b) Right-click and select *Add Permission*.
 - c) Add the vCenter Server EMS user (if not already present) and assign the newly created Neverfail Continuity Engine role.

Note: You may need to bind the role at the host level (in Hosts and Cluster View) as well as the Datastore permissions tab level (in Datastores & Datastore Clusters).

Pre-Install Requirements

The following provides a listing of pre-requisites that must be addressed prior to attempting an installation of Neverfail Continuity Engine.

Server	Action	
Engine Management Service	Engine Management Service installation is supported on the following operating systems:	<input type="checkbox"/>
	<ul style="list-style-type: none"> • Microsoft Windows 2008 R2, 2012, 2012 R2, 2016 • Microsoft Windows Client Edition 7.0, 8.0, 8.1, and 10.0 	
	Note: Connectivity with VMware vCenter Server is NOT required for deployment of Neverfail Continuity Engine but is recommended for fully automated deployments.	
	vCenter Server Administrator level user credentials (equivalent with Administrator@vsphere.local) or a user configured with minimal permissions listed in the previous section. Where possible, we recommend vCenter Server Administrator level user credentials (equivalent with Administrator@vsphere).	<input type="checkbox"/>
	For P2V installation, VMware Converter 5.5 must be available and configured prior to attempting installation of the Primary server.	<input type="checkbox"/>
	Engine Management Service (EMS) supports most browsers used to connect to the EMS UI but requires that the latest version of Adobe Flash Player be installed.	<input type="checkbox"/>
Primary Server	Engine Management Service requires elevated permissions in order to be installed.	<input type="checkbox"/>
	Note: Engine Management Service will be configured to use NeverfailEngine log on account. This account is member of the local administrators group. Neverfail recommends changing the account's password after the EMS installation is completed (after the password is changed, Neverfail Engine Management Web Services service should be reconfigured to use the new password).	
	Neverfail Engine requires that Microsoft™ .Net Framework 4.0 or later be installed prior to installation.	<input type="checkbox"/>
	Neverfail Engine requires that the Windows Feature SMB 1.0/CIFS File Sharing Support be enabled prior to installing Neverfail Engine.	<input type="checkbox"/>

Server	Action	
	If the Primary server has a pending reboot, it must be resolved prior to the deployment of Neverfail Engine on to the server.	<input type="checkbox"/>
	Obtain and use local administrator permissions to perform Neverfail Engine installation.	<input type="checkbox"/>
	If UAC is enabled on the target server, you must use the built-in local Administrator account or, for domain member servers, you can alternatively use a Domain User account that is a member of the local Administrators group.	
	If UAC is not enabled, you may use any account with membership in the local Administrators group on the target server.	
	Note: <i>Neverfail Engine services are required to be run under the Local System account.</i>	
	The server to be protected by Neverfail Engine can NOT be any of the following:	<input type="checkbox"/>
	<ul style="list-style-type: none"> A server running Engine Management Service A server configured as a Domain Controller, Global Catalog, DHCP, or DNS 	
	Important: <i>These roles and services must be removed before proceeding with installation.</i>	
	The Primary server can be Virtual or Physical with the Secondary and Tertiary server (if deployed) as either Virtual or Physical as well.	<input type="checkbox"/>
	Verify that all services to be protected have all three <i>Recovery</i> settings configured to <i>Take no Action</i> .	<input type="checkbox"/>
	Verify no other critical business applications except those to be protected by Neverfail Engine are installed on the server.	<input type="checkbox"/>
	Verify that there is a minimum of 2GB of available RAM in addition to any other memory requirements for the Operating System or installed applications. 512MB of RAM must remain available to Neverfail Engine at all times.	<input type="checkbox"/>
Secondary Server	Verify that a minimum 2GB of free disk space is available on the drive where Neverfail Engine is to be installed.	<input type="checkbox"/>
	Note: <i>Although Neverfail Engine requires only 2GB of available disk space on the drive to receive the Neverfail Engine installation, once installed, the size of each Send and Receive queue is configured by default for 10GB. For Trio configurations the send and receive queues will by default require 20GB per server. You must ensure that sufficient disk space is available to accommodate the send and receive queues or modify the queue size configuration to prevent MaxDiskUsage errors.</i>	
	Apply the latest Microsoft security updates and set Windows Updates to <i>manual</i> .	<input type="checkbox"/>
	All applications that will be protected by Neverfail Engine must be installed and configured on the Primary server prior to installing Neverfail Engine.	<input type="checkbox"/>
	Verify that all services to be protected are running or set to <i>Automatic</i> prior to installation.	<input type="checkbox"/>
	Note: <i>During installation, protected services are set to manual to allow Neverfail Engine to start and stop services depending on the role of the server. The target state of the services is normally running on the active server and stopped on the passive.</i>	
	Register this connection's address in DNS must be disabled on all NICs on the target server.	<input type="checkbox"/>
	Note: <i>If deploying in a DR configuration, replace the existing DNS "A" record for the Public IP address with a static record and configure the TTL to 45 seconds. Otherwise, after installation, re-enable Register this connections's address in DNS.</i>	
	File and Printer Sharing must be enabled and allowed access through all firewalls on the Primary target server prior to deployment.	<input type="checkbox"/>
	Verify that the Server service is running prior to deployment to the target server.	<input type="checkbox"/>
	When installing in a P2V environment, the specifications of the Secondary Neverfail Engine virtual machine must match the Primary physical server as follows:	<input type="checkbox"/>
	<ul style="list-style-type: none"> Similar CPU Identical Memory Sufficient disk space to host VM disks to match the Primary server 	

Server	Action	
	The Secondary Neverfail Engine virtual machine must have sufficient priority in resource management settings so that other virtual machines do not impact its performance.	
IP Addressing	<p>IP Address requirements:</p> <p>Public:</p> <ul style="list-style-type: none"> 1 each Public IP address - Engine Management Service 1 each Public IP address - Primary Server 1 each Public IP address - Secondary Server (only when deployed for DR) <p>Note: When deployed for HA or as part of a trio, the Primary and Secondary server will share the Public IP address.</p> <p>1 each Public IP address - Tertiary Server (only when deployed in a trio)</p> <p>Channel:</p> <ul style="list-style-type: none"> 1 each Channel IP address - per server when deployed in a pair 2 each Channel IP addresses - per server when deployed in a trio 	<input type="checkbox"/>
LAN	<p>When deployed in a LAN environment, Neverfail Engine requires that both servers use the same Public IP address. Each server also requires a <u>unique</u> Neverfail Channel IP address.</p> <p>Note: After deployment, on the Public NIC, go to the Network Properties for TCP/IP4 and under Advanced Properties, select Register this connection's address in DNS for the Public NIC.</p>	<input type="checkbox"/>
WAN	<p>When deployed in a WAN environment, persistent static routing configured for the channel connection(s) where routing is required.</p> <p>Note: This requirement can be avoided if the channel IP addresses are in the same subnet as the Public IP address in which case the default gateway can be used for routing.</p>	<input type="checkbox"/>
	At least one Domain Controller at the Disaster Recovery (DR) site.	<input type="checkbox"/>
	<ul style="list-style-type: none"> If the Primary and DR site uses the same subnet: <ul style="list-style-type: none"> During installation, follow the steps for a LAN or vLAN on the same subnet. Both the Primary and Secondary servers in the pair use the same Public IP address. If the Primary and DR site use different subnets: <ul style="list-style-type: none"> During installation, follow the steps for a WAN. The Primary and Secondary servers in the Neverfail Engine pair require a separate Public IP address and an Neverfail Channel IP address. Provide a user account with rights to update DNS using the <code>DNSUpdate.exe</code> utility provided as a component of Neverfail Engine through the Engine Management Service User Interface tasks or Neverfail Engine Manager Applications > Tasks > User Accounts. Neverfail recommends integrating Microsoft DNS into AD so that <code>DNSUpdate.exe</code> can identify all DNS Servers that require updating. 	<input type="checkbox"/>
Firewalls	<p>If using Windows Firewall, Engine Management Service can automatically configure the necessary ports for traffic. In the event that other than Windows Firewall is being used, configure the following specific ports to allow traffic to pass through:</p> <ul style="list-style-type: none"> From VMware vCenter Server -> Engine Management Service <ul style="list-style-type: none"> TCP 443 / 9727 / 9728 / Ephemeral port range From VMware vCenter Server -> The protected virtual machine <ul style="list-style-type: none"> TCP 443 / Ephemeral port range From Engine Management Service -> VMware vCenter Server <ul style="list-style-type: none"> TCP 443 / 9727 / 9728 / Ephemeral port range 	<input type="checkbox"/>

Server	Action
	<ul style="list-style-type: none"> From Engine Management Service -> The protected virtual machine <ul style="list-style-type: none"> TCP 7 / 445 / 135-139 / 9727 / 9728 / Ephemeral Port Range From the Protected Virtual Machine -> Engine Management Service <ul style="list-style-type: none"> TCP 7 / 445 / 135-139 / 9727 / 9728 / Ephemeral Port Range From the Protected Virtual Machine -> VMware vCenter Server <ul style="list-style-type: none"> TCP 443 / Ephemeral port range From Protected Virtual Machines -> VProtected Virtual Machines in Duo/Trio and back <ul style="list-style-type: none"> TCP 7 / 52267 / 57348 / Ephemeral port range From Management Workstation -> VProtected Virtual Machines in Duo/Trio and back <ul style="list-style-type: none"> TCP 52267 / 57348 / Ephemeral port range
	For more detailed information, see KB-2907 Firewall Configuration Requirements for Neverfail Continuity Engine.
	Note: <u>The default dynamic ephemeral port range for Windows 2008 and 2012 is ports 49152 through 65535.</u>

Server Deployment Architecture Options

The selected server architecture affects the requirements for hardware and the technique used to clone the Primary server.

Virtual-to-Virtual

Virtual-to-Virtual is the supported architecture if applications to be protected are already installed on the production (Primary) server running on a virtual machine. Benefits to this architecture include reduced hardware cost, shorter installation time, and use of the VMware Cloning for installation.

The Secondary virtual machine will be an exact clone of the Primary server and thus automatically meet the minimum requirements for installation of the Secondary server.

Each virtual machine used in the Virtual-to-Virtual pair should be on a separate ESX host to guard against failure at the host level.

Physical-to-Virtual

The Physical-to-Virtual architecture is used when the environment requires a mix of physical and virtual machines. This architecture is appropriate to avoid adding more physical servers or if you plan to migrate to virtual technologies over a period of time.

The Secondary Neverfail Engine virtual machine will be created from the Primary server.

- The specifications of the Secondary Neverfail Engine virtual machine must match the Primary physical server as follows:
 - Similar CPU
 - Identical Memory
- The Secondary Neverfail Engine virtual machine must have sufficient priority in resource management settings so that other virtual machines do not impact its performance.

Physical-to-Physical

The Physical-to-Physical architecture is used in environments where both the Primary and Secondary servers are physical servers. Use of Physical-to-Physical limits installation options as it requires using Neverfail Continuity Engine's manual cloning during the installation process. This architecture requires attention to detail when preparing for installation as both hardware and software must meet specific prerequisites.

Primary Server

The Primary server must meet the hardware and software requirements as specified in the [Pre-Install Requirements](#).

Secondary Server

The Secondary server operates as a near clone of the Primary server and must meet the following requirements.

- **Hardware**

Hardware should be equivalent to the Primary server to ensure adequate performance when the server is in the active role:

- Similar CPU
- Similar memory
- Identical number of NICs to the Primary server
- Drive letters must match the Primary server
- Available disk space must be greater than or equal to the Primary server

- **Software**

Software on the Secondary server must meet the following requirements.

- OS version and Service Pack version must match the Primary server
- OS must be installed to the same drive letter and directory as on the Primary server
- Machine name must be different from the Primary server prior to installing Neverfail Continuity Engine
- Set up in a workgroup prior to installing Neverfail Continuity Engine
- System date, time, and time zone settings must be consistent with the Primary server

Cloning Technology Options

Cloning the Primary server to create a nearly identical Secondary or Tertiary server involves different technologies depending on the selected server architecture.

Automated Cloning Technologies

The following cloning technologies are supported for creating cloned images for use as a Secondary or Tertiary server during the installation of Neverfail Engine:

- VMware vCenter virtual machine cloning is used when deploying a standby HA or standby DR server in a Virtual-to-Virtual environment.

- The VMware vCenter Converter is automatically used when cloning in a Physical-to-Virtual environment.

Note: *VMware Converter must be configured prior to attempting installation of the Secondary server.*

Manual Cloning Technologies

The following cloning technologies are supported with this version of Neverfail Engine:

- Using Windows Server Backup for Manual Cloning
- Using VMM for Hyper-V to Hyper-V for Manual Cloning
- Using SCVMM for Hyper-V to Hyper-V for Manual Cloning
- Using Paragon PPR for Manual Cloning
- Using XenCenter for Xen-to-Xen Manual Cloning
- Using virt manager for KVM-to-KVM Manual Cloning

Application Component Options

Neverfail Engine supports any of the plug-ins listed below:

Supported Plug-ins

- Neverfail for Exchange
 - ForeFront
 - Symantec Mail Security
- Neverfail for File Server
- Neverfail for IIS
- Neverfail for SharePoint Server
- Neverfail for SQL Server
- Neverfail for VMware vCenter Server
- Neverfail for VMware vSphere 6.0 Plug-in Suite
- Neverfail System Plug-in

Additionally, Neverfail Engine supports the Neverfail for Business Application Plug-in which may be installed post deployment.

Networking Configuration

Networking requirements are contingent upon how Neverfail Engine is to be deployed. To deploy as a High Availability (HA) solution, a LAN configuration is required. To deploy Neverfail Engine for Disaster Recovery (DR), a WAN configuration is required. To deploy in a Trio, both a LAN and a WAN configuration are used. Each network configuration has specific configuration requirements to ensure proper operation.

Note: *Neverfail recommends that the Neverfail Channel be configured on the same network as the Public network. If required to isolate for replication, the Neverfail Channel can be configured on a different subnet than the Public network.*

When Neverfail Engine is installed using a single NIC configuration, upon completion of installation, Neverfail recommends that you add an additional NIC to each server (Primary/Secondary/Tertiary) in order to provide network redundancy and then move the Neverfail Channel configuration to the newly added NICs. For more information about adding additional NICs to Neverfail Engine, see [Adding an Additional Network Interface Card](#) in this guide.

Local Area Network (LAN)

When deployed in a LAN environment, Neverfail Engine requires that both servers use the same Public IP address. Each server also requires a Neverfail Channel IP address.

Wide Area Network (WAN)

Neverfail Engine supports sites with different subnets. In this scenario, the Primary and Secondary servers in the Neverfail Engine Pair or Secondary and Tertiary in a Trio will require unique Public IP addresses in each subnet and a unique Neverfail Channel IP address in each subnet for each server.

WAN Requirements

WAN deployments require the following:

- Persistent static routing configured for the channel connection(s) where routing is required

Note: *This requirement can be avoided if the channel IP addresses are in the same subnet as the Public IP address in which case the default gateway can be used for routing.*

- One NIC (minimum)
- At least one Domain Controller at the Disaster Recovery (DR) site
- If the Primary and DR site uses the same subnet:
 - During install, follow the steps for a LAN or VLAN on the same subnet
 - Both the Primary and Secondary servers in the pair use the same Public IP address
- If the Primary and DR site use different subnets:
 - During install, follow the steps for a WAN
 - The Primary and Secondary servers in the Neverfail Engine pair require a separate Public IP address and a Neverfail Channel IP address
 - Provide a user account with rights to update DNS using the `DNSUpdate.exe` utility provided as a component of Neverfail Engine through the Engine Management Service User Interface tasks or Neverfail Advanced Management Client **Applications > Tasks > User Accounts**
 - Neverfail recommends integrating Microsoft DNS into AD so that `DNSUpdate.exe` can identify all DNS Servers that require updating
 - At least one Domain Controller at the DR site
 - Refer to the following articles in the Neverfail Knowledge Base:

Knowledge base article KB-1425 – Configuring DNS with Neverfail Continuity Engine in a WAN Environment

Knowledge base article KB-1599 – Configuring Neverfail Continuity Engine to Update BIND9 DNS Servers Deployed in a WAN

Bandwidth

Neverfail Engine includes automatic bandwidth optimization in WAN environments. This feature compresses data transferred over the Neverfail Channel, optimizing the traffic for low bandwidth connections causing some additional CPU load on the active server.

Determine the available bandwidth and estimate the required volume of data throughput to determine acceptable latency for the throughput. Additionally, the bandwidth can affect the required queue size to accommodate the estimated volume of data. Neverfail recommends making a minimum of 1Mbit of spare bandwidth available to Neverfail Engine.

Latency

Latency has a direct effect on data throughput. Latency on the link should not fall below the standard defined for a T1 connection (2-5ms for the first hop).

Neverfail SCOPE Data Collector Service can assist in determining the available bandwidth, required bandwidth, and server workload. For more information about Neverfail SCOPE Data Collector Service, contact Neverfail Professional Services.

Network Interface Card (NIC) Configuration

Neverfail Engine supports use of either multiple NICs or a single NIC.

This release of Neverfail Engine adds very flexible support for configuring NICs with Public and Channel connections. The following scenarios are some supported:

- **Single NIC Installation** : Neverfail Engine is installed on a server having a single NIC, which is shared by both the Public Network and the Neverfail Channel. This can simplify the install process by avoiding down-time when adding a NIC.
- **Adding a NIC post-installation** . Using a single NIC results in a potential single point of failure. To prevent a single point of failure, additional NICs can be added post-installation, and the Public and Neverfail Channel IP addresses distributed across these. See *Adding a Network Card*.
- **Multiple NIC Installation.** Neverfail Engine can be installed on a server with multiple NICs. You can choose which NIC will be used for the Neverfail Channel connection.

Primary Server

The Primary server is configured with the following connections:

- A Public network connection configured with a static Public IP address, network mask, gateway address, preferred DNS server address, and secondary (if applicable) DNS server address.
- Neverfail Channel connection(s) configured with a static IP address in the same or a different subnet than the Public IP address, and with a different IP address than the Secondary server channel, and network mask. No gateway or DNS server address is configured where a dedicated NIC is used. NetBIOS will be filtered on the passive server to prevent server name conflicts.
- The *Register this connection's addresses in DNS* check box must be cleared on the Neverfail Channel connection(s) prior to installing Neverfail Engine.

Secondary/Tertiary Server

The Secondary/Tertiary server will have the same number of NICs as the Primary server, with the same names and will be configured as follows:

- A Public connection configured with a static IP address, network mask, gateway address, preferred DNS server address, and secondary (if applicable) DNS server address.

Note: *If deploying as a pair in a WAN, the Public IP address of the Secondary server may be in a different subnet than the Primary server.*

Note: *If configured in a trio, the Primary and Secondary servers are configured for LAN deployment and the Tertiary server is configured for a WAN deployment.*

- Neverfail Channel network connection(s) configured on the same or a separate dedicated NIC with a static IP address in the same or a different subnet than the Secondary/Tertiary Public IP address, and with a different IP address than the Primary or Secondary (for Tertiary) server's Neverfail Channel NIC, and a network mask. A gateway address and DNS address are not configured by the user. NetBIOS will be filtered to prevent server name conflicts.
- The *Register this connection's addresses in DNS* check box must be cleared on the Neverfail Channel connection(s) prior to installing Neverfail Engine.

Note: *Neverfail recommends that this network change be made during a scheduled downtime to minimize risk of a system outage. Once this is done, you should immediately replace the dynamic "A" record for the Neverfail Engine protected server with a static entry with a TTL of 45 seconds.*

Firewall Configuration Requirements

When firewalls are used to protect networks, you must configure them to allow traffic to pass through specific ports for Neverfail Engine installation and management. If using Windows Firewall, Engine Management Service can automatically configure the necessary ports for traffic. In the event that other than Windows Firewall is being used, configure the following specific ports to allow traffic to pass through:

- Ports 9727 and 9728 for managing Neverfail Engine from the Engine Management Service
- Port 52267 for the Client Connection port
- Port 57348 for the Default Channel port

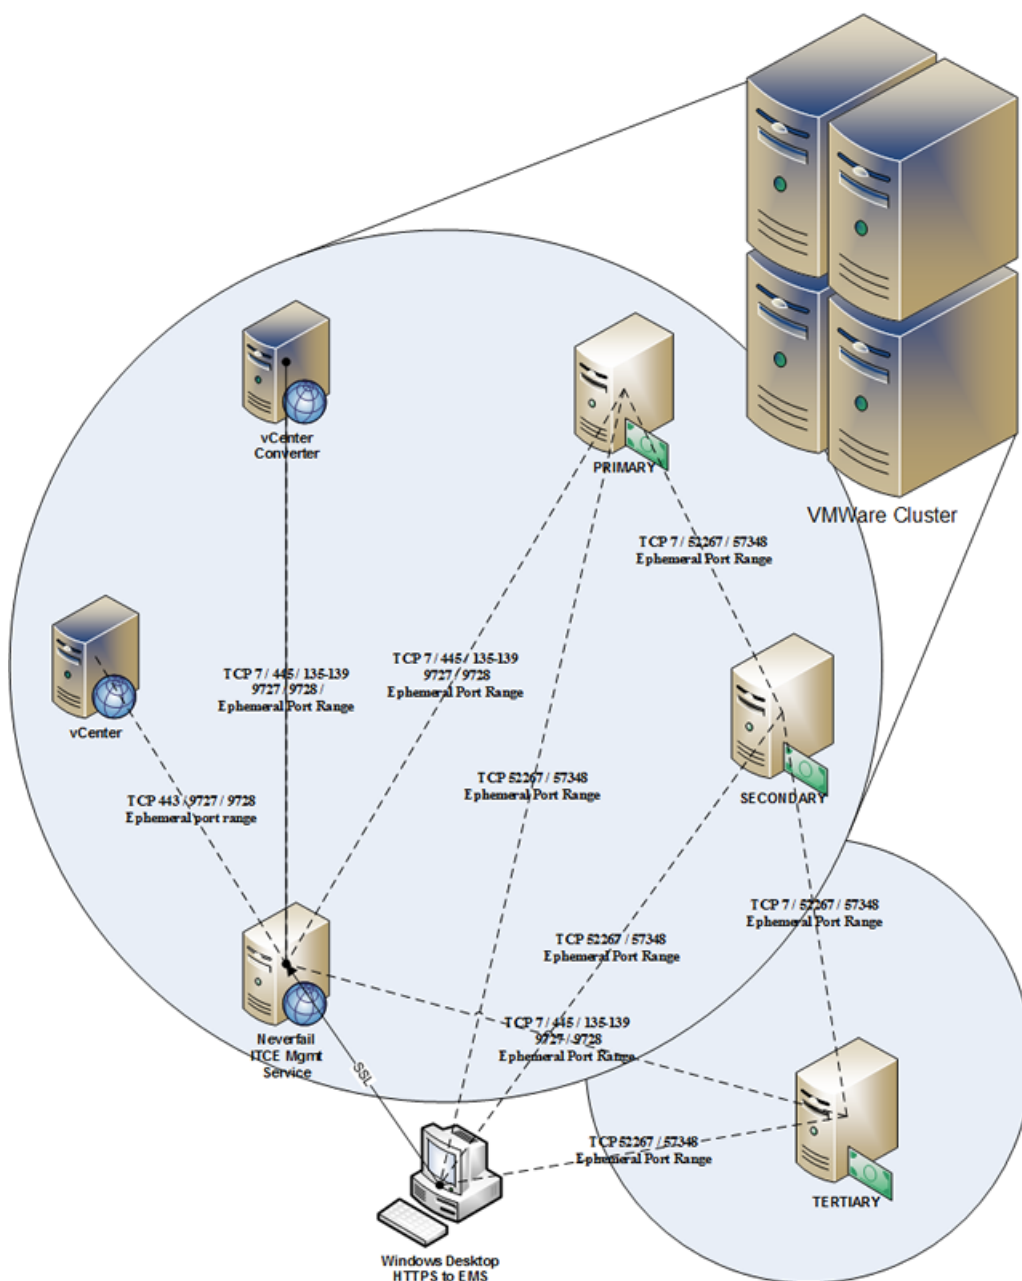


Figure 3: Firewall Ports diagram

Important: When installing on Windows Server 2008 R2, Microsoft Windows may change the connection type from a Private network to an Unidentified network after you have configured the firewall port to allow channel communications resulting in the previously configured firewall changes to be reset for the new network type (Unidentified).

The firewall rules must be recreated to allow traffic to pass through for the Client Connection port and the Default Channel port. Neverfail recommends that the firewall be configured to allow the Client to connect to the Client Connection port by process, `nfgui.exe`, rather than by a specific port. To enable Channel communications between servers, change the Network List Manager Policy so that the Neverfail Channel network is identified as a Private Network, and not the default Unidentified

Network, and configure the firewall to allow traffic to pass through on Port 57348, the Default Channel port.

Anti-Malware Recommendations

Consult with and implement the advice of your anti-malware provider, as Neverfail Continuity Engine guidelines often follow these recommendations. Consult the Artisan Knowledge Base for up to date information on specific anti-malware products.

Do not use file level anti-malware to protect application server databases, such as Microsoft SQL Server databases. The nature of database contents can cause false positives in malware detection, leading to failed database applications, data integrity errors, and performance degradation.

Neverfail recommends that when implementing Neverfail Continuity Engine, you do not replicate file level anti-malware temp files using Neverfail Engine.

The file level anti-malware software running on the Primary server must be the same as the software that runs on the Secondary server. In addition, the same file level anti-malware must run during both active and passive roles.

Configure file level anti-malware to use the Management IP address on the passive server(s) for malware definition updates. If this is not possible, manually update malware definitions on the passive server(s).

Exclude the following Neverfail directories from file level anti-malware scans (`C:\Program Files\Neverfail\` is the default installation directory):

- `C:\Program Files\Neverfail\r2\logs`
- `C:\Program Files\Neverfail\r2\log`

Any configuration changes made to a file level anti-malware product on one server (such as exclusions) must be made on the other server as well. Neverfail Engine does not replicate this information.

Chapter 3

Installing Neverfail Continuity Engine

This chapter discusses the installation process used to implement Neverfail Continuity Engine on Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012R2 when the Secondary or Tertiary server is virtual. Prior to installing Neverfail Continuity Engine, you should identify the deployment options you want so that during the installation process you are prepared to select the required options to achieve your configuration goals.

After selecting implementation options, begin the installation process. During the installation process, Engine Management Service performs a variety of checks to ensure the target server meets the minimum requirements for a successful installation. Should the target server fail one of the checks, a critical stop or warning message appears. You must resolve critical stops before you can proceed with setup.

Topics:

- [Installing Neverfail Continuity Engine](#)
- [Using the Engine Management Service User Interface](#)
- [Post Installation Configuration](#)

Installing Neverfail Continuity Engine

Prerequisites

Prior to attempting installation of Neverfail Continuity Engine Management Service, ensure that the server meets all of the pre-requisites stated in [Pre-Install Requirements](#).

Procedure

To install the Neverfail Continuity Engine Management Service:

1. Having verified all of the environmental prerequisites are met, download the Neverfail Continuity Engine .msi file to an appropriate location.

Note: *Install on any server running Windows Server 2008 R2 64-bit or later with connectivity to a VMware vCenter Server 5.1 or later or a Desktop Edition of Windows OS 7, 8.x, or 10.*

2. While logged in as the Local Administrator, double-click the `Neverfail-CE-[n]-[n]-[nnnnn]-x64.msi` file to initiate installation of the Neverfail Continuity Engine Management Service.

The *Welcome* page is displayed.

3. Click **Next**.
The *End User License Agreement* page is displayed.
4. Review the *End User License Agreement* and select *I accept the terms in the License Agreement*. Click **Next**.
The *Firewall Modification* screen is displayed.
5. If using something other than Windows Firewall, manually configure Firewall Rules to allow TCP on Ports 9727 and 9728 at this time. If using Windows Firewall, the *Inbound Firewall Rules* are created automatically and no actions are necessary. Click **Next**.
The *Ready to install Neverfail Continuity Engine* screen is displayed.
6. Click **Install**.
The *Installing Neverfail Continuity Engine* screen is displayed. When the installation has finished installing the appropriate components, the *Completed the Neverfail Continuity Engine Setup Wizard* screen is displayed.

Note: *If you are upgrading from version 7.x, you may be asked if a service should be stopped. Press to allow the service to be stopped automatically.*

7. Click **Finish**.
Once installation of the Neverfail Continuity Engine Management Service is complete, the Neverfail Continuity Engine Management Service User Interface will launch automatically.
8. Login to the Neverfail Continuity Engine Management Service user interface using a local administrator account. If you have upgraded from an earlier version, the *Protected Servers* pane should display your list of servers.

Deploying Neverfail Engine on the Primary Server

Prerequisites

Prior to deploying Neverfail Engine on the target Primary server, ensure that the server meets all of the pre-requisites stated in [Pre-Install Requirements](#). During the installation process, Engine Management Service will install Neverfail Engine on the target servers identified in the cluster and validate that the servers meet the minimum requirements for a successful installation.

Procedure

To install Neverfail Engine on the Primary server:

1. Login to the Neverfail Continuity Engine Management Service UI and select the *Management* drop-down. Click on **Deploy > Deploy to a Primary server**.
The *Deploy Engine* page is displayed.
2. Enter the DNS name or IP address of the target (Primary) server, or select a virtual server from the inventory. Enter the name of a user account with full administrator permissions and click **Next**.
The *Validating Install* step is displayed. The Engine Management Service automatically configures Windows firewalls to allow installation to continue and communications via the Neverfail Channel and Neverfail Engine.
3. Once the *Validating Install* step completes and displays that the server is a valid target, click **Next**.
The *Select Public (Principal) IP Address* step is displayed.
4. Validate the Public IP address displayed and ensure the check box is selected for addresses that should be available for client connection. Click **Next**.
The *Ready to Complete* step is displayed.

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.
The *Status* page is displayed.

Automated Deployment of Stand-by Servers with Automatic Cloning

1. You have the following options:
 - If the Primary server is physical, go to [Step 2](#)
 - If the Primary server is virtual, go to [Step 4](#)
2. Click on the **Converter** button. The *Configure Connection to VMware vCenter Converter* page is displayed. Provide the URL where the VMware vCenter Converter resides and provide the Username and Password with local Administrator permissions on the machine where VMware vCenter Converter is installed. Click **Next**.
The *Ready to Complete* step is displayed.
3. Review the URL and if accurate, click **Finish**.

Note: If VMware vCenter Server is configured before connecting to VMware vCenter Converter, the success or failure of connecting to the VMware Converter is indicated as a vSphere Task and also by the icon shown next to the **Converter** button.

4. Navigate to **Management > Deploy**.
5. Select one of the following depending on the environment you intend to support:
 - **Add a stand-by server for high availability**, go to [Step 6](#)
 - **Add a stand-by server for disaster recovery**, go to [Step 12](#)
 - **Create Secondary and Tertiary stand-by VMs for HA and DR**, go to [Step 19](#)

Note: You can also create a stand-by VM for Disaster Recovery for an existing High Availability pair, and vice-versa.

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

6. Select clone type step – select to use automated cloning (recommended). Click **Next**.
The *Select channel IP addresses* step is displayed.
7. 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**.

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

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

8. Select the Datacenter and Host where the Secondary server will be created and click **Next**.
The *Select Storage* step is displayed.

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

The *Select storage (optional)* step is displayed.

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

The *Ready to complete* step is displayed.

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

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

11. Once complete, perform Post Installation Configuration tasks as listed in the *Neverfail Continuity Engine Installation Guide*.
12. On the Neverfail Engine Management Service user interface, click the **Add a stand-by server for Disaster Recovery**.

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

13. Select either of the following:

- The public (principal) IP address will be identical to the Primary server.
- The public (principal) IP address will be different than the Primary server - you must add credentials to be used for updating DNS.

Click **Next**. The *Select Channel IP Addresses* step is displayed.

14. 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**.

The *Select Clone Type* step is displayed.

15. 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. Click **Next**.

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

The *Select Host* step is displayed.

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

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

The *Select Storage* step is displayed.

17. Select the storage location for the virtual machine. Click **Next**.
18. Review the information on the *Ready to Complete* step and if accurate, click **Finish** to create the Secondary server.
Once cloning process is complete, automatic reconfiguration of the stand-by server will take place requiring only a few minutes to finish. Once complete, perform *Post Installation Configuration* tasks as listed in this guide.

19. 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). On the Neverfail Engine Management Service, navigate to the **Management > Deploy** drop-down and select *Create Secondary and Tertiary VMs for HA and DR*. The *Create Secondary and Tertiary VMs for High Availability and Disaster Recovery* wizard is displayed.
20. Review the information in the step and then click **Next**.
The *Select host* step is displayed.
21. Click on the appropriate Datacenter to display all available hosts. Select the intended host for the Secondary server and then click **Next**.
The *Select storage* step is displayed.
22. Select the intended datastore for the Secondary VM, and then click **Next**.
The *Configure Tertiary VM* step is displayed.
23. Review the contents of the step and then click **Next**.
The *Select public IP address* step is displayed.
24. 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 in a separate subnet in the *Public IP Addresses* field. 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**.
The *Select VM move type* step is displayed.
25. Review the definitions of the options and then select whether the VM will be transferred manually or not. Click **Next**.
The *Select host* step is displayed.
26. Click on the appropriate Datacenter to display all available hosts. Select the intended host for the Tertiary server and then click **Next**.
The *Select storage* step is displayed.
27. Select the intended datastore for the Tertiary VM, and then click **Next**.
The *Configuring Channel Communications* step is displayed.
28. Review the contents of the step and then click **Next**.
The *Primary-Secondary* step is displayed.
29. Select the appropriate network adapter and then enter the channel IP addresses for Primary-Secondary communications. Click **Next**.
The *Secondary-Tertiary* step is displayed.
30. Select the appropriate network adapter and then enter the channel IP addresses for Secondary-Tertiary communications. Click **Next**.
The *Tertiary-Primary* step is displayed.
31. Select the appropriate network adapter and then enter the channel IP addresses for Tertiary-Primary communications. Click **Next**.
The *Ready to complete* step is displayed.
32. 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.

Semi-Automatic Deployment of Stand-by Servers Leveraging Manual Cloning

1. Navigate to **Management > Deploy**.
2. Select one of the following depending on the environment you intend to support:
 - **Add a stand-by server for high availability**, go to [Step 3](#)
 - **Add a stand-by server for disaster recovery**, go to [Step 7](#)

Note: You can also create a stand-by VM for Disaster Recovery for an existing High Availability pair, and vice-versa.

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

3. Select clone type step – select to use manual cloning. Click **Next**.
The *Select channel IP addresses* step is displayed.
4. 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**.
The *Ready to complete* step is displayed.
5. Review the information on the *Ready to Complete* step and if accurate, click **Finish** to prepare the Secondary server for manual cloning using a third-party tool.

During the pre-condition check, the following status messages will display.

- Shutting down Neverfail Software on all Nodes of (HOSTNAME)
- Reconfiguring Engine to participate in an extended cluster
- Waiting for server to become Active
- Completed reconfiguration of Engine
- PRIMARY server ready to be cloned. Please clone the PRIMARY

Once cloning process is complete, start the new stand-by server. The servers will connect and begin replication automatically.

6. Once complete, perform Post Installation Configuration tasks as listed in the *Neverfail Continuity Engine Installation Guide*.
7. On the Neverfail Engine Management Service user interface, click the **Add a stand-by server for Disaster Recovery**.
The *Add a stand-by server for disaster recovery* page is displayed.
8. Select either of the following:
 - The public (principal) IP address will be identical to the Primary server.
 - The public (principal) IP address will be different than the Primary server - you must add credentials to be used for updating DNS.

Click **Next**. The *Select Channel IP Addresses* step is displayed.

9. 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**.
The *Select Clone Type* step is displayed.
10. Select the *Select manual cloning* option. Click **Next**.
The *Ready to Complete* step is displayed.
11. Review the information on the *Ready to Complete* step and if accurate, click **Finish** to prepare the Secondary server for manual cloning using a third party.
During the pre-condition check, the following status messages will display.
 - Shutting down Neverfail Software on all Nodes of (HOSTNAME)
 - Reconfiguring Engine to participate in an extended cluster
 - Waiting for server to become Active

- Completed reconfiguration of Engine
- PRIMARY server ready to be cloned. Please clone the PRIMARY

Once cloning process is complete, start the new stand-by server. The servers will connect and begin replication automatically. Once complete, perform *Post Installation Configuration* tasks as listed in this guide.

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.

Configure Connection to VMware vCenter Server

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.

Procedure

To configure a connection to VMware vCenter Server:

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

Configure Connection to VMware vCenter Server

1) Configure vCenter
2) 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

Privilege	Allowed

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

Back Next Finish Cancel

Figure 4: Configure vCenter

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

The screenshot shows a dialog box titled "Configure Connection to VMware vCenter Server". On the left, a sidebar lists two steps: "1) Configure vCenter" and "2) Ready to complete", with the second step selected. The main area displays the configuration details:

- VMware vCenter Server URL:** https://192.168.0.3/sdk
- Account:** administrator@vsphere.local
- A green checkmark icon followed by the text: "Configuration successfully updated."
- Account Privileges:** A table listing privileges and their status.
- VMware vCenter Configuration:** A text box providing additional information about privileges and cloning VMs.

At the bottom right, there are four buttons: "Back", "Next", "Finish", and "Cancel".

Privilege	Allowed
Extension.Update	true
Task.Create	true
Extension.Register	true
Global.LogEvent	true
Extension.Unregister	true
Task.Update	true

VMware vCenter Configuration

If any of the privileges listed are not available to the account, tasks and events may not be visible in vCenter.

You will also require privileges for cloning VMs. See KB 2901 for more information.

Figure 5: Ready to Complete

Configure VMware vCenter Converter

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.

Prerequisites

VMware vCenter Converter 5.5 or later must be installed manually.

Procedure

To configure the VMware vCenter Converter:

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

The screenshot shows the 'Configure Connection to VMware vCenter Converter' wizard. The left pane has two steps: '1) Configure Converter' (selected) and '2) Ready to complete'. The main area contains three input fields: 'Enter the URL for the VMware vCenter Converter' with the value 'https://192.168.0.28:443/converter/sdk', 'Enter the name of an administrator account on the VMware vCenter Converter server' with the value 'Administrator', and 'Enter the password for the account' with a masked password. Below these is a blue box titled 'VMware vCenter Converter' containing text about its purpose and requirements, followed by a link 'Obtain VMware vCenter Converter'. At the bottom are 'Back', 'Next', 'Finish', and 'Cancel' buttons.

Figure 6: Configure 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**.

The screenshot shows the 'Configure Connection to VMware vCenter Converter' wizard at Step 2: 'Ready to complete'. The left pane shows '2) Ready to complete' as the active step. The main area displays a summary of the configuration: 'VMware vCenter Converter URL' is 'https://192.168.0.28:443/converter/sdk'. Below this is a green box titled 'VMware vCenter Converter Configuration' with a green checkmark icon, containing the text 'Configuration updated. Connection to VMware vCenter Converter will require up to 30s to validate.' At the bottom are 'Back', 'Next', 'Finish', and 'Cancel' buttons.

Figure 7: Ready to Complete

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

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.

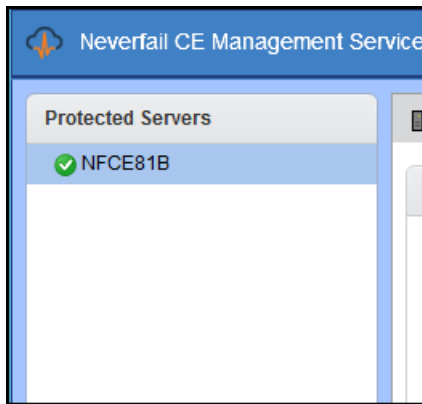


Figure 8: Protected Servers

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.

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.

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.

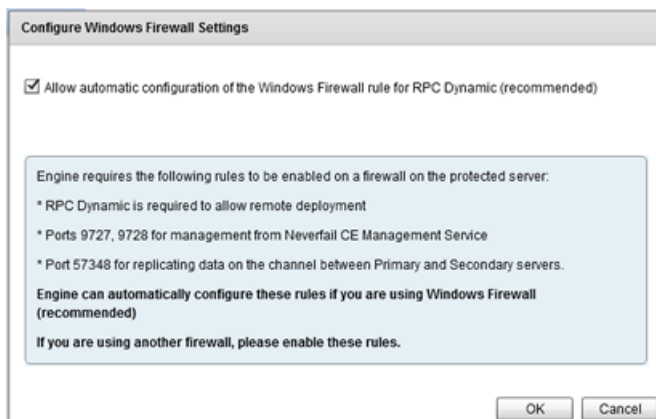


Figure 9: Configure Windows Firewall Settings

Deploy to a Primary Server

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

Prerequisites

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.

Procedure

To Deploy Neverfail Engine:

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

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

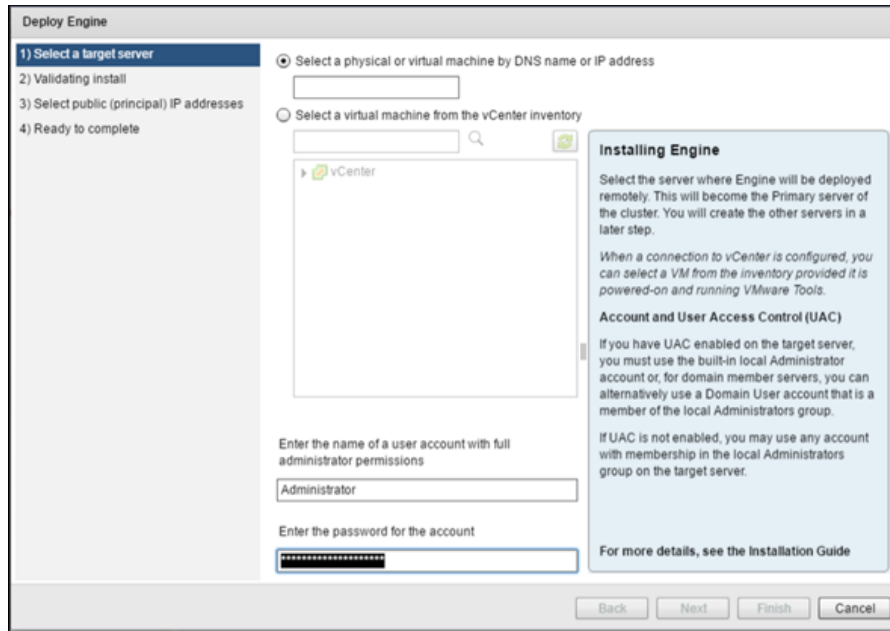


Figure 10: Deploy Neverfail Engine page

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.

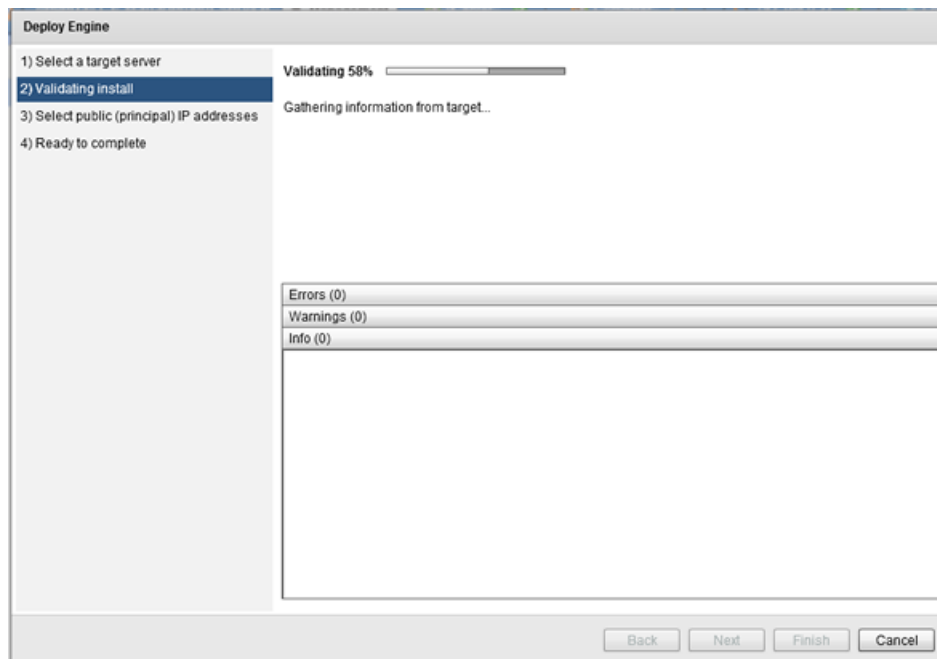


Figure 11: Validating Install step

- 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.

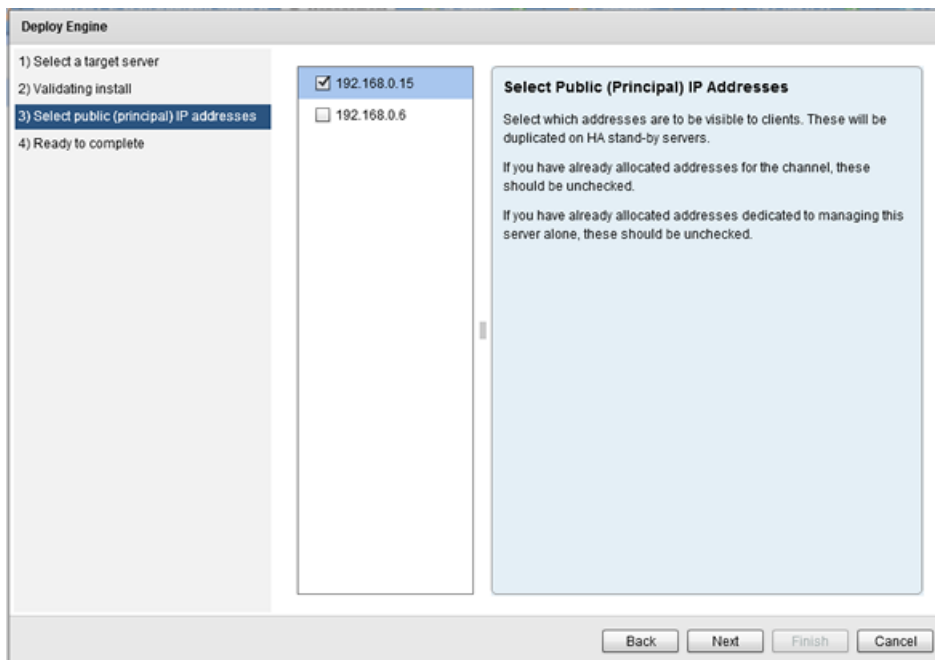


Figure 12: Select public (principal) IP addresses step

- 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.

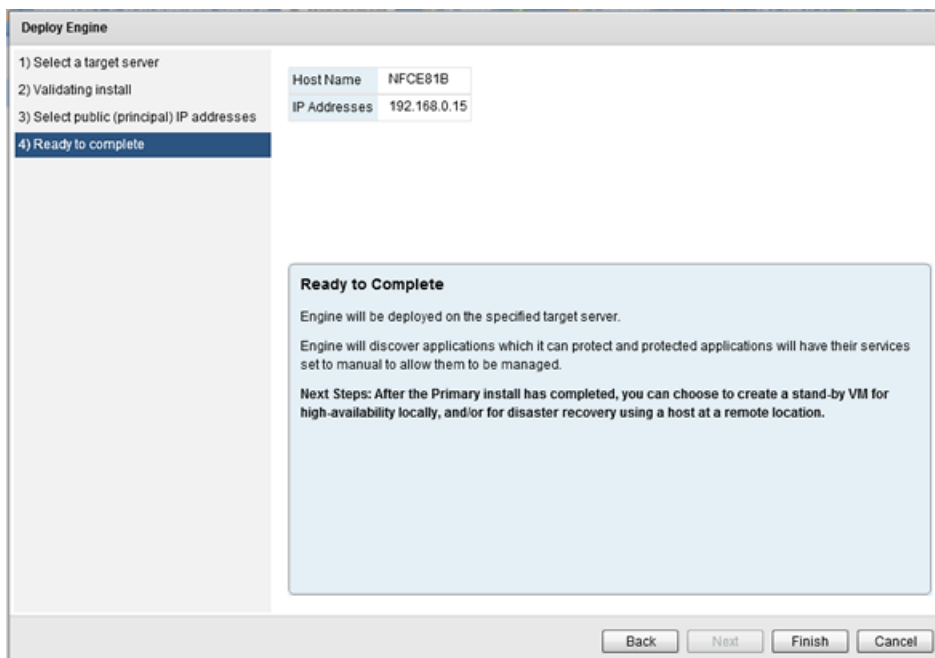


Figure 13: Ready to complete step

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 .

Upgrade the Selected Server

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

1. From the **Management** drop-down, navigate to **Deploy > Upgrade the selected server**.
The *Upgrade Engine* page is displayed.

Figure 14: Upgrade Engine

2. Enter the name of the local built-in Administrator account and password. After confirming that no users are logged into the Primary, Secondary (or Tertiary) servers, select the check box.
3. Select to either upgrade all server nodes or only a specific server in the cluster. Click **Next**.

Note: *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 15: Validating upgrade step

4. Once validation is complete, click **Next**.
The *Ready to complete* step is displayed.




Figure 16: Ready to complete step

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

Uninstall from the Selected Server

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

Procedure

To uninstall from the selected server:

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 17: 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.

Note: This option is only available if you attempt to uninstall a v8.1 cluster from Neverfail Continuity Engine Management Service v8.1

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.

Note: 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 (built-in) 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.

Add a Stand-by Server for High Availability

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.

Procedure

To add a stand-by VM for high availability:

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**.

The screenshot shows a web-based wizard titled "Add a Stand-by Server for High Availability". On the left is a vertical list of steps: 1) Select clone type (highlighted), 2) Select channel IP addresses, 3) Select a host (optional), 4) Select storage (optional), 5) Provide additional network settings (optional), and 6) Ready to complete. The main content area has two radio buttons: "Select automated cloning (recommended, requires vCenter)" which is selected, and "Select manual cloning". Below these is a light blue box titled "Select Clone Type" containing explanatory text and two options: "a) A connection to vCenter is configured. The stand-by server will be created as a VM. The source server VM and the target host are both managed by this instance of vCenter." and "b) Connections to vCenter and VMware Converter are configured. The stand-by server will be created as a VM." It also mentions that manual cloning can be performed using a third-party tool. At the bottom right are buttons for "Back", "Next", "Finish", and "Cancel".

Figure 18: 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**.

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

Figure 19: Select Channel IP Addresses step

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

4. Select the Datacenter and Host where the Secondary server will be created and click **Next**. The *Select Storage* step is displayed.

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

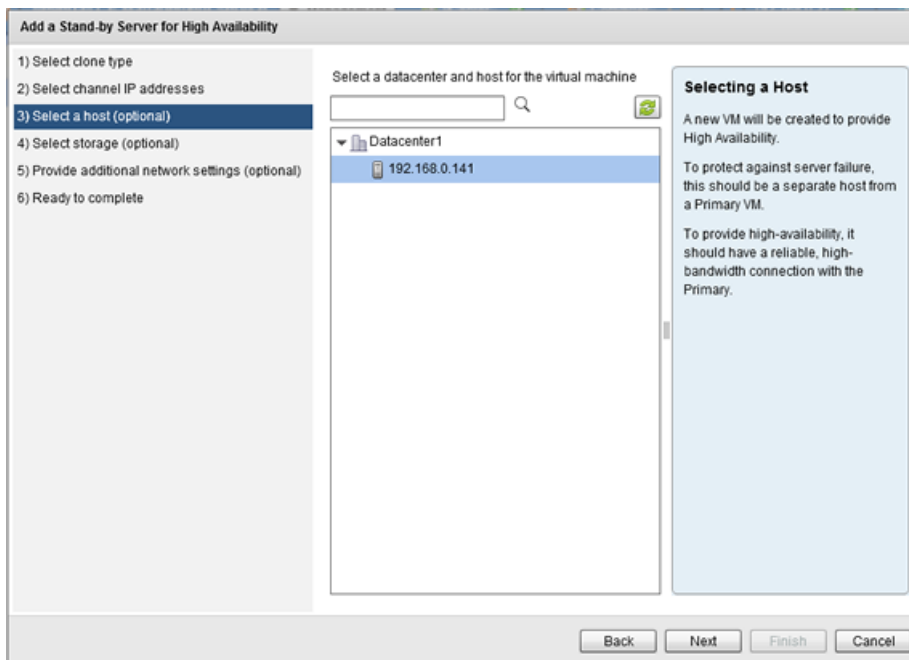


Figure 20: Select Host step

The *Select storage (optional)* step is displayed.

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

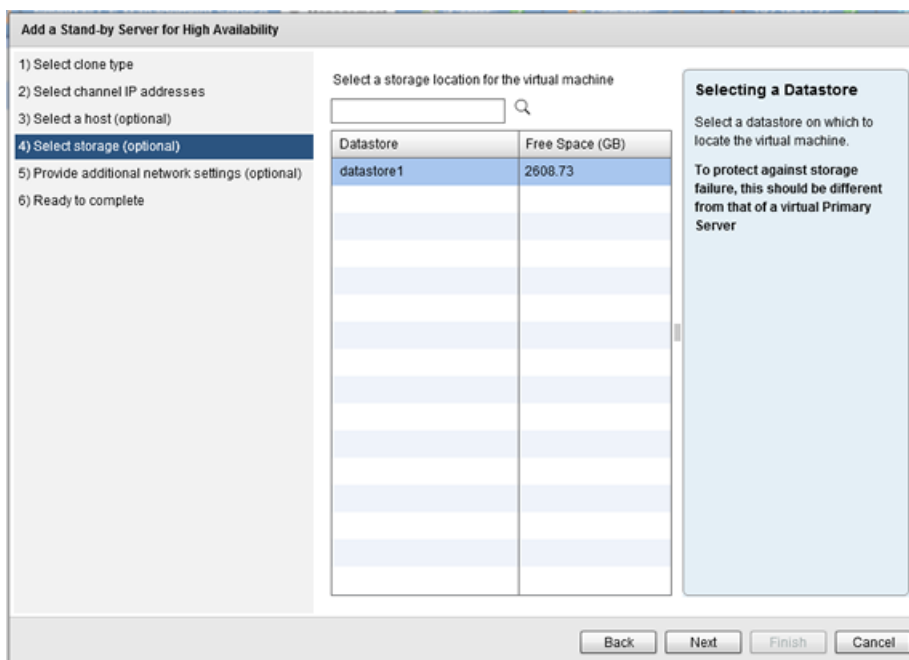


Figure 21: Select Storage step

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

The *Ready to complete* step is displayed.

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

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

The screenshot shows a wizard window titled "Add a Stand-by Server for High Availability". On the left, a list of steps is shown, with "6) Ready to complete" selected and highlighted in blue. The main area on the right contains a table of configuration details and a "Ready to Complete" message box.

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.

At the bottom right of the window are four buttons: "Back", "Next", "Finish", and "Cancel".

Figure 22: Ready to Complete step

- Once complete, perform [Post Installation Configuration](#) tasks listed in this guide.

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

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).

Procedure

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

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

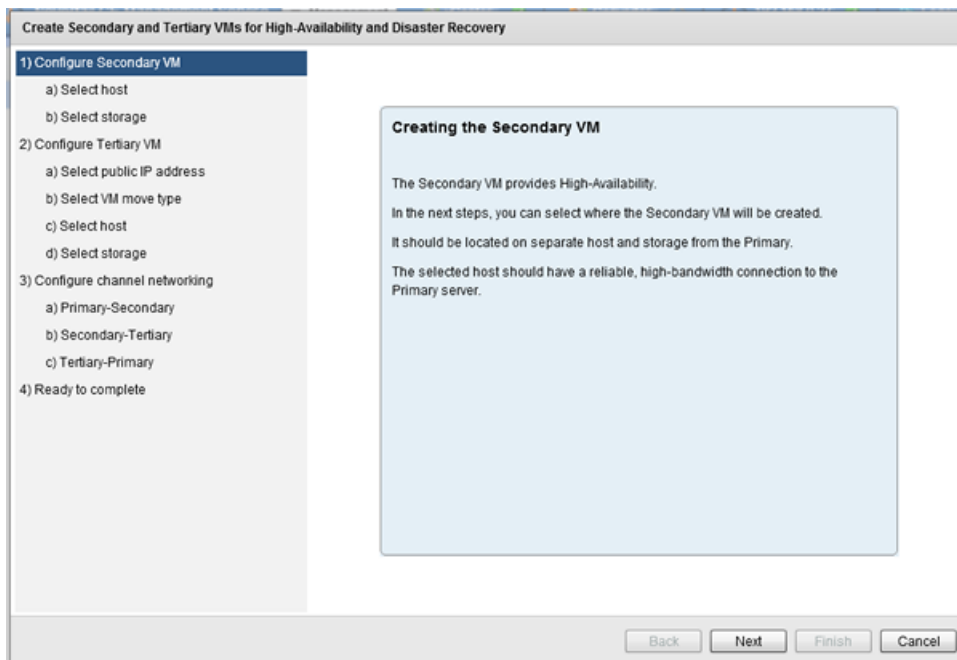


Figure 23: Configure Secondary VM step

2. Review the information in the step and then click **Next**. The *Select host* step is displayed.




Figure 24: Select host step

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

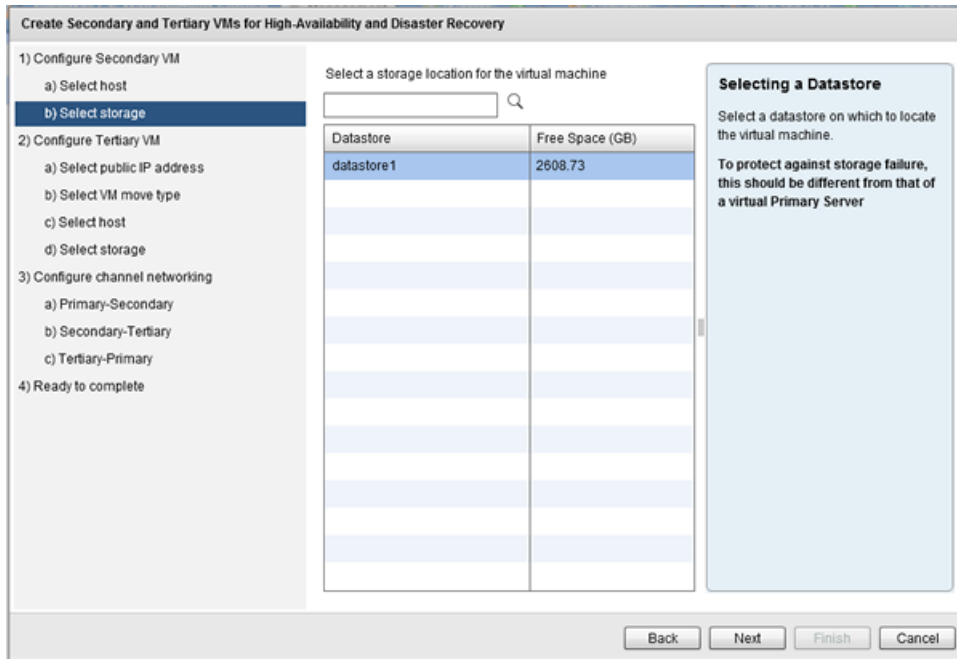


Figure 25: Select storage step

4. Select the intended datastore for the Secondary VM, and then click **Next**. The *Configure Tertiary VM* step is displayed.

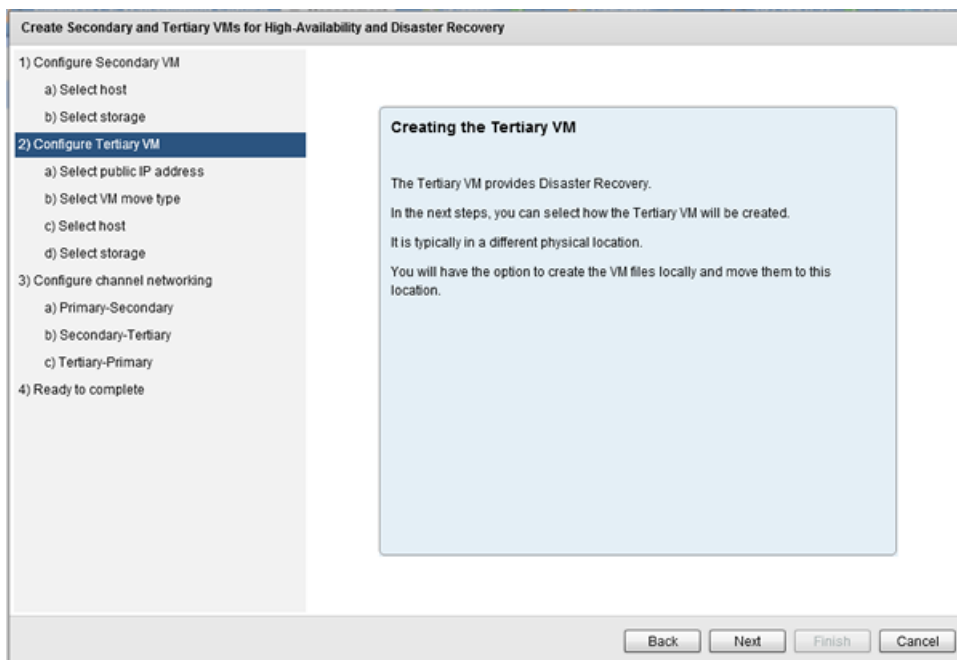


Figure 26: Configure Tertiary VM step

- Review the contents of the step and then click **Next**. The *Select public IP address* step is displayed.

Figure 27: Select public IP address step

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 in a separate subnet in the *Public IP Addresses* field. 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**. The *Select VM move type* step is displayed.

Figure 28: Select VM move type step

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

The *Select host* step is displayed.

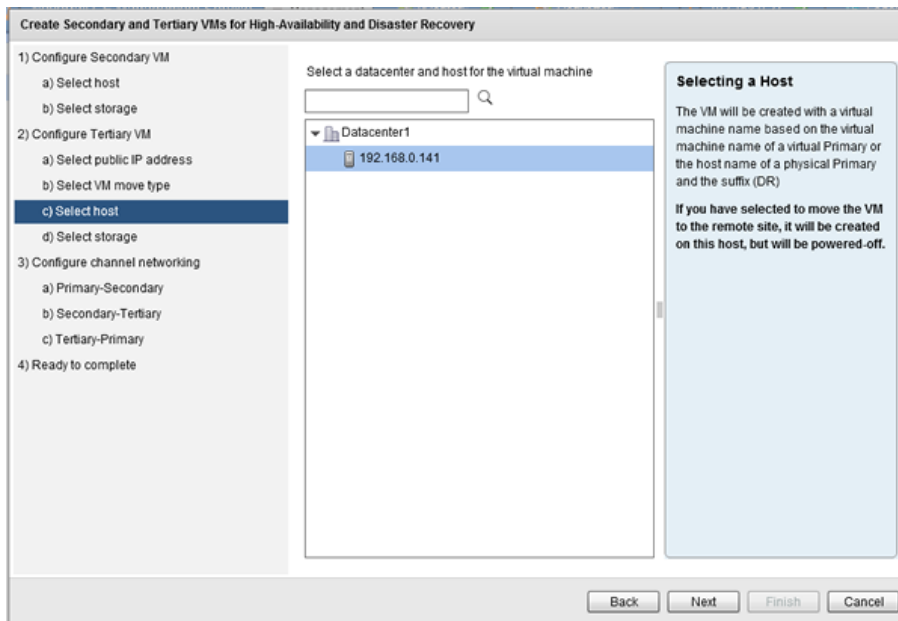


Figure 29: Select host step

- Click on the appropriate Datacenter to display all available hosts. Select the intended host for the Tertiary server and then click **Next**.
The *Select storage* step is displayed.

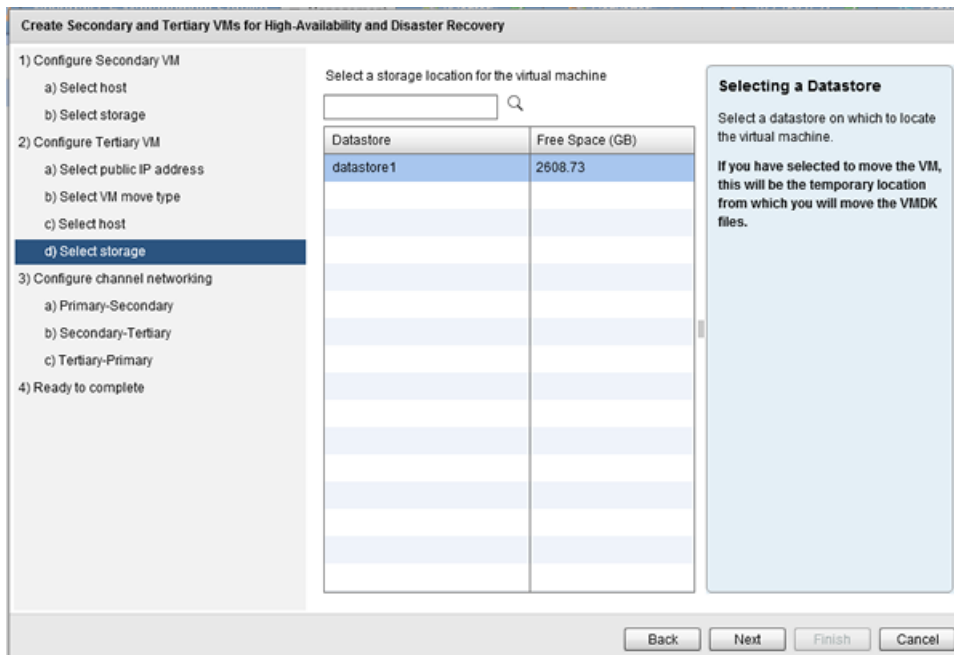


Figure 30: Select storage step

9. Select the intended datastore for the Tertiary VM, and then click **Next**. The *Configuring Channel Communications* step is displayed.

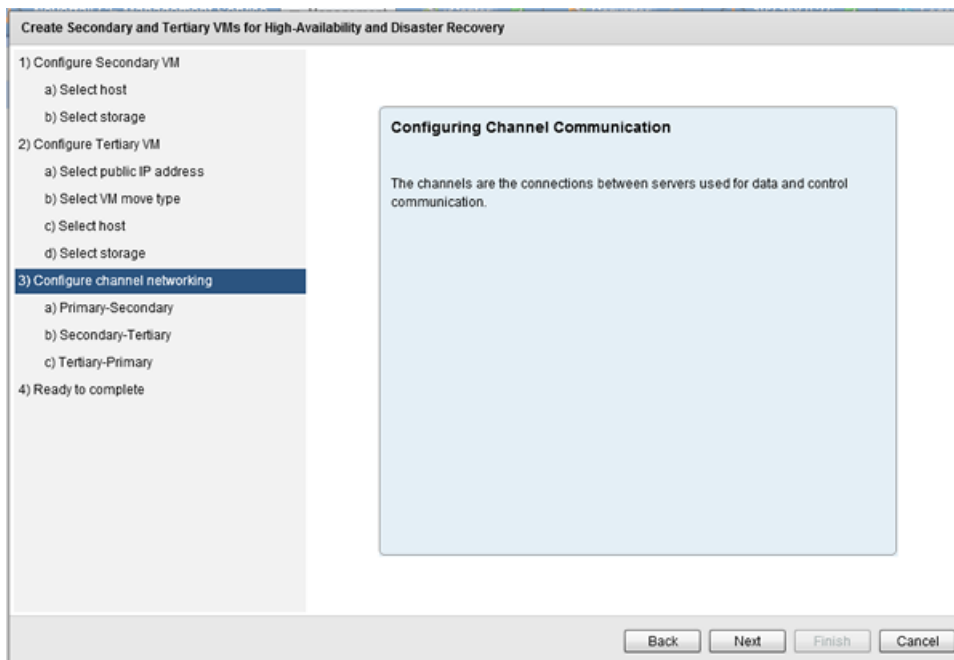


Figure 31: Configure channel networking step

10. Review the contents of the step and then click **Next**.
The *Primary-Secondary* step is displayed.

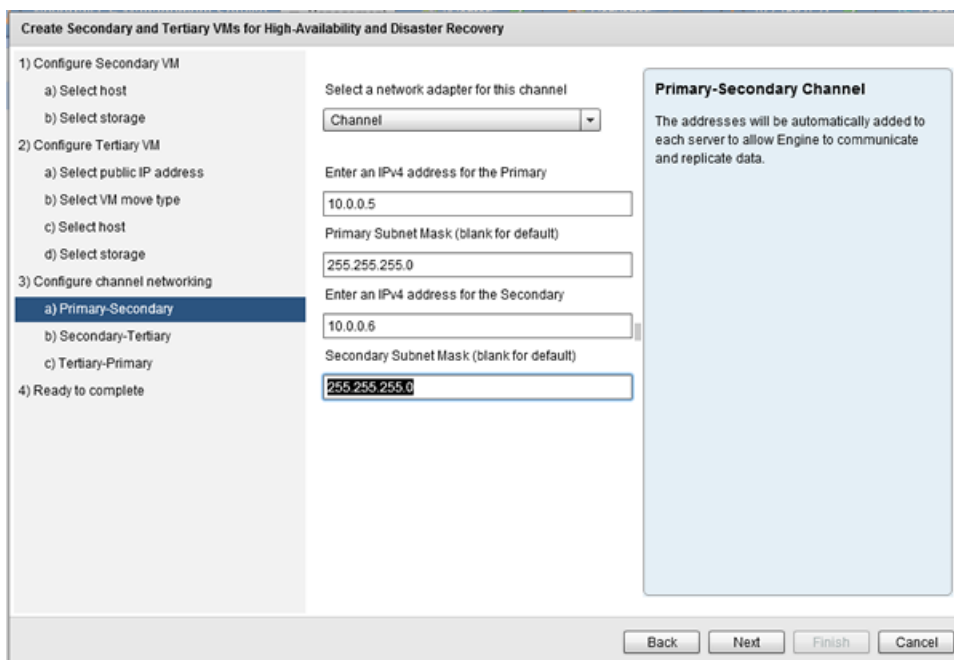


Figure 32: Primary-Secondary step

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

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

Select a network adapter for this channel

Public

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

Subnet Mask (blank for default)

255.255.255.0

Secondary-Tertiary Channel

The addresses will be automatically added to each server to data replication and cluster communication

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

Back Next Finish Cancel

Figure 33: Secondary-Tertiary step

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

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

Select a network adapter for this channel

Channel

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

Primary Subnet Mask (blank for default)

255.255.255.0

Tertiary-Primary Channel

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

Figure 34: Tertiary-Primary step

13. Select the appropriate network adapter and then enter the channel IP addresses for Tertiary-Primary communications. Click **Next**. The *Ready to complete* step is displayed.

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
Secondary Datacenter	Datacenter1
Secondary Host	192.168.0.141
Secondary Datastore	datastore1
Tertiary Datacenter	Datacenter1
Tertiary Host	192.168.0.141
Tertiary Datastore	datastore1
Tertiary Public IP Address	192.168.10.7
Location for Tertiary VM	Use Tertiary host location
Gateway	192.168.10.1
Preferred DNS	192.168.10.2
Alternate DNS	

P-S Channel IP Address	10.0.0.5
P-S Subnet Mask	255.255.255.0
S-P Channel IP Address	10.0.0.6
S-P Subnet Mask	255.255.255.0
S-T Channel IP Address	10.0.1.6
S-T Subnet Mask	255.255.255.0
T-S Channel IP Address	10.0.1.7
T-S Subnet Mask	255.255.255.0
T-P Channel IP Address	10.0.2.7
T-P Subnet Mask	255.255.255.0
P-T Channel IP Address	10.0.2.5
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

Figure 35: Ready to complete step

- 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.

Add a Stand-by Server for Disaster Recovery

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.

Procedure

To add a stand-by server for disaster recovery:

- 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:
 - The public (principal) IP address will be identical to the Primary server.
 - The public (principal) IP address will be different than the Primary server - you must add credentials to be used for updating DNS.

Click **Next**.

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: Public

Public IP Addresses: 192.168.0.15 Add... Remove

Enter the gateway: 192.168.0.1

Enter the preferred DNS server: 192.168.0.2

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

Figure 36: Select Public IP Address step

The *Select Channel IP Addresses* step is displayed.

3. 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**.

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: Channel

Enter an IPv4 address for the Primary: 10.0.0.5

Primary Subnet Mask (blank for default): 255.255.255.0

Enter an IPv4 address for the Secondary: 10.0.1.6

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

Back Next Finish Cancel

Figure 37: Select Channel IP Addresses step

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**.

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

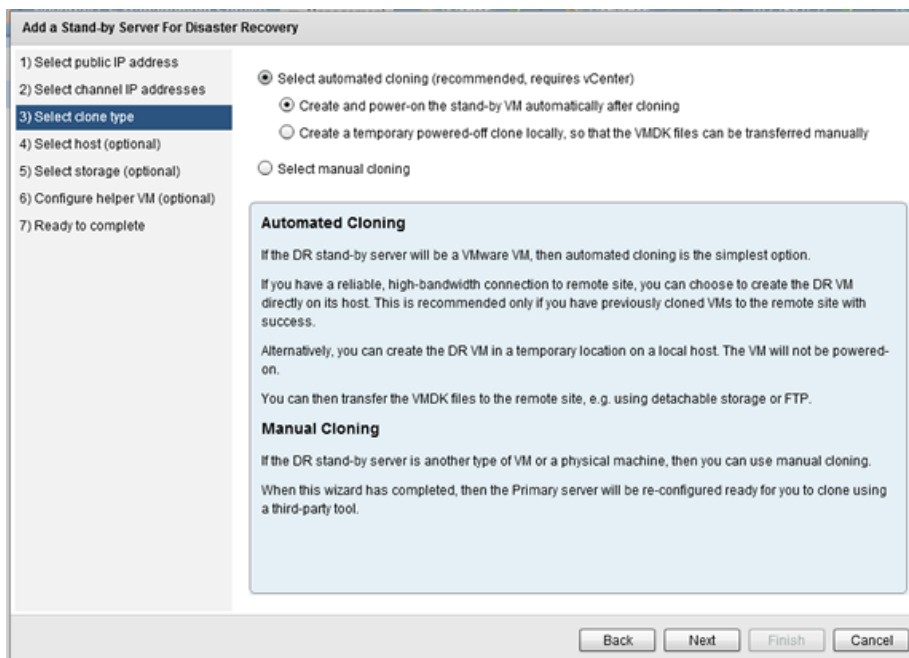


Figure 38: Select Clone Type step

The *Select Host* step is displayed.

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

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

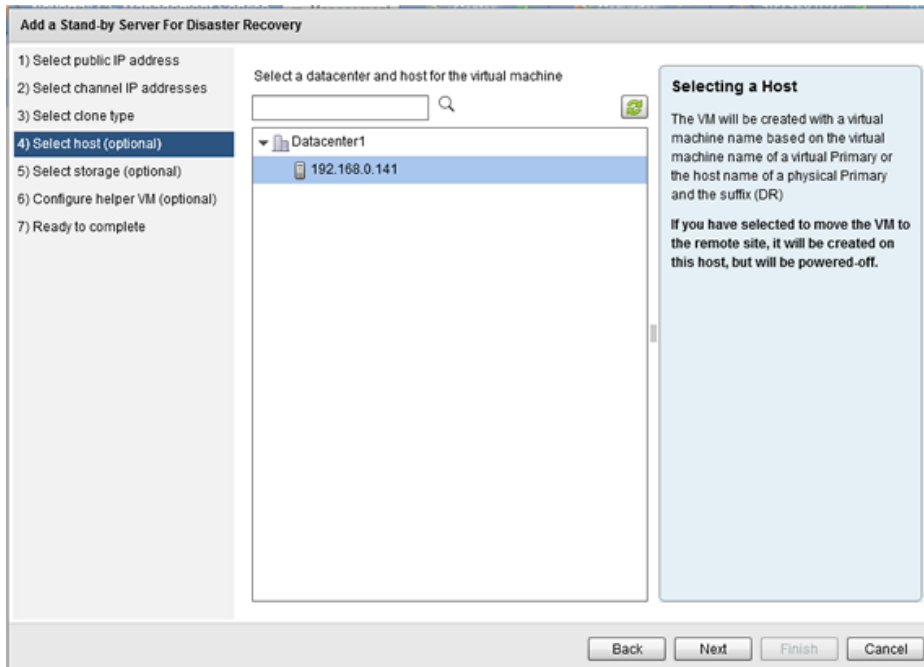


Figure 39: Select Host step

The *Select Storage* step is displayed.

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




Figure 40: Select Storage step

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

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

The screenshot shows a wizard window titled "Add a Stand-by Server For Disaster Recovery". On the left, a list of steps is shown, with "7) Ready to complete" selected and highlighted in blue. The main area contains a table of configuration details for the secondary server. To the right of the table is a text box titled "Ready to complete" with instructions. At the bottom right are four buttons: "Back", "Next", "Finish", and "Cancel".

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

Figure 41: Ready to Complete step

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.

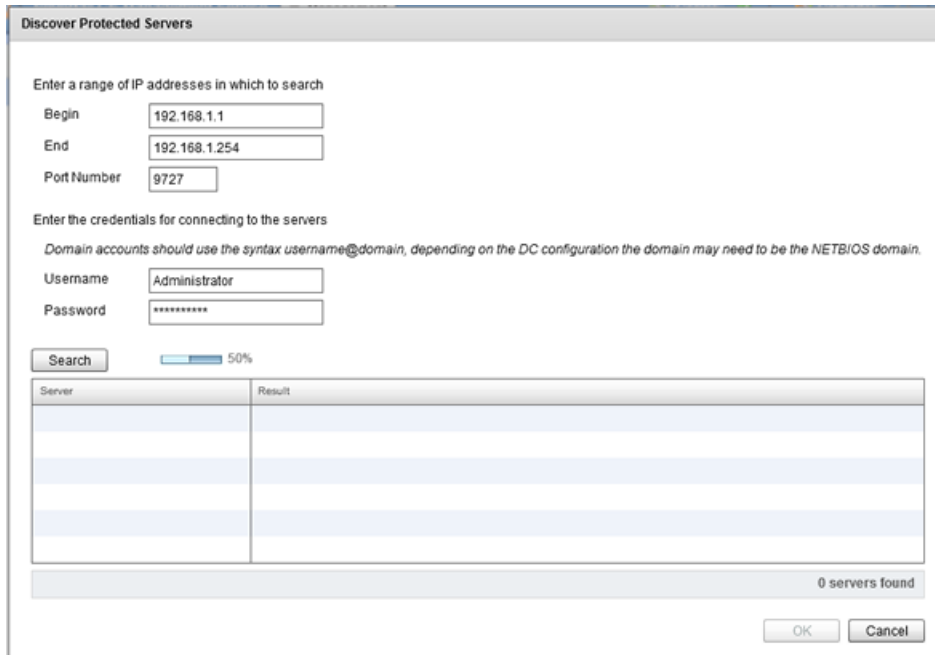
Discover Protected Servers

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

Procedure

To discover protected servers:

- From the **Management > Manage** drop-down pane, click **Discover Protected Servers**. The *Discover Server* dialog is displayed.



Discover Protected Servers

Enter a range of IP addresses in which to search

Begin

End

Port Number

Enter the credentials for connecting to the servers

Domain accounts should use the syntax username@domain, depending on the DC configuration the domain may need to be the NETBIOS domain.

Username

Password

50%

Server	Result

0 servers found

Figure 42: Discover Protected Servers dialog

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

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

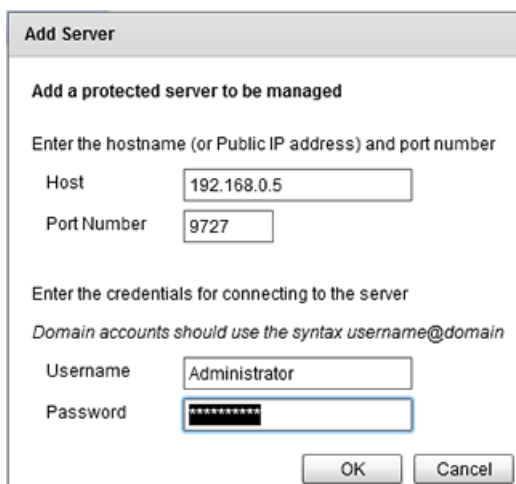
4. Click **Search** to run Neverfail Engine server discovery.
The Neverfail Continuity Engine Management Service displays all Neverfail Continuity Engine clusters discovered. Discovered items will be added automatically to the Protected Servers pane in the background.
5. Click **OK** or **Cancel** to dismiss the Discover Protected Servers dialog.

Add a Protected Server

Procedure

To add a protected server:

1. Neverfail Continuity Engine Management Service allows you to add individual protected servers which may be part of a cluster. Click **Add a Protected Server** in the **Management > Manage** drop-down pane to add a server.
The **Add Server** dialog is displayed.



Add Server

Add a protected server to be managed

Enter the hostname (or Public IP address) and port number

Host: 192.168.0.5

Port Number: 9727

Enter the credentials for connecting to the server

Domain accounts should use the syntax username@domain

Username: Administrator

Password: [REDACTED]

OK Cancel

Figure 43: Add Server dialog

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

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

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

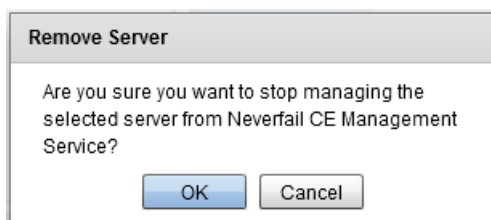
Remove the Selected Server

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

To remove the selected server:

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.



Remove Server

Are you sure you want to stop managing the selected server from Neverfail CE Management Service?

OK Cancel

Figure 44: Remove Server dialog

You are prompted to verify that you want to remove the selected server from management by the Neverfail Continuity Engine Management Service.

3. Click **OK**.

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

Download the Advanced Management Client

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

To download the Advanced Management Client:

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



Figure 45: Download Advanced Management Client

2. Select a target location for the downloaded file using the dialog navigation features.

3. Click **Save**.

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.

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.

Procedure

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 46: VMware vSphere

Create VMware SRM Plan Step for Selected Server

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.

Prerequisites

- 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
```

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.

- 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.

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

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

Alternate IP addresses:
- 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:
- 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.
- 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

Figure 47: Create SRM Plan Step

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

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

- 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.

Note: 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.

- 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*.
- 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.
- 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.

Note: 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.

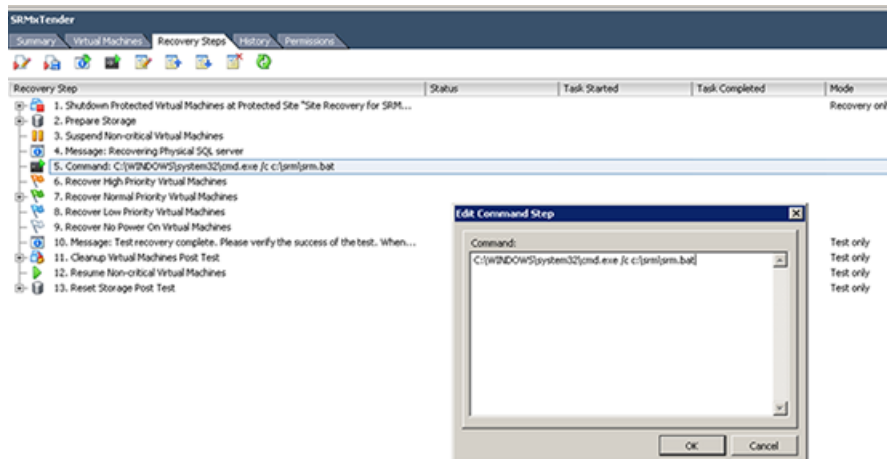


Figure 48: SRM Edit Command Step

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

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

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

15. Click **OK**.

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

License

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

Configure an Internet Proxy Server for Licensing

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.

Procedure

To configure for use with an internet proxy:

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

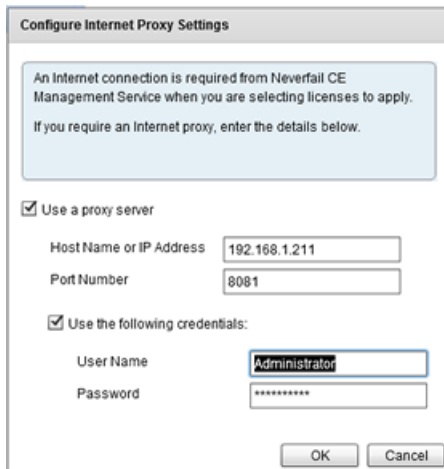


Figure 49: Configure Internet Proxy Settings

License the Selected Server

Licensing is performed via the Neverfail Continuity Engine Management Service.

To license Neverfail Engine:

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.

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

Apply License

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

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

Enter the email address you use for your Neverfail account
 ☒ Remember email address

Enter the password

☐ Manually enter a license key

Licensing Engine on NFCE81B (Signature: 4ZKS4FG3)

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.](#)

Figure 50: Apply License page

2. If there is an Internet connection from the Neverfail Continuity Engine Management Service, select *Apply a License from your Neverfail account*, enter your Neverfail credentials, press **Next** and continue from step 4.
3. 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 **Apply**. If the key is successfully applied, click **Finish**, otherwise review the error message.

Apply License

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

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

Enter the email address you use for your Neverfail account
 ☐ Remember email address

Enter the password

☒ Manually enter a license key

Licensing Engine on NFCE81B (Signature: 4ZKS4FG3)

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.](#)

Figure 51: Manual License Entry

4. 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**.
5. Review the *Ready to Complete* summary information and Click **Next**.
6. On the *Apply License* step, click **Finish**.

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.

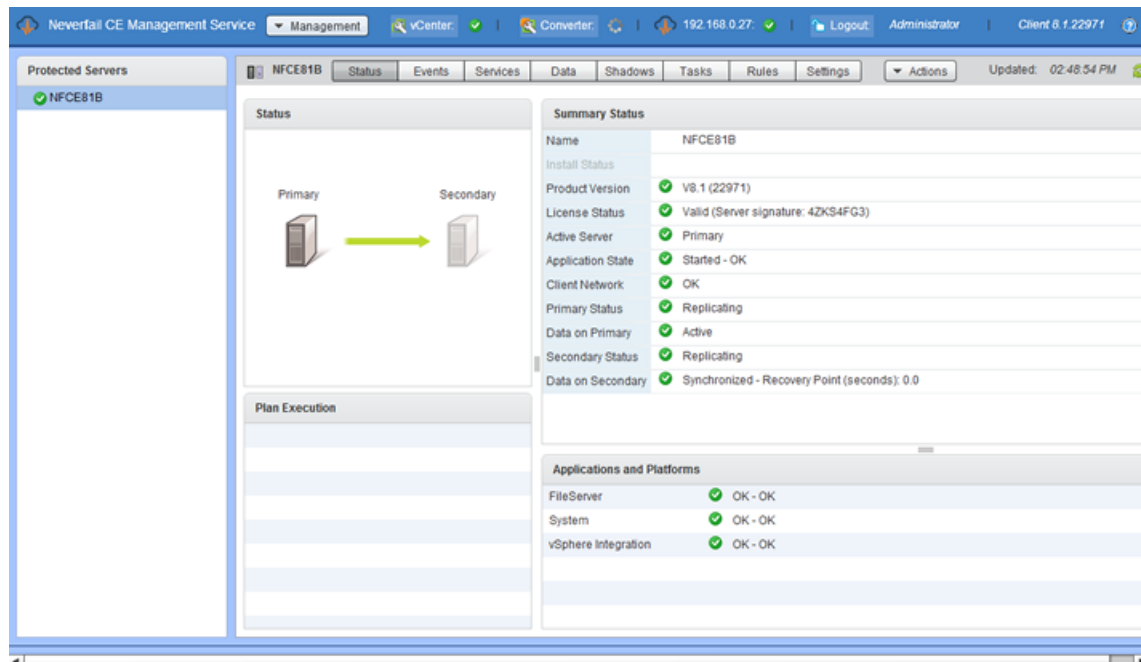


Figure 52: Summary Page

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).

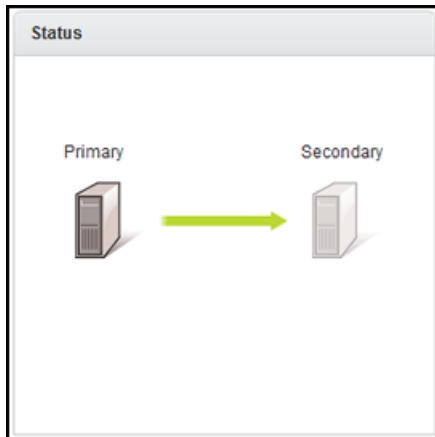


Figure 53: Status Pane

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.

Summary Status	
Name	NFV8.abcd.local
Install Status	
Product Version	✓ V8.0 (22727)
License Status	ⓘ Expires in 29 days (Server signature: YDCLYR4B)
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

Figure 54: Summary Status pane

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 55: Plan Execution pane

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.

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

Figure 56: Applications and Platforms

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.

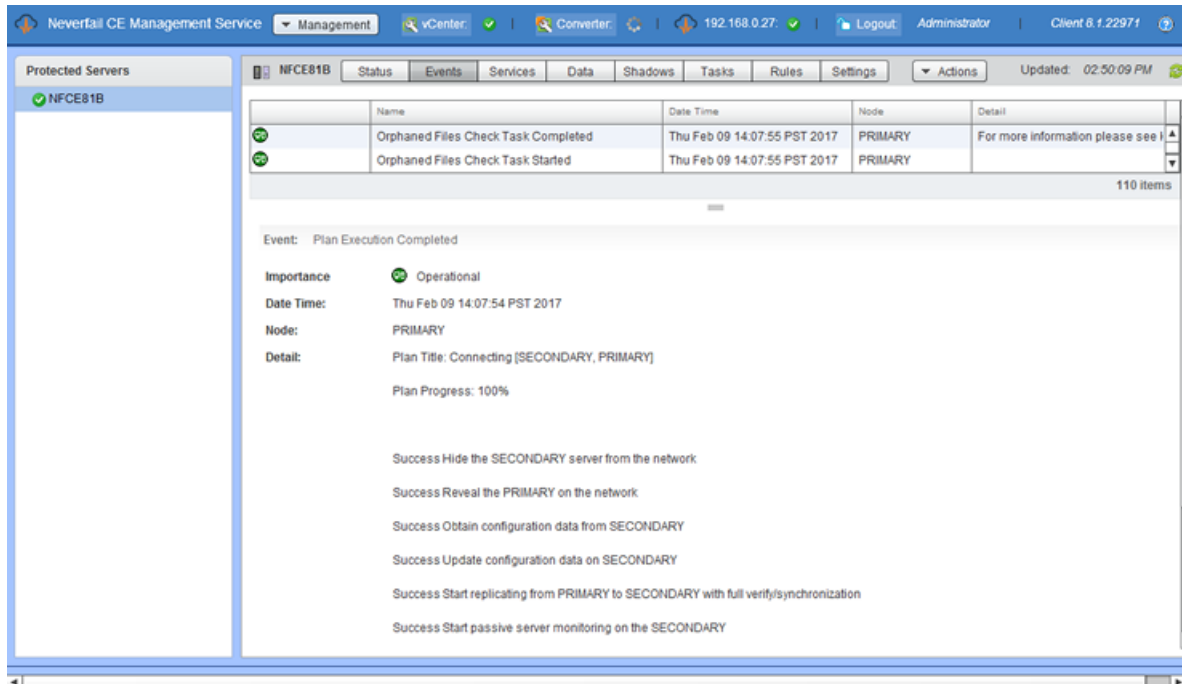


Figure 57: Events page

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.
	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.

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 of the server.

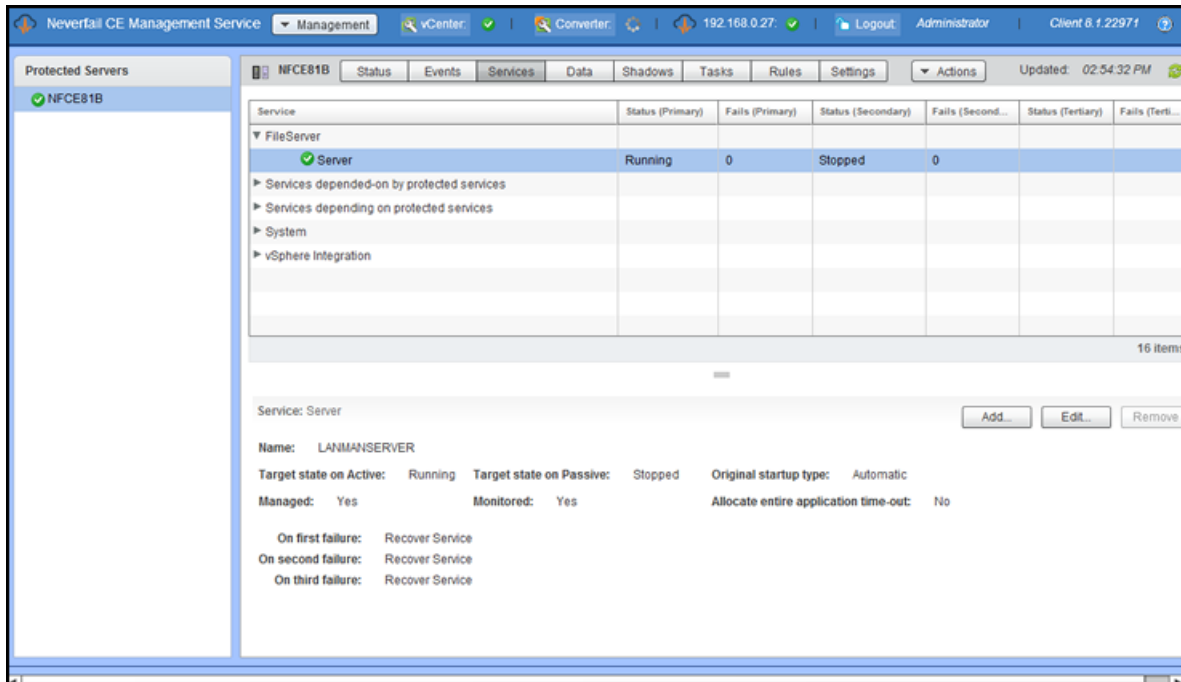


Figure 58: Applications: Services page

The target state of protected services can be specified for the active and passive server(s), and is typically *Running* on the active and *Stopped* on the passive(s). Services are protected when they are in a Running state in Engine Management Service or set to Automatic in Windows Services, and otherwise are logged as unprotected. 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).

Add a Service

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.

Procedure

To add a service:

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

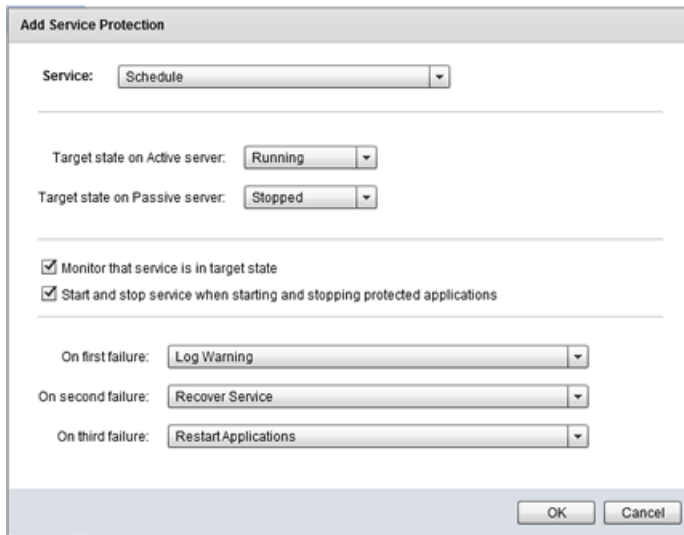


Figure 59: Add Service

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.

Note: *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.

Note: When dependent services are involved, actions to take on failure should match the protected service.

If a service fails and the failure option is set to Restart Applications, all applications are restarted.

Edit a Service

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

Procedure

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.

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.

The screenshot shows the 'Edit Service Protection' dialog box. At the top, it says 'Service: 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'. 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 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, there is a checked checkbox 'Allocate entire application time-out when recovering service' and two buttons: 'OK' and 'Cancel'.

Figure 60: Edit Service Protection

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**.

Configure Service Recovery Options for Protected Services

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**.

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

Figure 61: Edit Service Protection

Remove a Service

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

Procedure

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

- 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.

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**.

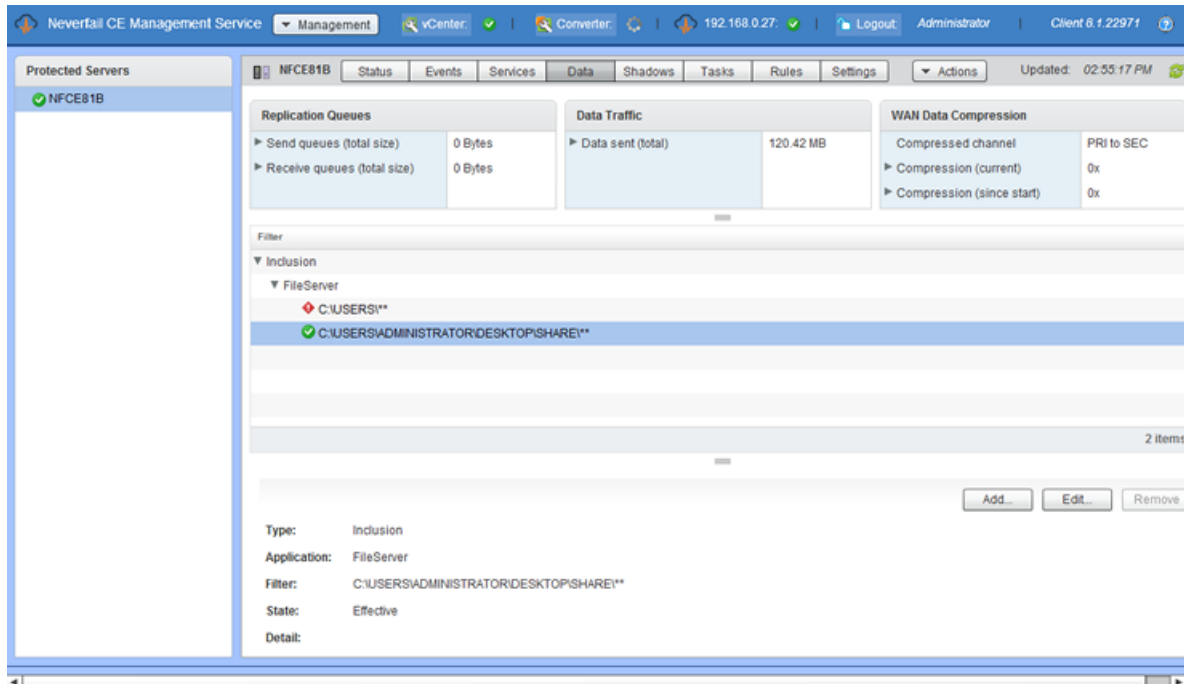


Figure 62: Data page

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.

Add Filters

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.

Procedure

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

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




Figure 63: 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.
3. Click **OK** to accept the changes, or **Cancel** to dismiss the dialog without making any changes.

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

Note: Neverfail 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.

Add an Exclusion Filter

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

Procedure

1. Filters to exclude files and folders from protection and replication are defined by clicking **Add** button on the **Data** page of the Neverfail Continuity Engine Management Service.




Figure 64: Add Exclusion Filter

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

Edit Filters

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

Procedure

To Edit a user defined Inclusion/Exclusion Filter:

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




Figure 65: 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**.
The file filter is changed and becomes active.

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

Remove Filters

Procedure

To Remove a user defined filter:

Note: *Plug-in filters can not be removed.*

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

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*.

Best Practices for Using Volume Shadow Copy Service & DRM

The Volume Shadow Copy Service (VSS) component of Windows 2008 and Windows 2012 takes shadow copies and allows you to configure the location and upper limit of shadow copy storage.

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

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

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

In accordance with the following guidelines from Microsoft:

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

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

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

Configure Shadow Creation Options

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

Procedure

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

- **Create a shadow every:**

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

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

If the check box is cleared, then no shadows are automatically created on the active. If it is selected, then a Shadow is taken each day at the time selected from the drop down list. The Shadow is taken with “application co-operation”, which means that if the application protected by Neverfail 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.

- **Only between the hours:**

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

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

- **Only on the days:**

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

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

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

Figure 66: Shadow Creation Options

Configure the Shadow Copy Schedule

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

Procedure

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

1. Navigate to the **Shadows** page and click **Configure**. The *Configure Shadow Schedule* dialog appears.

Figure 67: Configure Shadow Schedule

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

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

Note: *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.*

Configure Shadow Keep Options

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

Procedure

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

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

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

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

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

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

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

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

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

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

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

Figure 68: Shadow Keep Options

Manually Create Shadow Copies

Shadow Copies can be created manually using the steps below:

Procedure

- 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.

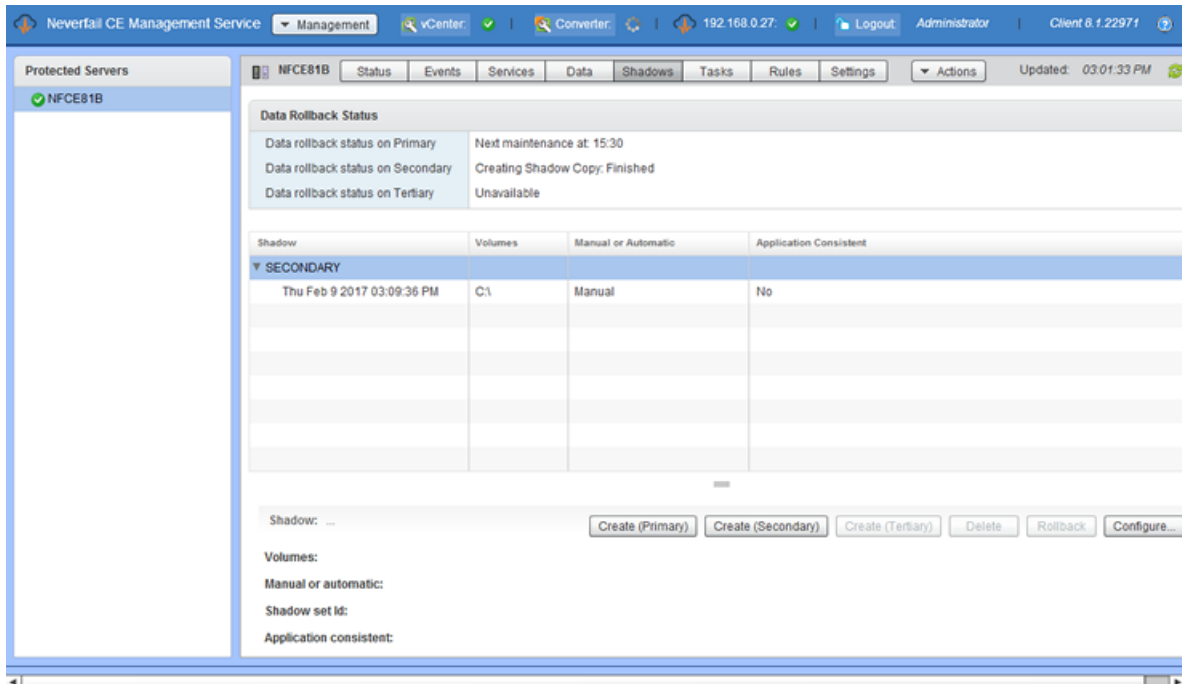


Figure 69: Create Shadow Dialog

Delete a Shadow Copy

Procedure

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

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

Roll Back Protected Data to a Previous Shadow Copy

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

Procedure

- 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**.
- 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.

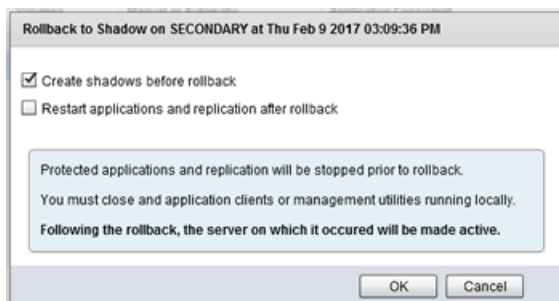


Figure 70: Rollback to Shadow dialog

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.

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.

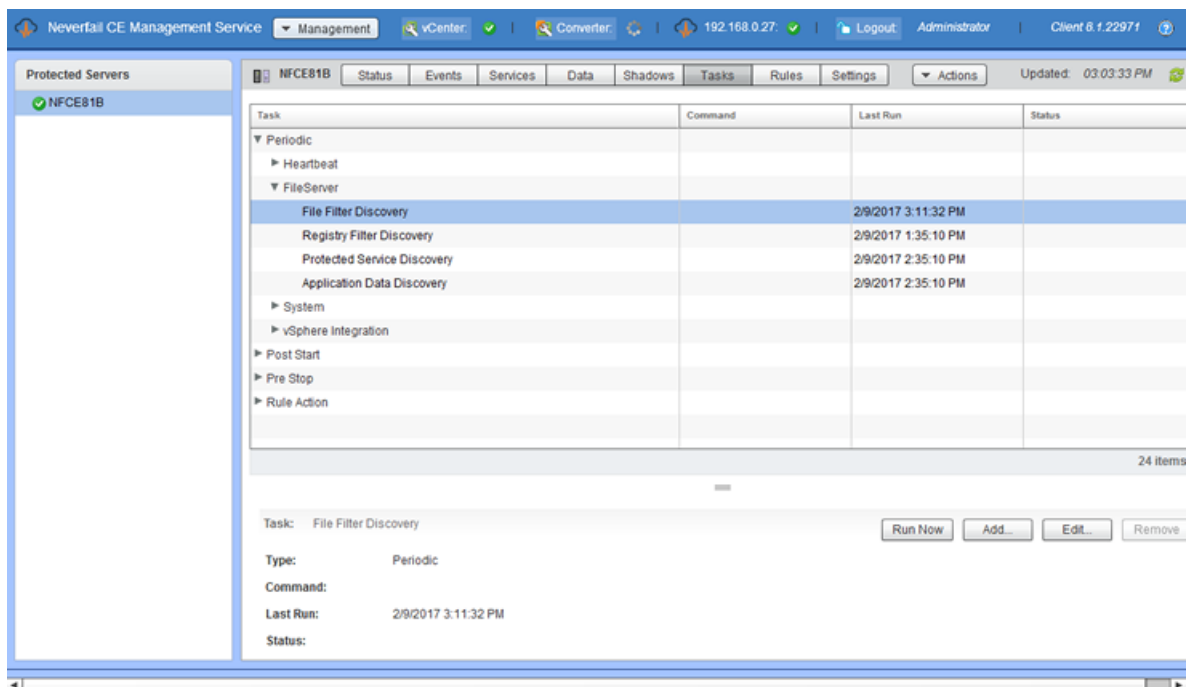


Figure 71: Tasks page

Run Now

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:

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.

Add Task

Tasks can be added from the Tasks page of the Neverfail Continuity Engine Management Service.

To add a User Defined task:

1. Click **Add** at the lower right of the pane. The *Add Task* dialog appears.

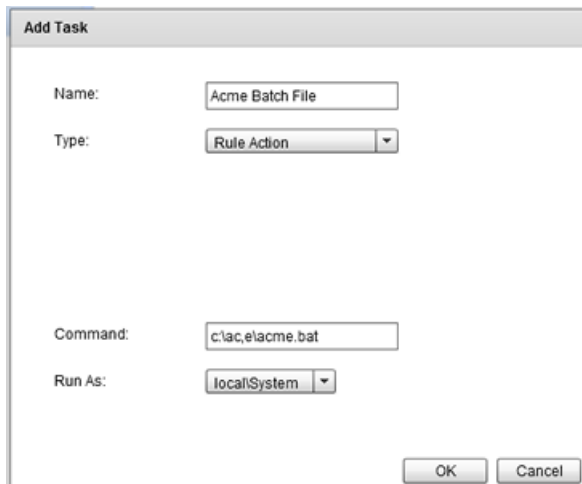
The image shows a Windows-style dialog box titled "Add Task". It contains four labeled fields: "Name:" with a text box containing "Acme Batch File"; "Type:" with a drop-down menu showing "Rule Action"; "Command:" with a text box containing "c:\acme\acme.bat"; and "Run As:" with a drop-down menu showing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

Figure 72: Add Task

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

Note: *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.

Note: *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.

Edit Task

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

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.

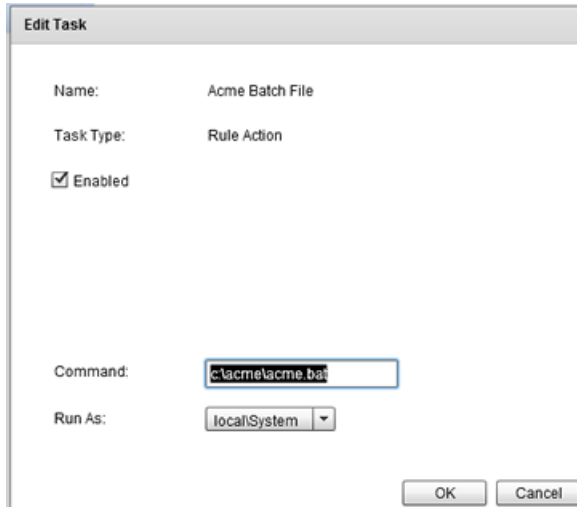


Figure 73: Edit Task

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

Remove Task

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

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

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.

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.

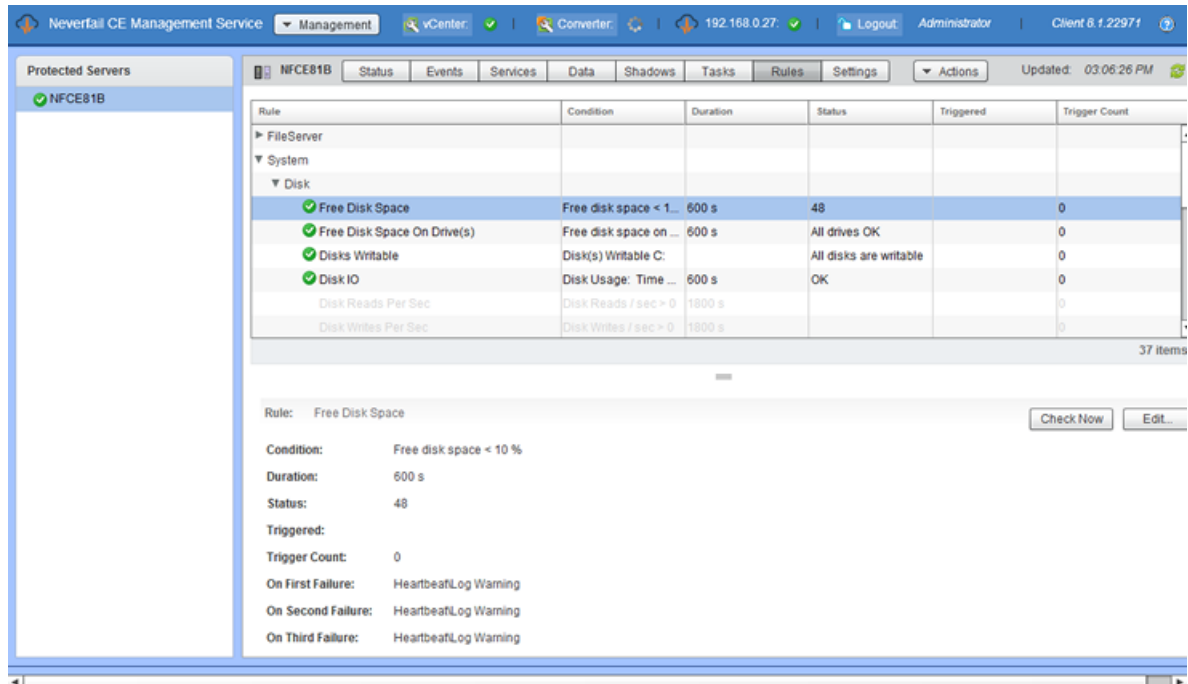


Figure 74: Rules page

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.

Edit a Rule

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

To Edit a rule:

1. 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 75: Edit Rule dialog

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
- **vSphere Integration\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.

Settings

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

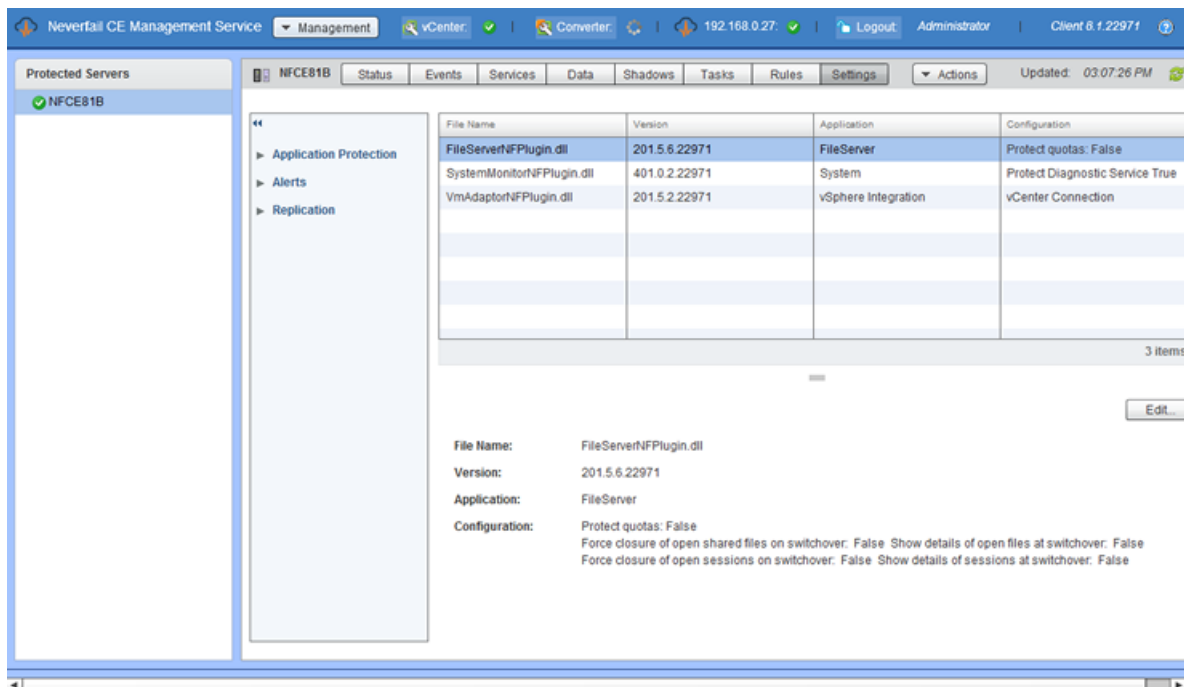


Figure 76: Settings page

Configure Plug-ins

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:

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

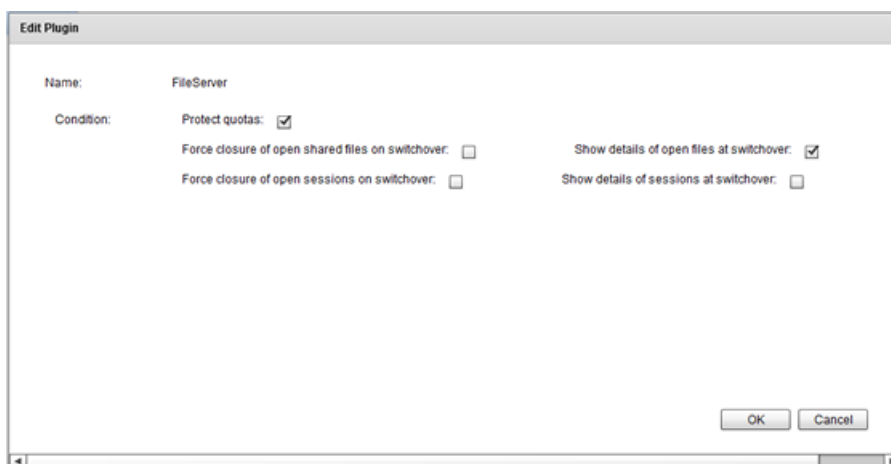


Figure 77: Edit Plug-in dialog

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.

Alert Settings

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.

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 78: Alert Settings

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.

Note: 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.
-

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.

Table 1: Neverfail Engine Variables

<i>Variables</i>	<i>Values</i>
\$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.

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.

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

Figure 79: Edit Alert Triggers

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.

Outgoing Mail Server for Primary Server	smtp1.acme.com
Outgoing Mail Server for Secondary Server	smtp2.acme.com
Outgoing Mail Server for Tertiary Server	
Send Mail As	engine@acme.com
Mail Server Requires Authentication	<input checked="" type="checkbox"/>
Username	administrator
Password	*****

OK Cancel

Figure 80: Email Settings

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. 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.

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 81: Edit WAN Compression dialog

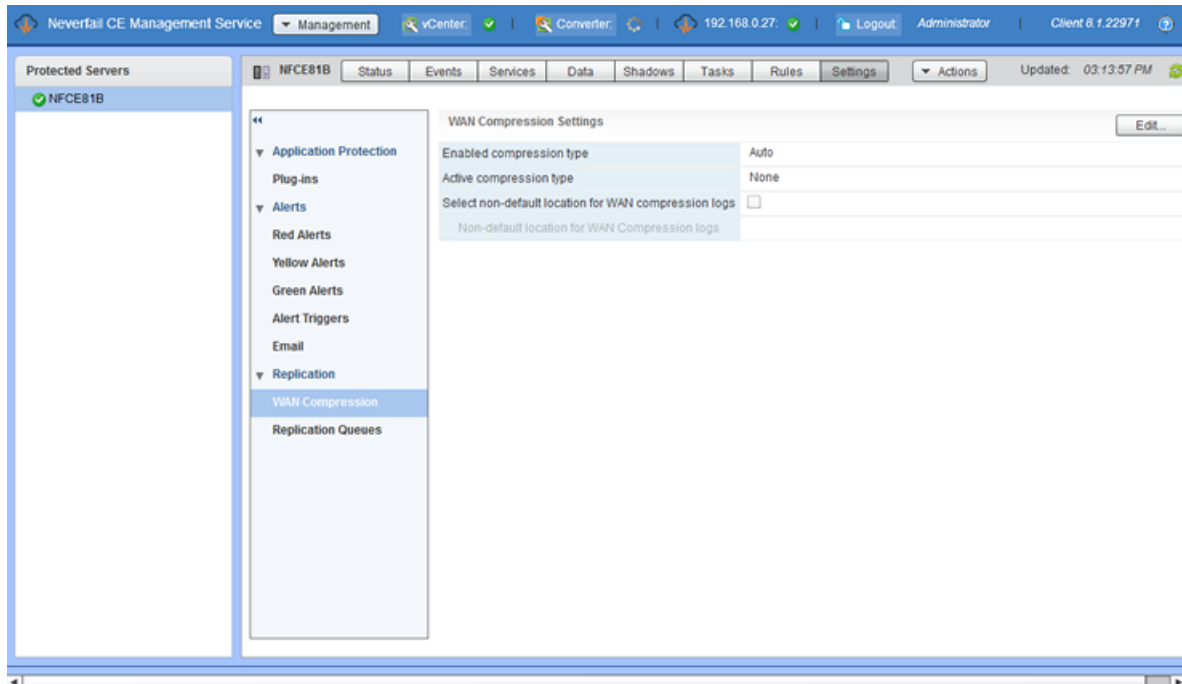


Figure 82: WAN Compression page

Replication Queue Settings

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

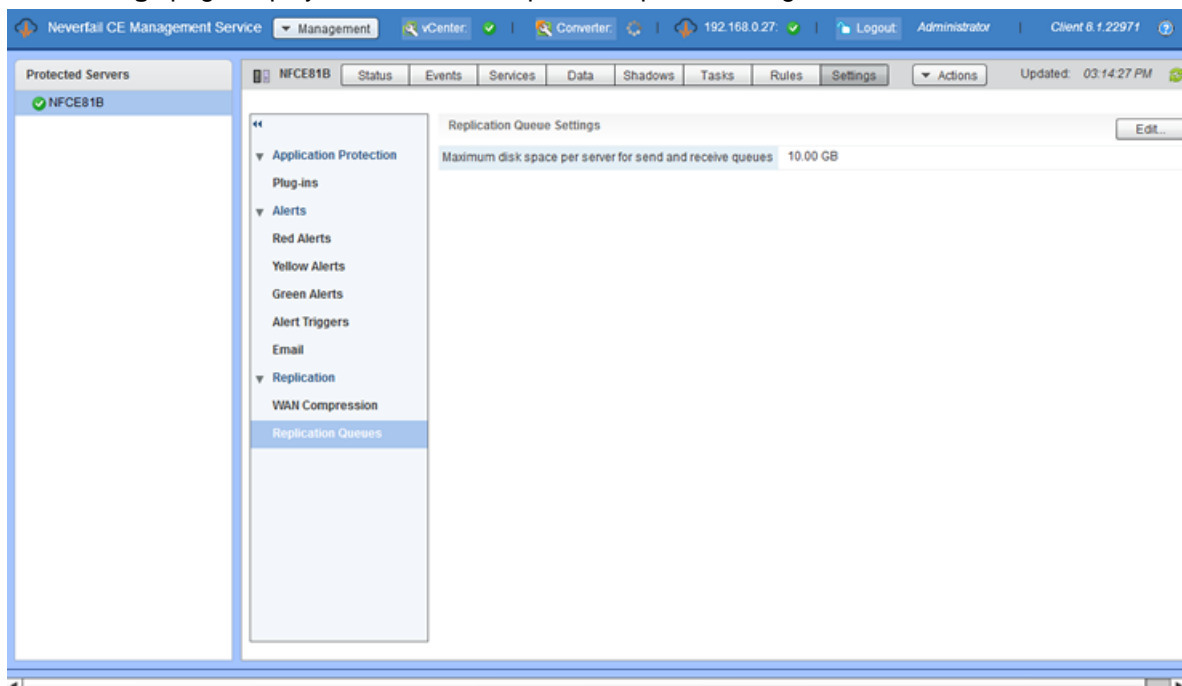


Figure 83: Configured Queue Size

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 84: Edit Replication Queue Settings dialog

Actions

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

The Neverfail Continuity Engine Management Service allows administrators to manage Neverfail Engine clusters similar to the Neverfail Advanced Management Client. The Neverfail Continuity Engine Management Service provides the ability to perform the main operations, comprising a Switchover, Start/Stop Replication, Start/Stop Applications, Create Shadows, Check file and registry system, and Startup/Shutdown of Neverfail Engine.




Figure 85: 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 86: 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.

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.

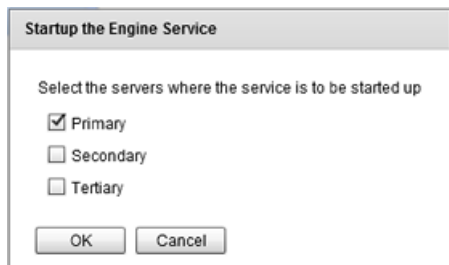


Figure 87: Startup Services

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 88: Shutdown

Post Installation Configuration

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

Configure the VmAdapter Plug-in

After installation of Neverfail Engine is complete:

Procedure

Configure the VmAdapter Plug-in:

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.

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

8. Click **OK**.

Adding an Additional Network Interface Card

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.

Procedure

To 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. Neverfail recommends that the NIC be added on the passive (Secondary) server, and then a switchover be performed making the Secondary server active, and then adding an additional NIC to the passive (Primary) server.

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

1. Shutdown Neverfail Engine on the passive server.
2. Navigate to **Start -> Control Panel -> Administrative Tools -> Services**.
3. Select the *Neverfail Engine Service* and change the *Start up* to *Manual*.
4. Add a virtual NIC to the Secondary server.
5. Restart the server.
6. Navigate to **Control Panel -> Network and Internet -> Network and Sharing -> Change Adapter Settings**.
7. Right-click the newly added NIC and select *Properties*.
8. Right-click the newly added NIC and select *Internet Protocol Version 4 (TCP/IPv4)* and click **Properties**.
9. Configure the NIC so that it does not use DHCP by temporarily entering an unused IP address (for example, 1.1.1.1).
10. Click **OK -> Ok -> Close**.
If the NIC is not enabled, enable it now.
11. 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.
12. Navigate to **Start -> Control Panel -> Administrative Tools -> Services**.
13. Select the *Neverfail Engine* service and change the *Start up* to *Automatic*.
14. Start Neverfail Engine on the passive (Secondary) server.
15. Perform a switchover to make the Secondary server active and the Primary server passive.
16. Shutdown Neverfail Engine on the (Primary) passive server.
17. Navigate to **Start -> Control Panel -> Administrative Tools -> Services**.
18. Select the *Neverfail Engine* service and change the *Start up* to *Manual*.
19. Add a virtual NIC to the Primary server.
20. Restart the server.
21. Right-click the newly added NIC and select *Properties*.
22. Select *Internet Protocol Version 4 (TCP/IPv4)* and click **Properties**.
23. Configure the NIC so that it does not use DHCP by temporarily entering a unused IP address (for example, 2.2.2.2).

24. Click **OK** -> **Ok** -> **Close**.
If the NIC is not enabled, enable it now.
25. 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.
26. Start Neverfail Engine on the passive (Primary) server.
27. Allow the server to synchronize. Once synchronized, perform a switchover.

Appendix

A

Installation Verification Testing

Testing a Neverfail Engine Pair

Important: The following procedure provides information about performing Installation Verification testing on a Neverfail Continuity Engine pair or trio to ensure proper installation and configuration. Additionally, this procedure provides step-by-step procedures to perform a controlled switchover in the event of an application failure and failover in the event of network or hardware failure resulting in excessive missed heartbeats.

Note: In this document, the term “Pair” refers to a Neverfail Engine pair. Refer to the for more information about Neverfail Engine Pairs.

Exercise 1 - Auto-switchover

Neverfail Continuity Engine monitors Neverfail services and the system environment to ensure that protected services are available for end users. To monitor services and the system environment, Neverfail Engine uses plug-ins which are designed for Neverfail services and the system.

If a protected service or the system begins to operate outside of preconfigured thresholds, Neverfail Engine can automatically switch to make the passive server the active server in the pair that provides services for end users.

Important: These exercises are examples and should be performed in order. Neverfail recommends against attempting to test failover on a properly operating pair by methods such as unplugging a power cord. At the moment power is lost, any data not written to the passive server is lost. Neverfail recommends that all actions intended to verify operation of the passive server be performed as a switchover rather than a failover.

Starting Configuration

Prior to initiating the Installation Verification process in a pair, Neverfail Engine must be configured with the Primary server as active and the Secondary server as passive. Additionally, the following prerequisites must be met:

- The Secondary server must be synchronized with the Primary server.
- All protected services must be operating normally.
- If installed in a LAN environment, using the Neverfail Advanced Management Client, verify that *Failover from Primary server to Secondary server if channel heartbeat is lost for failover timeout* is selected from the **Server: Monitoring > Configure Failover** dialog (default setting).
- If installed in a WAN environment, using the Neverfail Advanced Management Client, you must manually select *Failover from Primary server to Secondary server if channel heartbeat is lost for failover timeout* in the **Server: Monitoring > Configure Failover** dialog.

Important: Prior to starting the Installation Verification process, ensure that a known good backup of the Primary server exists and examine the Windows event logs for recent critical errors.

Neverfail provides an executable, `nfavt.exe`, to emulate conditions that result in auto-switchover so you can verify that your Neverfail Engine installation performs as expected. This section guides you through the steps necessary to perform this verification.

Steps to Perform

Important: If you encounter errors and or find it necessary to back out the changes made by this exercise, you can stop at any point and perform the steps described in the [Back-out Procedure \(Auto-switchover\)](#) to return the Pair to its original operating configuration and state.

Table 2: Perform the following procedure to verify Auto-Switchover in a Pair configuration.

Machine ID	Activity	Results
Primary	Open a command prompt.	
	Change directory to C:\Program Files\Neverfail\R2\Bin	
	Execute <code>nfavt.exe</code> When prompted, "Are you sure you wish to continue", click Continue .	Service is switched to the Secondary server and Neverfail Engine shuts down on the Primary.
Secondary	Login to the Engine Management Service	
	In the <i>Status</i> pane of the Engine Management Service, review the status of the server pair.	The <i>Status</i> pane indicates that the Secondary server is active.
	Verify all protected applications have started on the Secondary.	Services are running on the Secondary.
	Verify data is present.	Data is present.

Successful completion of this procedure leaves the Neverfail Engine pair in the state necessary to perform the second part of the Installation Verification process, detailed in [Exercise 2 - Data Verification](#).

Back-out Procedure (Auto-switchover)

Important: Do not perform this back-out procedure if you intend to continue the Installation Verification process.

If for any reason you find it necessary to back out of this exercise, you can stop at any point and return the pair to the state it was in at the beginning of this exercise by performing the following steps:

1. Shut down Neverfail Engine and protected services on all servers.
2. Complete the following on both servers:
 - a. Open the *Configure Server* wizard.
 - b. Select the *Machine* tab.
 - c. Select the *Primary* server as active.
 - d. Click **Finish**.
3. On the Secondary server, right-click the taskbar icon and select *Start Neverfail Engine*.
4. Verify that the Secondary server is passive (**S/-**).
5. On the Primary server, right-click the taskbar icon and select *Start Neverfail Engine*.
6. After Neverfail Engine starts, login to the Engine Management Service.
7. Verify that applications have started and replication to the passive server has resumed.

Exercise 2 - Data Verification

The Data Verification exercise validates that data is synchronized between the servers resulting in current data on the active server following the Auto-switchover exercise performed previously. The objective is to take a working active server (the Secondary server) and synchronize it with the passive (Primary server). This exercise also demonstrates that all the correct services stopped when the Primary server became passive.

Starting Configuration

Neverfail Engine is running on the Secondary active server. Login to the Secondary server and using the *System Tray* icon, verify that the server status displays **S/A**. Neverfail Engine is not running on the Primary server which is set to passive. Login to the Primary server and using the *System Tray* icon, verify that the server status displays **-/-** to indicate that Neverfail Engine is not running.

Steps to Perform

Table 3: Perform the following steps to verify that data is synchronized following Auto-switchover in a Pair configuration.

Activity	Results
On the Primary server, right-click the taskbar icon and select <i>Start Neverfail Engine</i> .	Neverfail Engine successfully starts.
Login to the Engine Management Service.	
In the <i>Protected Servers</i> pane of the Engine Management Service, select the server pair.	The <i>Summary</i> screen is displayed.
Review the <i>Status</i> pane and verify the connection from the Secondary (active) to Primary (passive).	The <i>Status</i> pane shows a connection from the Secondary server to the Primary server.

Activity	Results
View the <i>System Summary</i> pane and wait for both the <i>File System</i> and the <i>Registry</i> status to display as <i>Synchronized</i> . Access the Neverfail Engine logs and confirm that no exception errors occurred during the synchronization process.	Data replication resumes from the Secondary server back to the Primary server. Both the <i>File System</i> & <i>Registry</i> status become <i>Synchronized</i> .

Successful completion of this procedure leaves the Neverfail Engine Pair in the state necessary to perform the final part of the Installation Verification process, detailed in [Exercise 3 - Switchover](#).

Exercise 3 - Switchover

The Switchover exercise demonstrates the ability to switch the functionality and operations of the active server on command to the other server in the pair using the Neverfail Engine. Perform this exercise only after successfully completing the Auto-switchover and Data Verification Exercises.

Starting Configuration

Neverfail Engine is running on the Secondary active server and Primary passive server. Using the Engine Management Service *Summary* page, verify that the Secondary server is active and that the Primary server is passive.

Steps to Perform

Table 4: Perform the following steps to switch functionality and operations on command from the active server to the ready-standby server.

Activity	Results
Using the Engine Management Service, review the <i>Summary</i> pane to verify that both the <i>File System</i> and <i>Registry</i> status are <i>Synchronized</i> .	
Navigate to the Actions drop-down and click on Make Primary Server Active .	The Engine Management Service <i>Summary Status</i> pane displays the applications stopping on the active server. Once all applications are stopped, the active server becomes passive and the passive server becomes active. The <i>Summary Status</i> pane shows the applications starting on the newly active server. Both the <i>File System</i> and <i>Registry</i> status are <i>Synchronized</i> .
Confirm application performance and availability meets previously defined criteria. Verify that client applications are running as expected after the switchover process.	Services continue to be provided as before the switchover occurred. You may need to refresh or restart some client applications as a result of a switchover.

Successful completion of this procedure indicates a successful outcome from the Installation Verification process.

Testing a Neverfail Engine Trio

Important: The following procedure provides information about performing Installation Verification testing on a Neverfail Continuity Engine trio to ensure proper installation and configuration. Additionally, this procedure provides step-by-step procedures to perform a controlled switchover in the event of

an application failure and failover in the event of network or hardware failure resulting in excessive missed heartbeats.

Note: *In this document, the term "Trio" refers to a Neverfail Engine trio. Refer to the [Glossary](#) for more information about Neverfail Engine trios.*

Exercise 1 - Auto-switchover

Neverfail Continuity Engine monitors services and the system environment to ensure that protected services are available for end users. To monitor services and the system environment, Neverfail Engine uses plug-ins which are designed for Neverfail services and the system.

If a protected service or the system begins to operate outside of preconfigured thresholds, Neverfail Engine can automatically switch to and make active the passive server in the pair to provide services for end users.

Important: *These exercises are examples and should be performed in order. Neverfail recommends against attempting to test failover on a properly operating Cluster by methods such as unplugging a power cord. At the moment power is lost, any data not written to the passive server is lost. Neverfail recommends that all actions intended to verify operation of the passive server be performed as a switchover rather than a failover.*

Neverfail provides an executable, `nfavt.exe`, to emulate conditions that result in auto-switchover so you can verify that your Neverfail Continuity Engine installation performs as expected. This section guides you through the steps necessary to perform this verification.

Starting Configuration

Prior to initiating the Installation Verification process in a Trio, Neverfail Engine must be configured with the Primary server as active, the Secondary server as 1st passive, and the Tertiary server as 2nd passive. All servers must be synchronized with the Primary server, and all protected applications must be operating normally.

Important: *Prior to starting the Installation Verification process, ensure that a known good backup of the Primary server exists and examine the Windows event logs for recent critical errors.*

Neverfail provides an executable, `nfavt.exe`, to emulate conditions that result in auto-switchover so you can verify that your Neverfail Engine installation performs as expected. This section guides you through the steps necessary to perform this verification.

Prior to initiating this procedure, download `nfavt.exe` from the Neverfail to
<installation_location>\Neverfail\R2\Bin

Steps to Perform

Important: *If you encounter errors and or find it necessary to back out the changes made by this exercise, you can stop at any point and perform the steps described in the [Back-out Procedure \(Auto-switchover\)](#) to return the Pair to its original operating configuration and state.*

Table 5: Perform the following procedure to verify Auto-Switchover in a Pair configuration.

Machine ID	Activity	Results
Primary	Open a command prompt.	
	Change directory to C:\Program Files\Neverfail\R2\Bin	
	Execute <code>nfa_vt .exe</code> When prompted, "Are you sure you wish to continue", click Continue .	Service is switched to the Secondary server and Neverfail Engine shuts down on the Primary.
Secondary	Login to the Engine Management Service.	
	In the <i>Servers</i> pane of the Engine Management Service, select the server Cluster.	The <i>System Overview</i> screen indicates that the Secondary server is active.
	Verify all protected applications have started on the Secondary.	Services are running on the Secondary.
	Verify data is present and is replicating to the Tertiary server.	Data is present and replicating.
Tertiary	Verify that the Tertiary server is passive and in-sync	The <i>System Overview</i> page indicates that the Tertiary server is passive and in-sync

Successful completion of this procedure leaves the Neverfail Engine trio in the state necessary to perform the second part of the Installation Verification process, detailed in [Exercise 2 - Managed Switchover](#).

Back-out Procedure (Auto-switchover)

Important: Do not perform this back-out procedure if you intend to continue the Installation Verification process.

If for any reason you find it necessary to back out of this exercise, you can stop at any point and return the Cluster to the state it was in at the beginning of this exercise by performing the following steps:

1. Shut down Neverfail Engine and protected services on all servers.
2. Complete the following on all three servers:
 - a. Open the *Configure Server* wizard.
 - b. Select the *Machine* tab.
 - c. Select the *Primary* server as active.
 - d. Click **Finish**.
3. On the Secondary and Tertiary servers, right-click the taskbar icon and select *Start Neverfail Engine*.
4. Verify that the Secondary and Tertiary servers are passive (**S/-** and **T/-**).
5. On the Primary server, right-click the taskbar icon and select *Start Neverfail Engine*.
6. After Neverfail Engine starts, login to the Engine Management Service.
7. Verify that applications have started and replication to the passive server has resumed.

Exercise 2 - Managed Switchover

Neverfail Engine provides manual control over switching the active server role to another server in the Cluster. On command, Neverfail Engine gracefully stops replication and the protected applications on

the currently active server and then starts the protected applications and replication on the server selected to assume the active role.

Use this exercise to validate seamless switching of the active server role to another server in the Cluster. At the end of this section are instructions on how to back out of the exercise (such as if errors are encountered) and return the Cluster to its original operating configuration and state.

Starting Configuration

Neverfail Engine is running on the Secondary active server (**S/A**) and Tertiary server (**T/-**). Neverfail Engine is not running on the Primary server (**-/-**)

Steps to Perform

Important: If you encounter errors and or find it necessary to back out the changes made by this exercise, you can stop at any point and perform the steps described in the Back-out Procedure (Managed Switchover) below to return the Cluster to its original operating configuration and state.

Table 6: Perform the following steps to verify Managed Switchover in a Trio configuration.

Machine ID	Activity	Results
Secondary	Login to the Engine Management Service.	
	Click Rollback .	The <i>Rollback</i> screen is displayed.
	Under <i>Shadows</i> , click Create . In the <i>Create Shadow</i> dialog, select <i>Secondary</i> , and then click OK .	A rollback point is created prior to testing Secondary to Tertiary switchover.
	In the <i>Servers</i> pane of the Engine Management Service, select the server Cluster.	The <i>System Overview</i> screen is displayed.
	In the <i>System Overview</i> page, select the Tertiary server and then click Make Active .	Neverfail Engine performs a managed switchover to the Tertiary server and makes the Tertiary server active.
Tertiary	Login to the Engine Management Service.	
	In the <i>Servers</i> pane of the Engine Management Service, select the server Cluster.	The <i>System Overview</i> screen is displayed.
	Verify that all protected applications have started.	Services are running on the Tertiary server.
	Verify that data is present and replicating to the Secondary server.	Data is present and replicating.
Secondary	Verify that the Secondary server is passive and in-sync.	The <i>System Overview</i> screen indicates that the Secondary server is passive and in sync.

Successful completion of this procedure leaves the Cluster in the state necessary to perform the third part of the Installation Verification process, detailed in [Exercise 3 - Data Verification](#).

Back-out Procedure (Managed Switchover)

Important: Do not perform this back-out procedure if you intend to continue the Installation Verification process.

If for any reason you find it necessary to back out of this exercise, you can stop at any point and return the Cluster to the state it was in at the beginning of this exercise by performing the following steps:

1. Shut down Neverfail Engine and protected applications on the Secondary and Tertiary servers.
2. Complete the following on the Tertiary server:
 - a. Open the *Configure Server* wizard.
 - b. Select the *Machine* tab.
 - c. Select the *Secondary* server as active.
 - d. Click **Finish**.
 - e. Right-click the taskbar icon and select *Start Neverfail Engine*.
 - f. Verify that the Tertiary server is passive (T/–) and then shut down Neverfail Engine.
3. On the Secondary, right-click the taskbar icon and select *Start Neverfail Engine*.
4. After Neverfail Engine starts on the Secondary server, login to the Engine Management Service.
5. Click **Rollback**.
6. Under *Shadows*, select the previously created shadow on the Secondary server and click **Rollback**.
7. The *Rollback Shadow* dialog is displayed. Select *Restart applications and replication automatically after rollback*, and then click **OK**.
8. The *Rollback Status & Control* dialog is displayed. Click **Yes**.
9. Once the rollback is complete, verify applications have started and are operating as expected.
10. On the Tertiary server, right-click the taskbar icon and select *Start Neverfail Engine*.
11. Verify that replication to the passive server has resumed.

Exercise 3 - Data Verification

The Data Verification exercise validates that data is synchronized between the servers resulting in current data on the active server following a Managed Switchover. The objective is to take a working active server (the Secondary server) and synchronize it with the passive (Tertiary server).

Starting Configuration

Neverfail Engine is running on the Secondary and Tertiary servers. Using the *System Tray* icon, verify that the server status displays **S/A**. Neverfail Engine is not running on the Primary server which is set to passive. Using the *System Tray* icon, verify that the server status displays **-/-** to indicate that Neverfail Engine is not running.

Important:

If you encounter errors and or find it necessary to back out the changes made by this exercise, you can stop at any point and perform the steps described in the [Back-out Procedure \(Data Verification\)](#) below to return the Cluster to its original operating configuration and state.

Steps to Perform

Table 7: Perform the following steps to verify that data is synchronized following Managed Switchover in a Trio configuration.

Machine ID	Activity	Results
Primary	Right-click the taskbar icon and select <i>Start Neverfail Engine</i>	Neverfail Engine successfully starts.

<i>Machine ID</i>	<i>Activity</i>	<i>Results</i>
	Login to Engine Management Service.	
	In the <i>Servers</i> pane of the Engine Management Service, select the server Cluster.	The <i>System Overview</i> screen is displayed.
	Click on the Primary server icon to select the <i>Primary</i> server and verify that it is in a synchronized state.	Ensure that the full system check is complete.
Tertiary	Login to the Engine Management Service.	
	Click Rollback .	The <i>Rollback</i> screen is displayed.
	Under <i>Shadows</i> , click Create . In the <i>Create Shadow</i> dialog, select <i>Tertiary</i> , and then click OK .	A rollback point is created prior to testing Tertiary to Primary switchover.
Primary	In the <i>System Overview</i> screen, select the <i>Primary</i> server and click Make Active.	Neverfail Engine performs a managed switchover to the Primary server and makes the Primary server active.
	Verify that all protected applications have started.	Services are running on the Primary server.
	Verify that data is present.	Data is present on the Primary server and is synchronized.
	Verify that all three servers are connected and replicating.	

Successful completion of this procedure indicates a successful outcome from the Installation Verification process.

Back-out Procedure (Data Verification)

Important: Do not perform this back-out procedure if you intend to continue the Installation Verification process.

If for any reason you find it necessary to back out of this exercise, you can stop at any point and return the Cluster to the state it was in at the beginning of this exercise by performing the following steps:

1. Shut down Neverfail Engine and protected applications on all servers.
2. Complete the following on the Primary and Secondary servers:
 - a. Open the *Configure Server* wizard.
 - b. Select the *Machine* tab
 - c. Select the *Tertiary* server as active.
 - d. Click **Finish**.
 - e. Right-click the taskbar icon and select *Start Neverfail Engine*.
 - f. Verify that the Primary and Secondary servers are passive (**P/–** and **S/–**).

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 offsite facility, and replicating the data through a WAN link.

DNS (Domain Name System) Server

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

Domain

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

Domain Controller (DC)

The server responsible for maintaining privileges to domain resources; sometimes called AD controller in Windows 2003 and above domains.

Dualed

A way to provide higher reliability by dedicating more than one NIC for the Neverfail Channel on each server.

Failover

Failover is the process by which the 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.