



Installation Guide

Neverfail Continuity Engine

Notice

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

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

Copyright (c) 2024 Neverfail, LLC. All rights reserved.

Contents

Introduction

Neverfail Continuity Engine Concepts

Communications

Continuity Engine Switchover and Failover Processes

Implementation

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

Anti-Malware Recommendations

Installing Neverfail Continuity Engine

Installing Continuity Engine Management Service

Deploying Continuity Engine on the Primary Server

Automated Deployment of Stand-by Servers with Automatic Cloning

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

Manually installing Continuity Engine without using Continuity Engine Management Service

Post Installation Configuration

Installation Verification Testing

Testing a Continuity Engine Pair

Testing a Continuity Engine Trio

Glossary

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:

- **About This Book** (this chapter) provides an overview of this guide and the conventions used throughout.
- **Introduction** presents an overview of Neverfail Continuity Engine concepts including the Switchover and Failover processes.
- **Implementation** discusses supported platforms 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.
- **Installing Neverfail Continuity Engine** describes the installation process, guides you through installation on the Primary, Secondary, and Tertiary (if deployed) servers, and through post-installation configuration.
- **Installation Verification Testing** provides a quick, simple procedure to verify that Neverfail Continuity Engine is properly installed and initially configured.
- **Glossary** provides definitions for terms used in this document.

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 <https://www.neverfail.com/services-and-support/>.

Online and Telephone Support

Use online support located at <https://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 <https://www.neverfail.com/services-and-support/> .

Neverfail Professional Services

Neverfail Professional Services courses offer extensive hands-on labs, case study examples, and course materials designed for use as on-the-job reference tools. Courses are available on site, in the classroom, and live online. For the day-to-day operations of Neverfail Continuity Engine, Neverfail Professional Services provides offerings to help you optimize and manage your Neverfail Continuity Engine servers. To access information about education classes, certification programs, and consulting services, go to <https://www.neverfail.com/services-and-support/> .

Neverfail Continuity Engine Documentation Library

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

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

Conventions

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

Convention	Specifying
Bold	Window items including buttons.
<i>Italics</i>	Book and CD titles, variable names, new terms, and field names.
Fixed font	File and directory names, commands and code examples, text typed by you.

Convention	Specifying
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.

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.

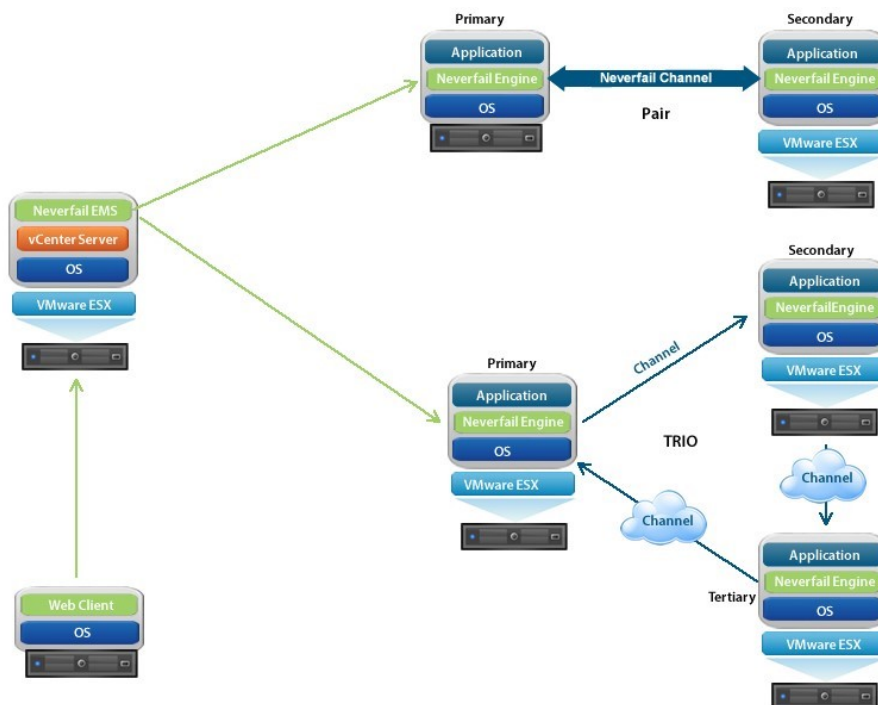
- **Neverfail Continuity Engine Concepts**
- **Communications**
- **Neverfail Engine Switchover and Failover Processes**

Neverfail Continuity Engine Concepts

Overview

Neverfail Continuity Engine consists of the Engine Management Service that is used to deploy and manage the Continuity 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 Continuity Engine with the ability to view Continuity Engine status and perform most routine Continuity Engine operations from a single pane of glass.



Neverfail describes the organization of Continuity Engine servers based upon Clusters, Cluster status, and relationships between Clusters. Neverfail refers to a Cluster of two servers as a Continuity Engine Pair or a Cluster of three servers as a Continuity Engine Trio. Installing Continuity Engine on the servers and assigning an identity to the servers results in a Continuity 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 Continuity Engine is deployed on a Pair or Trio of servers, Continuity 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, Continuity Engine operates as a Continuity 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, Continuity Engine can assist replication by utilizing WAN Compression with the built-in WAN Acceleration feature.

Architecture

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



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 defined in the same or a different subnet than the Public IP address or Neverfail Channel IP address subnet and is always available for administrators to access the server.

Continuity Engine Switchover and Failover Processes

Continuity 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).

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.

- **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**
- **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 Platforms

Neverfail Continuity Engine Management Service

- Windows Server 2016 Standard/Datacenter
- Windows Server 2019 Standard/Datacenter
- Windows Server 2022 Standard/Datacenter
- Desktop Edition of Windows OS 7, 8.x, 10 or 11.

Neverfail Continuity Engine

- Windows Server 2016 Standard/Datacenter
- Windows Server 2019 Standard/Datacenter
- Windows Server 2022 Standard/Datacenter
- Windows 10 IoT/Enterprise
- Windows 11 IoT/Enterprise

Unsupported Platforms

Neverfail Continuity Engine Management Service

- A server where Neverfail Continuity Engine is currently installed
- On an IA-64 Itanium Platform

Neverfail Continuity Engine

- A server where Engine Management Service is currently installed
 - 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

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 > Change Configuration**
- **Virtual Machine > Interaction > Configure CD Media**
- **Virtual Machine > Interaction > Connect Devices**
- **Virtual Machine > Interaction > Power On**
- **Virtual Machine > Interaction > Power Off**
- **Virtual Machine > Interaction > Reset**
- **Virtual Machine > Edit 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.
 1. Select the top level for vCenter Server, then click the **Permissions** tab.
 2. Right-click and select *Add Permission*.
 3. 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.

Engine Management Service

1. Engine Management Service installation is supported on the following operating systems:
 - Microsoft Windows 2016, 2019, 2022
 - Microsoft Windows Client Edition 7, 8.x, 10, and 11

Note: Connectivity with VMware vCenter Server is NOT required for deployment of Neverfail Continuity Engine but is recommended for fully automated deployments.

Important: If installing on Domain Controller Server, some additional steps are required (detailed in [Installing Neverfail Continuity Engine](#)).

Important: If installing on French Localized Windows Server, some additional steps are required (detailed in [Installing Neverfail Continuity Engine](#)).

2. A minimum of 512MB of RAM must be available for Neverfail Engine Management Web Services in addition to any other memory requirements for the Operating System or installed applications.
3. Engine Management Service requires 1GB of disk space for its files on the drive where it is to be installed.
4. 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).
5. For P2V installation, a supported version of VMware Converter must be available and configured prior to attempting installation of the Primary server.
6. Engine Management Client (EMS) supports most browsers used to connect to the EMS UI.
7. Engine Management Service requires elevated permissions in order to be installed.

Note: Engine Management Service will be configured to use **Neverfail Engine** 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).

Primary Server

1. Continuity Engine requires that Microsoft .Net Framework 4.0 or later be installed prior to installation.
2. Continuity Engine requires SMB2 file transfer protocol to be enabled prior to installation.
3. If the Primary server has a pending reboot, it must be resolved prior to the deployment of Continuity Engine on to the server.
4. Obtain and use local administrator permissions to perform Continuity Engine installation.
 - If *User Account Control: Run all administrators in Admin Approval Mode = 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 *User Account Control: Run all administrators in Admin Approval Mode = Disabled*, you may use any account with membership in the local Administrators group on the target server.

Note: Continuity Engine services are required to be run under the Local System account.

5. The server to be protected by Continuity Engine can NOT be any of the following:
 - A server running Engine Management Client
 - A server configured as a Domain Controller, Global Catalog, DHCP, or DNS

Note: These roles and services must be removed before proceeding with installation.

6. The Primary server can be Virtual or Physical with the Secondary and Tertiary server (if deployed) as either Virtual or Physical as well.
7. Verify that all services to be protected have all three *Recovery* settings configured to *Take no Action*.
8. Verify no other critical business applications except those to be protected by Continuity Engine are installed on the server.
9. Verify that 1GB RAM is available for *Neverfail Engine Service* and *Neverfail Engine Web Services* in addition to any other memory requirements for the Operating System or installed applications.
10. Verify that a minimum 2GB of free disk space is available on the drive where Continuity Engine is to be installed.

Note: Although Continuity Engine requires only 2 GB of available disk space on the drive to receive the Continuity 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.

11. Apply the latest Microsoft security updates and set Windows Updates to *manual*.
12. All applications that will be protected by Continuity Engine must be installed and configured on the Primary server prior to installing Continuity Engine.
13. Verify that all services to be protected are running or set to Automatic prior to installation.

Note: During installation, protected services are set to manual to allow Continuity 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.

14. Register this connection's address in DNS should be disabled.

Note: If enabled, Continuity Engine will disable Register this connection's address in DNS during installation. It is recommended to verify and configure (if needed) the DNS records on your DNS server.

15. File and Printer Sharing must be enabled and allowed access through all firewalls on the Primary target server prior to deployment.
16. Verify that the Server service is running prior to deployment to the target server.

Secondary Server

1. When installing in a P2V environment, the specifications of the Secondary Continuity Engine virtual machine must match the Primary physical server as follows
 - Similar CPU
 - Identical Memory
 - Sufficient disk space to host VM disks to match the Primary server
 - Same type/mode for Boot Firmware (UEFI, BIOS)

The Secondary Continuity Engine virtual machine must have sufficient priority in resource management settings so that other virtual machines do not impact its performance.

IP Addressing

1. IP Address requirements:
 - Public:
 - 1 each Public IP address - Engine Management Client
 - 1 each Public IP address - Primary Server
 - 1 each Public IP address - Secondary Server (only when deployed for DR)

Note: When deployed for HA or as part of a trio, the Primary and Secondary server will share the Public IP address.

- 1 each Public IP address - Tertiary Server (only when deployed in a trio)
 - Channel:
 - 1 each Channel IP address - per server when deployed in a pair
 - 2 each Channel IP addresses - per server when deployed in a trio

LAN

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

WAN

1. When deployed in a WAN environment, 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.

2. If the Primary and DR site uses the same subnet:

- 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:

- During installation, follow the steps for a WAN.
- The Primary and Secondary servers in the Continuity 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 DNSUpdate.exe utility provided as a component of Continuity Engine through the Engine Management Client User Interface tasks or Neverfail Engine Manager **Applications > Tasks > User Accounts**.
- Neverfail recommends integrating Microsoft DNS into AD so that DNSUpdate.exe can identify all DNS Servers that require updating.

Firewalls

1. 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:

- From VMware vCenter Server -> Engine Management Service
 - TCP 443 / 9727 / 9728 / Ephemeral port range
- From VMware vCenter Server -> Engine Server node
 - TCP 443 / Ephemeral port range
- From Engine Management Service -> VMware vCenter Server
 - TCP 443 / 9727 / 9728 / Ephemeral port range
- From Engine Management Client -> Engine Server node
 - TCP 7 / 445 / 135-139 / 9727 / 9728 / Ephemeral Port Range
- From Engine Server node -> Engine Management Service
 - TCP 7 / 445 / 135-139 / 9727 / 9728 / Ephemeral Port Range
- From Engine Server node -> VMware vCenter Server
 - TCP 443 / Ephemeral port range
- From Engine Server node -> Engine Server node in Duo/Trio and back
 - TCP 7 / 52267 / 57348 / Ephemeral port range or NFServerR2.exe (recommended)
- From Advanced Management Client -> Engine Server node in Duo/Trio and back
 - TCP 52267 / 57348 / Ephemeral port range or NFServerR2.exe (recommended)

For more detailed information, see KB-2907 Firewall Configuration Requirements for Neverfail Continuity Engine.

Note: The default dynamic ephemeral port range starting from Windows 2008 is from 49152 to 65535.

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 Continuity Engine virtual machine will be created from the Primary server.

- The specifications of the Secondary Continuity Engine virtual machine must match the Primary physical server as follows:
 - Similar CPU
 - identical Memory
 - Same type/mode for the Boot Firmware (UEFI, BIOS)
- The Secondary Continuity 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.

The Primary server must meet the hardware and software requirements as specified in the **Pre-Install Requirements**.

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
- 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.

- Same type/mode for the Boot Firmware (UEFI, BIOS)
- OS version and Service Pack version must match the Primary server
- OS must be installed to the same driver letter and directory as on the Primary server
- 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 Continuity 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 Continuity 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

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

- Neverfail for Exchange
 - ForeFront
 - Symantec Mail Security
- Neverfail for File Server
- Neverfail for IIS
- Neverfail for SharePoint Server
 - Office Online Server
- Neverfail for SQL Server
- Neverfail for MySQL Server
- Neverfail for VMware vCenter Server
- Neverfail for VMware vSphere 6.x Plug-in Suite
 - VMware vCenter Server
 - VMware Platform Services Controller
 - VMware Authentication Proxy 6.0
 - VMware Update Manager
 - VMware vPostgres
 - VMware Composer 5.x, 6.x and 7.x
- Neverfail System Plug-in
- Neverfail for Mitel MiContact Center
- Neverfail for Progress MOVEit Plugin Suite
 - MOVEit Automation Server
 - MOVEit Transfer Server
 - MOVEit Mobile Server
 - MOVEit Analytics Agent
- Neverfail for SolarWinds Orion Network Management
- Neverfail for Mitel MiContact Center
- Neverfail for Atlassian Confluence

- Neverfail for Atlassian Jira
- Neverfail for Oracle Database
 - Oracle Management Agent
- Neverfail for PostgreSQL
- Neverfail for Veeam Backup and Replication 9.5 Plug-ins Suite
 - Veeam Backup and Replication
 - Veeam Backup Enterprise Manager
 - Veeam Availability Console
 - Veeam Backup Proxy
 - Veeam WAN Accelerator
 - Veeam Cloud Gateway
- Neverfail for Apache Tomcat

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

Networking Configuration

Networking requirements are contingent upon how Continuity Engine is to be deployed. To deploy as a High Availability (HA) solution, a LAN configuration is required. To deploy Continuity 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 Channel can 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 Continuity 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 Continuity Engine, see **Adding an Additional Network Interface Card** in this guide.

Local Area Network (LAN)

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

Wide Area Network (WAN)

Continuity Engine supports sites with different subnets. In this scenario, the Primary and Secondary servers in the Continuity 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 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)
- 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
 - Primary site nodes and DR node can use either the same or different Public IP addresses
 - Provide a user account with rights to update DNS using the DNSUpdate.exe utility provided as a component of Continuity Engine through the Engine Management Service User Interface tasks or Neverfail Advanced Management Client **Applications > Tasks > User Accounts**. The necessary permissions required for DNSUpdate utility to operate correctly should be at least those indicated in the procedure Granting User the Rights to Run the DNSUpdate Tasks, from below.
 - Neverfail recommends integrating Microsoft DNS into AD so that DNSUpdate.exe can identify all DNS Servers that require updating
 - 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

Follow the procedure below to grant specific permissions to run DNSUpdate:

1. Create a dedicated domain account that will be used only for the DNSUpdate process.
2. Add the following necessary permissions:

Note: These steps should be performed on all the Microsoft DNS servers that will need to have records updated / zone refreshed during a switchover or a failover.

1. Membership in the *BUILTIN\Distributed COM Users* group.
2. Membership in the *DNSAdmins* group (domain wide) OR equivalent via ACLs on the DNS server / zones.
3. Remote Enable permissions for the *ROOT\MicrosoftDNS WMI* namespace. Follow the steps below to do this:
 1. Go to **Start > Run** and type *wmimgmt.msc*, then click **OK**.
 2. Right-click on *WMI Control (Local)* and select *Properties*.
 3. Select the *Security* tab.
 4. Expand *ROOT*, navigate to *MicrosoftDNS* and select the namespace.
 5. Click on the **Security** button at the bottom right of the window. This action edits the security settings for the *RootMicrosoftDNS WMI* namespace.
 6. Click **Advanced**.
 7. Add the designated *DNSUpdate* user to the list, and select *Allow* for at least the *Remote Enable* permission.
 8. Click **OK** (on all windows opened previously) to save the new permissions.
4. **Only for DNS Serves running on Windows 2003:**
 1. From **Start > All Programs > Administrative Tools**, open **DNS**.
 2. Right click the name of the DNS server and select *Properties*.
 3. Select the *Security* tab.
 4. Add the *DNSAdmins* group to the list and give it *Full Control*.
 5. Click **OK** on all windows open previously to save the new security settings.
5. Test the *DNSUpdate* task, while being run under the new user, by performing a switchover / switchback.

Continuity 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 Continuity Engine.

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

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

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

- **Single NIC Installation:** Continuity 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:** Continuity Engine can be installed on a server with multiple NICs. You can choose which NIC will be used for the Neverfail Channel connection.

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 a 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.

- If *Register this connection's addresses in DNS* is found enabled on any of the used NICs during Continuity Engine installation a warning will be displayed and then the option will be automatically disabled. It is recommended to verify and configure (if needed) the DNS records on your DNS server.

The Secondary/Tertiary server NICs configuration:

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

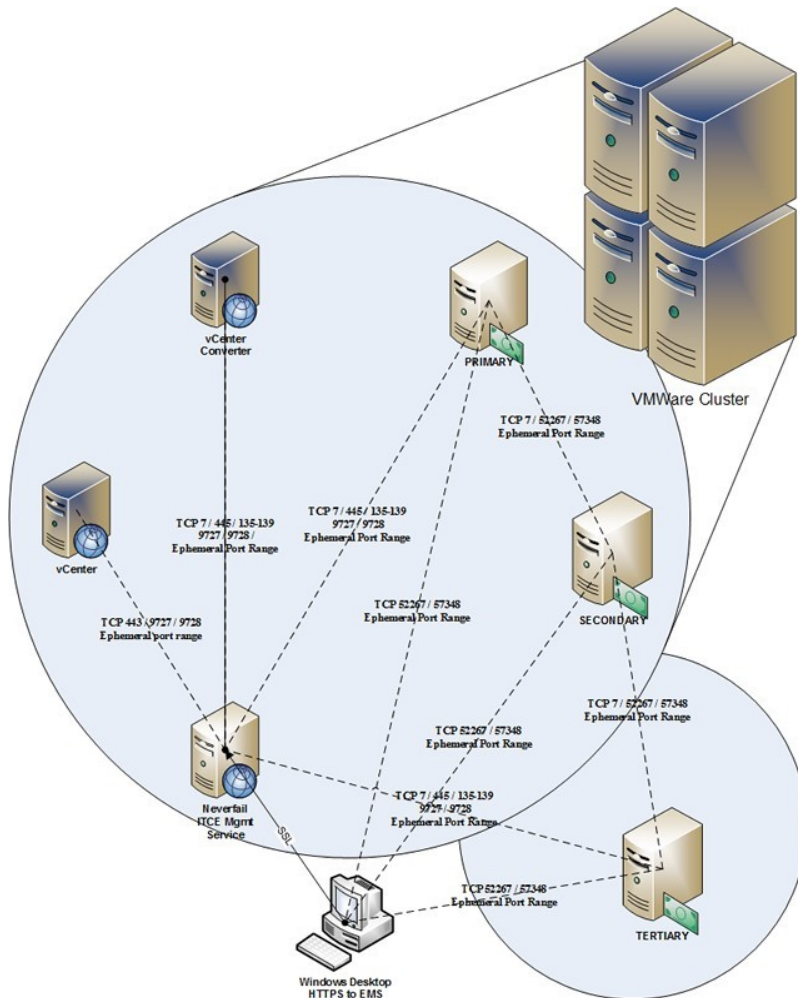
Note: If deploying as a pair in a WAN, the Public IP address of the Secondary server may be in a different subnet than the Primary server. If configured in a trio, the Primary and Secondary servers are configured for LAN deployment and the Tertiary server is configured for a WAN deployment.

- 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.
- If *Register this connection's addresses in DNS* is found enabled on any of the used NICs during Continuity Engine installation a warning will be displayed and then the option will be automatically disabled. It is recommended to verify and configure (if needed) the DNS records on your DNS server.

Firewall Configuration

When firewalls are used to protect networks, you must configure them to allow traffic to pass through specific ports for Continuity 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 Continuity Engine from the Engine Management Service
- Port 52267 for the Client Connection port
- Port 57348 for the Default Channel port



Note: 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 Continuity 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. Continuity Engine does not replicate this information.

Installing Neverfail Continuity Engine

This chapter discusses the installation process used to implement Neverfail Continuity Engine 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.

- **Installing Continuity Engine Management Service**
- **Deploying Continuity Engine on the Primary Server**
- **Automated Deployment of Stand-by Servers with Automatic Cloning**
- **Semi-Automatic Deployment of Stand-by Servers Leveraging Assisted Cloning**
- **Manually installing Continuity Engine, without using Continuity Engine Management Service**
- **Post Installation Configuration**

Installing Continuity Engine Management Service

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](#).

If installing Neverfail Continuity Engine Management Service on a French Localized Windows Server then the **Administrators** group must be created prior of starting the installation. This group's presence is required for a successful creation and configuration of the **NeverfaiEngine** user log on account.

Installing or upgrading the Neverfail Continuity Engine Management Service.

Note: Neverfail Continuity Engine Management Service upgrades in-place: you just need to install the .msi kit as indicated below, without uninstalling the previous version.

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

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 older version, you may be asked if a service should be stopped. Select **Close the applications and attempt to restart them.**

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. [Additional]: If installing on a **Domain Controller**: Add *NeverfailEngine* user to *Domain Admins* group

Before starting to use the Neverfail Continuity Engine Management Service, the *NeverfailEngine* user configured as Neverfail Engine Management Web Services log on account must be added to the *Domain Admins* group. There is no need to restart the Neverfail Engine Management Web Services.

9. [Additional]: If installing on a **French Localized Windows Server** following additional actions are required:

Before starting to use the Neverfail Continuity Engine Management Service, move the *NeverfailEngine* user configured as Neverfail Engine Management Web Services log on account from the *Administrators* group to the *Administrateurs* group. Then restart the Neverfail Engine Management Web Services. The *Administrators* group may be deleted if not needed anymore.

10. 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 Continuity Engine on the Primary Server

Prior to deploying Continuity 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 Continuity Engine on the target servers identified in the cluster and validate that the servers meet the minimum requirements for a successful installation.

Installing Continuity Engine on the Primary server: server protection flow

Once the server protection is initiated by clicking the **Protect server** or **Protect** buttons, the **Protect Server (Install Engine)** dialog is displayed.

Protect server (install Engine) [X]

Select a target server | Validating install | Select public IP address | Ready to complete

Installing Engine [i] [v]

Enter DNS name or IP address | Select from inventory

Target server

dns.name / 192.168.1.1

Username (with full administrator permissions) | Password

administrator | *****

Next

The next steps will guide you through the server protection flow:

1. Start with selecting a **target server** on which the Continuity Engine will be installed. This server will become the Primary instance of your protected cluster. There are two methods of selecting the target server; you can choose either one by clicking the corresponding tab:
 - **Enter DNS name or IP address** of the server you want to protect; when you already know the DNS name or IP address of your server, manually enter it in the **Target server** field.

- **Select from inventory** a server to protect. When a vCenter Server connection is configured, this will show you all the servers available in you vCenter inventory and allow you to search and select the one you want to protect. The **Target server** field is automatically populated with the selected server's DNS name.
2. Provide the administrator user name and password for the server you've selected for protection. Full administrative permissions are required for this account in order to perform the server protection underlying operations.

If User Access Control (UAC) is enabled on the target server, either use the target's built-in local Administrator account or, for Domain member target servers, the a domain user account member in the target's local Administrators group.

If UAC is not enabled, any user account that is member of the target's local Administrators group will work.

3. Click Next to proceed to the **Validating install** step. Here, EMS assesses the target server and validates the provided user account. Any issues with the target server or user account are shown upon the completion of the validation procedure.

The **Validating Result** shows errors, warnings or info about the targeted server. The output messages are displayed as a list sortable by message type. In case of errors, the Server Protection flow can not proceed until the errors are solved.

In case of warnings, the Server Protection flow can proceed only when the warning messages are acknowledged by the user.

The info messages can provide good insights on the targeted system status and may guide you towards optimizing your server.

4. Click Next to continue to the **Select public IP addresses** step. In this step you can select the IP addresses that will be used by clients to access your server. These IP addresses will be replicated on the passive high availability instance when added.

Make sure to unselect the IP addresses which are already assigned as channel addresses or dedicated management addresses.

5. Click Next to proceed to the **Ready to complete** step, in which the Protect server dialog will present summary of the settings previously done. Upon clicking the **Finish** button, Continuity Engine will be deployed on the targeted server and will automatically discover applications it can protect.

The applications found available for protection will have their services set to Manual start, in order to allow individual management.

6. EMS will show the Servers page where the Continuity Engine deployment triggered by the executed flow is listed in the **Operation in progress** section. Once the deployment is complete, your now protected server will be listed in the Protected Servers panel.

Automated Deployment of Stand-by Servers with Automatic Cloning

1. You have the following options:

- If the Primary server is a Managed* virtual machine, go to **Step 4**.
- Otherwise go to **Step 2**.

*Managed virtual machine = a VMware virtual machine managed by the vCenter Server configured in Engine Management Service.

2. In the **Settings** page of the EMS, locate the **vCenter Converter connection** and click on the **Configure** button. The *Configure vCenter Converter connection* dialog 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**.

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.

4. Navigate to the **Servers** page and select the server for which you want to deploy stand-by instances.

The **Status** panel shows the protected server as a single Production instance. Clicking the **+ Add standby servers** button will open the **Add standby servers** dialog.

5. Select one of the following depending on the environment you intend to support:

- Add a stand-by server for high availability, continue with **Adding a High Availability instance**
- Add a stand-by server for disaster recovery, continue with **Adding a Disaster Recovery instance**
- Create Secondary and Tertiary stand-by VMs for HA and DR, continue with **Adding both HA and DR instances to a protected server**

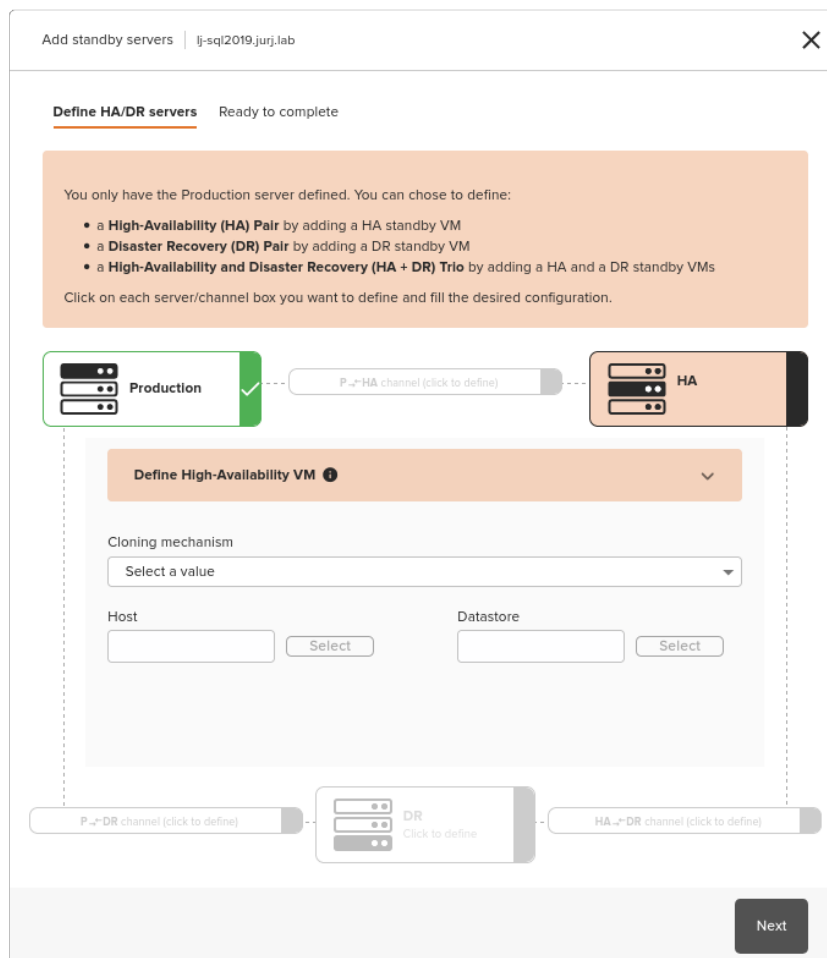
You can also create a stand-by VM for Disaster Recovery for an existing High Availability pair, and vice-versa.

Adding a High Availability instance

Adding a **High Availability (HA) instance** to your protected Production server enhances this protection. It allows Continuity Engine to provide a seamless transition to a perfectly synchronized standby server instance, in case your Production server is affected by any number of issues.

A HA instance is a perfect replica of your Production server, created by Continuity Engine as a virtual machine. The HA instance is added as a passive, secondary instance, which turns your production server in an active, primary instance of a protected cluster.

Clicking the **+ Add standby servers** button will open the **Add standby servers** dialog.



When adding a HA instance, you will need to configure the HA instance itself and the Production-HA Channel. The next steps will guide you through the process of adding a HA instance to your protected server:

1. The **Define HA/DR servers** step of the dialog presents the available options for your current configuration in the form of clickable boxes. Select the **HA box** to configure the High Availability instance for your server. The mid section of the dialog will adjust to provide settings for the selected option (HA in this scenario).
 - Select a value for **Cloning mechanism**. **Automated** cloning relies on configured vCenter Server and vCenter Converter connections in order to seamlessly create the standby HA instance VM.
 - Set the **Host** (the server hosting the clone) and the **Datastore** (clone storage) options. The Host and Datastore should be different than those of the production server, for better protection against server or storage failure.

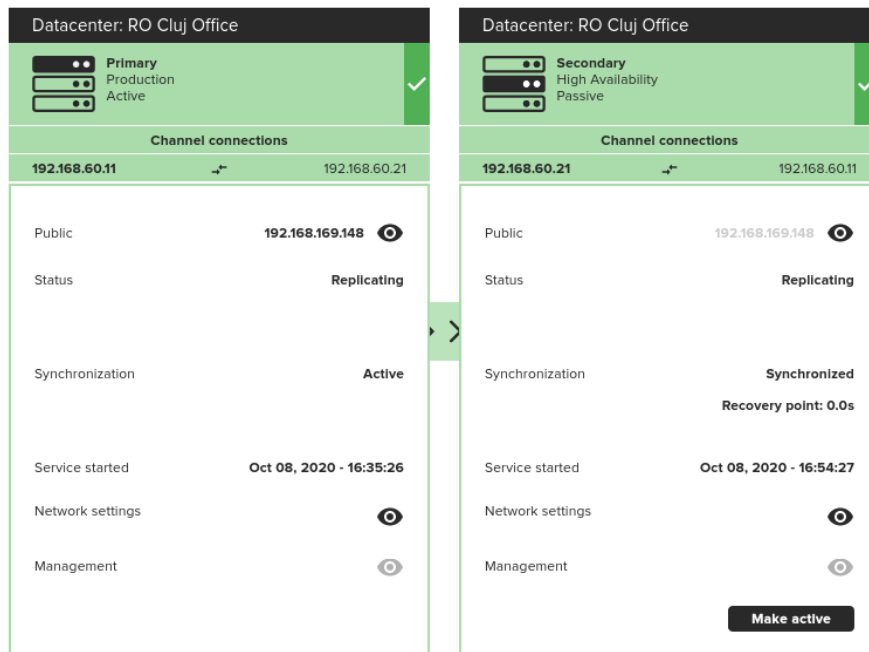
Once the HA instance has been defined, the **HA box** is marked as defined by a check mark on green background.

2. Select the **P - HA channel box** (between your production instance and the previously defined HA instance) to configure the Production-to-High-Availability Channel. The mid section of the dialog will present the channel options; proceed as follows:
 - Select the network adapter for the channel from the combo-box. The network adapters available on the Production server will be listed as options.
 - In the Production IPv4 address field, enter the IP address to be used on the Production server side of the P - HA channel connection.
 - In the Production subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.
 - In the HA IPv4 address field, enter the IP address to be used on the HA side of the P - HA channel connection.
 - In the HA subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.

Once the P - HA channel has been configured, the P - HA channel box is marked as defined by a check mark on green background.

3. Click Next to proceed to the **Ready to complete** step of the procedure. The summary of the configuration is displayed under the High-Availability VM configuration section. You can review your settings and go back to the previous step to edit the configuration if needed.
4. Click Finish to start the operation. The **Operation in progress** section will list the ongoing operation and its real time progress. Once the operation is complete, the HA instance will be listed in the Status view of the Server Details section. Replication between your production instance and the new HA instance will be started automatically.

Once replication is started, the two instances of your protected cluster will reflect the Replicating status. The Primary (Production) instance synchronization should be Active while the Secondary (HA) instance should be synchronized, with a 0 seconds recovery point.




Adding a Disaster Recovery instance

Unlike a High Availability instance, a DR instance usually does not reside in the same physical location as the Production or HA instances. It can be added as a single passive instance to a Production active instance (this current article) or as a third passive instance in a protected trio cluster, along with a HA passive instance (see the **Add both HA and DR instances** article).

Like a HA instance, the DR instance is a perfect replica of your Production server, created by Continuity Engine as a virtual machine. The DR instance is added as a passive, secondary instance, which turns your production server in an active, primary instance of a protected cluster.

Clicking the **+ Add standby servers** button will open the **Add standby servers** dialog.



When adding a DR instance, you will need to configure the DR instance itself and the Production-DR Channel. The next steps will guide you through the process of adding a DR instance to your protected server:

1. The **Define HA/DR servers** step of the dialog presents the available options for your current configuration in the form of clickable boxes. Select the **DR box** to configure the Disaster Recovery instance for your server. The mid section of the dialog will adjust to provide settings for the selected option (DR in this scenario).
 - Select the **Public IP Address** for the DR instance.
 - **Identical to Production:** If your DR instance will be hosted at the same site as your production instance (not recommended), using the same subnet, the same public IP address can be used.
 - **Different than Production:** If your DR instance will run on a different site (recommended), using a different subnet than your production site, the DR instance will require 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.

- Select the desired **Cloning mechanism**. For Automated cloning, you can choose between **Automated Powered-On** and **Automated Powered-Off**:
 - **Automated Powered-On** cloning relies on configured vCenter Server and vCenter Converter connections but also on a high-bandwidth connection to the remote site where your DR instance will be hosted. The new VMware VM will be created directly on the remote host and started automatically.
 - **Automated Powered-Off** cloning uses the same vCenter Server and Converter connections to create a VMware VM on a temporary host (the same host as the production instance, for example). The resulting DR instance will be powered off after creation and ready to be transferred at its final hosting site using FTP or removable media.
- Set the **Host** (the server hosting the clone) and the **Datastore** (clone storage) options. The Host and Datastore should be different than those of the production server, for better protection against server or storage failure.

Once the DR instance has been defined, the **DR box** is marked as defined by a check mark on green background.

2. Select the **P - DR channel box** (between your production instance and the previously defined DR instance) to configure the Production-to-Disaster-Recovery Channel. The mid section of the dialog will present the channel options; proceed as follows:
 - Select the network adapter for the channel from the combo-box. The network adapters available on the Production server will be listed as options.
 - In the Production IPv4 address field, enter the IP address to be used on the Production server side of the P - DR channel connection.
 - In the Production subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.
 - In the DR IPv4 address field, enter the IP address to be used on the DR side of the P - DR channel connection.
 - In the DR subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.

Once the P - DR channel has been configured, the P - DR channel box is marked as defined by a check mark on green background.

3. Click Next to proceed to the **Ready to complete** step of the procedure. The summary of the configuration is displayed under the Disaster Recovery VM configuration section. You can review your settings and go back to the previous step to edit the configuration if needed.
4. Click Finish to start the operation. The **Operation in progress** section will list the ongoing operation and its real time progress. Once the operation is complete, the DR instance will be listed in the Status view of the Server Details section. Replication between your production instance and the new HA instance will be started automatically.

Once replication is started, the two instances of your protected cluster will reflect the Replicating status. The Primary (Production) instance synchronization should be Active while the Secondary (DR) instance should be synchronized, with a 0 seconds recovery point.

The image shows two side-by-side screenshots of the Neverfail Continuity Engine interface, both titled "Datacenter: RO Cluj Office".

The left screenshot shows the Primary instance (Production) with a status of "Active". Below the instance details, there is a "Channel connections" section showing a connection between IP addresses 192.168.60.12 and 192.168.60.32. The instance details include:

- Public IP: 192.168.169.148
- Status: Replicating
- Synchronization: Active
- Service started: Oct 08, 2020 - 17:05:46
- Network settings and Management icons are visible.

The right screenshot shows the Secondary instance (Disaster Recovery) with a status of "Passive". Below the instance details, there is a "Channel connections" section showing a connection between IP addresses 192.168.60.32 and 192.168.60.12. The instance details include:

- Public IP: 192.168.169.148
- Status: Replicating
- Synchronization: Synchronized
- Recovery point: 0.0s
- Service started: Oct 08, 2020 - 17:15:32
- Network settings and Management icons are visible.
- A "Make active" button is present at the bottom right.

Adding both HA and DR instances to a protected server

Clicking the **+ Add standby servers** button will open the **Add standby servers** dialog.

When adding both HA and DR instances for a trio configuration, you will need to configure both HA and DR instances individually, as well as the channel connections between all three instances.

1. The **Define HA/DR servers** step of the dialog presents the available options for your current configuration in the form of clickable boxes. Select the **HA box** to configure the High Availability instance for your server. The mid section of the dialog will adjust to provide settings for the selected option (HA in this scenario).
 - Set the **Cloning mechanism** to **Automated**. Automated cloning relies on configured vCenter Server and vCenter Converter connections in order to seamlessly create the standby HA instance VM. This particular cloning mechanism is required in order to deploy both HA and DR instances in a single step.
 - Configure the **Host** (the server hosting the clone) and the **Datastore** (clone storage) options. The Host and Datastore should be different than those of the production server, for better protection against server or storage failure.

Once the HA instance has been defined, the **HA box** is marked as defined by a check mark on green background.

2. Select the **P - HA channel box** (between your production instance and the previously defined HA instance) to configure the Production-to-High-Availability Channel. The mid section of the dialog will present the channel options; proceed as follows:
 - Select the network adapter for the channel from the combo-box. The network adapters available on the Production server will be listed as options.
 - In the Production IPv4 address field, enter the IP address to be used on the Production server side of the P - HA channel connection.
 - In the Production subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.
 - In the HA IPv4 address field, enter the IP address to be used on the HA side of the P - HA channel connection.
 - In the HA subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.

Once the P - HA channel has been configured, the P - HA channel box is marked as defined by a check mark on green background.

3. Select the **DR box** to configure the Disaster Recovery instance for your server. The mid section of the dialog will adjust to provide settings for the selected option (DR in this scenario).
 - Select the **Public IP Address** for the DR instance.
 - **Identical to Production:** If your DR instance will be hosted at the same site as your production instance (not recommended), using the same subnet, the same public IP address can be used.
 - **Different than Production:** If your DR instance will run on a different site (recommended), using a different subnet than your production site, the DR instance will require 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.
 - Select the desired **Cloning mechanism**. You can choose between Automated Powered-On and Automated Powered-Off.
 - **Automated Powered-On** cloning relies on configured vCenter Server and vCenter Converter connections but also on a high-bandwidth connection to the

remote site where your DR instance will be hosted. The new VMware VM will be created directly on the remote host and started automatically.

- **Automated Powered-Off** cloning uses the same vCenter Server and Converter connections to create a VMware VM on a temporary host (the same host as the production instance, for example). The resulting DR instance will be powered off after creation and ready to be transferred at its final hosting site using FTP or removable media.
- Configure the **Host** (the server hosting the clone) and the **Datastore** (clone storage) options. The Host and Datastore should be different than those of the production or HA instances, for better protection against server or storage failure.

Once the DR instance has been defined, the **DR box** is marked as defined by a check mark on green background.

4. Select the **P - DR channel box** (between your production instance and the previously defined DR instance) to configure the Production-to-Disaster-Recovery Channel. The mid section of the dialog will present the channel options; proceed as follows:

- Select the network adapter for the channel from the combo-box. The network adapters available on the Production server will be listed as options.
- In the Production IPv4 address field, enter the IP address to be used on the Production server side of the P - DR channel connection.
- In the Production subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.
- In the DR IPv4 address field, enter the IP address to be used on the DR side of the P - DR channel connection.
- In the DR subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.

Once the P - DR channel has been configured, the P - DR channel box is marked as defined by a check mark on green background.

5. Select the **HA - DR channel box** (between your HA instance and the DR instance) to configure the High-Availability-to-Disaster-Recovery Channel. The mid section of the dialog will present the channel options; proceed as follows:

- Select the network adapter for the channel from the combo-box. The network adapters available on the Production server will be listed as options.
- In the HA IPv4 address field, enter the IP address to be used on the HA server side of the HA - DR channel connection.

- In the HA subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.
- In the DR IPv4 address field, enter the IP address to be used on the DR side of the HA - DR channel connection.
- In the DR subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.

Once the HA - DR channel has been configured, the HA - DR channel box is marked as defined by a check mark on green background.

- Click Next to proceed to the **Ready to complete** step of the procedure. The summary of the configuration is displayed under the High-Availability VM and Disaster Recovery VM configuration sections. You can review your settings and go back to the previous step to edit the configuration if needed.
- Click Finish to start the operation. The **Operation in progress** section will list the ongoing operation and its real time progress. Once the operation is complete, the HA and DR instances will be listed in the Status view of the Server Details section. Replication between your production instance and the new standby instances will be started automatically.

Once replication is started, the three instances of your protected cluster will reflect the Replicating status. The Primary (Production) instance synchronization should be Active while the Secondary (HA) and Tertiary (DR) instances should be synchronized, with a 0 seconds recovery point.

The image displays three screenshots of the Neverfail Continuity Engine interface, each showing the configuration and status of a different instance in a 'Datacenter: RO Cluj Office'.

- Primary Instance (Production):** Status is 'Active'. Channel connections are listed with IP addresses 172.16.155.21 and 172.16.155.26. Synchronization is 'Active'. Service started on Aug 21, 2020 - 15:47:23.
- Secondary Instance (High Availability):** Status is 'Replicating'. Channel connections are listed with IP addresses 172.16.155.22 and 172.16.155.23. Synchronization is 'Synchronized' with a 'Recovery point: 0.0s'. Service started on Aug 21, 2020 - 15:57:55. A 'Make active' button is visible at the bottom.
- Tertiary Instance (Disaster Recovery):** Status is 'Replicating'. Channel connections are listed with IP addresses 172.16.155.25 and 172.16.155.24. Synchronization is 'Synchronized' with a 'Recovery point: 0.0s'. Service started on Aug 11, 2020 - 16:46:40. A 'Make active' button is visible at the bottom.

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

1. Navigate to the **Servers** page and select the server for which you want to deploy stand-by instances. The **Server Details** page will open.
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 9**

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

3. On the EMS **Server Details** page, clicking the **+ Add standby servers** button will open the **Add standby servers** dialog.
4. Select the **HA box** to configure the High Availability instance for your server.
5. Set the **Cloning mechanism** to use **assisted** cloning.
6. 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.

7. Review the information on the *Ready to Complete* step and if accurate, click **Finish** to prepare the Secondary server for assisted cloning using a third-party tool.

During the pre-condition check, the following status messages will be displayed:

- 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.

8. Once complete, perform *Post Installation Configuration* tasks as listed in the *Neverfail Continuity Engine Installation Guide*.
9. On the EMS **Server Details** page, click the **+ Add standby servers** button.
The **Add standby servers** dialog is displayed.
10. Select the **DR box** to configure the Disaster Recovery instance for your server.
11. 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.
12. 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**.
13. Set the **Cloning mechanism** to use **assisted** cloning. Click **Next**.
The *Ready to Complete* step is displayed.
14. Review the information on the *Ready to Complete* step and if accurate, click **Finish** to prepare the Secondary server for assisted 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

Note: If adding a stand-by server for disaster recovery to an existing Continuity Engine HA Pair, then the SECONDARY server is the source server that should be cloned.

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.

Manually installing Continuity Engine without using Continuity Engine Management Service

This procedure should be used to install Continuity Engine on both Primary and Secondary nodes only when an Continuity Engine Management Service (EMS) based installation is not possible.

Carefully follow the steps listed below to manually install Continuity Engine.

1. On the management machine, install the Neverfail Continuity Engine Management Service. To do this, follow the [Installing Continuity Engine Management Service](#) steps.
2. Once the EMS is installed on the management machine, navigate to the following path: C:\ProgramData\Neverfail\VAD\catalog. Copy the following to the production server to be set as Primary:
 - NF-Engine-9-xxxxx.msi
 - ScopeCollector.msi
 - SystemMonitorNF plugin
 - VMAdaptor plugin
 - The plugins required for the applications that need to be protected
3. On the server to become Primary, install NF-Engine-9-xxxxx.msi.
4. Once the installation is complete, start Continuity Engine service (it will stop immediately as it is not completely configured).
5. Open the Registry Editor and locate the following registry entry: HKEY_LOCAL_MACHINE\Software\JavaSoft\Prefs\neverfail\current\Controller. Add the following:
 - */Group/Type* string value and set it /B/I/N/A/R/Y for Pair deployments
 - */Group/Type* string value and set it /T/E/R/N/A/R/Y for Trio deployments
6. Start the *Server Configuration Wizard* and make the required settings:
 - Public IP address
 - Channel(s) IP address routing, Default Channel Port = 57348, Channel IP SkipAs-Source Policy = Skip when Active (default)
 - License key (can be added at a later time using EMS)

-
- Management IP address - if used
7. Start the *Engine Service* on the Primary Server and restart the *Continuity Engine Webservices*.
 8. On the EMS machine, add the Primary server as protected server (either through **Discover Server** or **Add a protected server**).
 - If the Passive server(s) can be set up using the EMS's Automated Cloning, follow the indications in the dedicated deployment section and then return to the **Step 14** of this list.
 - If the Passive server(s) must be manually configured, continue with **Step 9**.
 9. On the Primary server, set the *Continuity Engine service* to manual.
 10. Clone the first passive (Secondary) server using Primary server as source. Choose not to start the clone once cloning is finished.
 11. Once the clone is ready, edit its VM setting and disconnect the NIC(s).
 12. Start the Secondary server, go to *Server Configuration Wizard* and make the required settings (set the Secondary as passive).
 13. Start the *Engine Service* and reconnect the NIC(s).
 14. For Trio configuration: repeat the **steps 10-13** for cloning the second passive (Tertiary) server.
 15. Once the servers are in sync, install *Neverfail Scope* on each server in the cluster by running the *ScopeCollector.msi*.
 16. Next, set Continuity Engine service to Automatic startup on all servers in the cluster.
 17. On the Primary server, login to Advanced Management Client and install the plugins:
 - SystemMonitorNF plugin
 - VMAdaptorFileServerNF plugin
 - Plugins for the applications that need to be protected
 18. License the installed Engine cluster.
-

Post Installation Configuration

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

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

Configure the VmAdapter Plug-in

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

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

The *Edit Plug-in* dialog is displayed.

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

This option requires vSphere HA Application monitoring for the cluster and VM.
7. Click **Save**.

Adding an Additional Network Interface Card

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

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

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

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

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

Split-brain Avoidance

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

Best practice

Neverfail recommends to always configure the split-brain avoidance by configuring the pinging over the "public" network using a management/auxiliary IP address on each server as the "from address". One can use in fact a management/auxiliary or any IP address not being used by Engine. If possible (to avoid the single-point-of-failure), this IP address should sit either in the Public network or in any other dedicated network, different from the channel network - This will assure an extra layer of protection against potential false failovers and split-brain syndrome. This configuration may be applied to both single and multi-NIC topologies.

Procedure

Split-brain avoidance configuration can be performed from EMS **Server Details > Monitoring > Server monitoring**. More details can be found in the **Neverfail Continuity Engine Administrators Guide**

Installation Verification Testing

Installation Verification testing is a procedure performed to validate the configuration of the server pair and its performance after installation.

- **Testing a Continuity Engine Pair**
 - **Exercise 1 Auto-switchover**
 - **Exercise 2 Data Verification**
 - **Exercise 3 Switchover**
- **Testing a Continuity Engine Trio**
 - **Exercise 1 Auto-switchover**
 - **Exercise 2 Managed Switchover**
 - **Exercise 3 Data Verification**

Testing a Continuity Engine Pair

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.

In this document, the term **Pair** refers to a Continuity Engine pair.

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, Continuity 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 pre-configured 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.

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	

Machine ID	Activity	Results
	In the <i>Server Details</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 that 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:
 1. Open the *Configure Server* wizard.
 2. Select the *Machine* tab.
 3. Select the *Primary* server as active.
 4. 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 (**SI-**).
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

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>Servers</i> page of the Engine Management Service, select the server pair.	The Server Details 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.
View the <i>Status</i> pane and wait for synchronization to be <i>Active</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 & Registry status</i> 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 Continuity Engine. Perform this exercise only after successfully completing the Auto-switchover and Data Verification Exercises.

Steps to Perform

Activity	Results
Using the Engine Management Service, review the <i>Summary Status</i> pane to verify that both the <i>File System</i> and <i>Registry status</i> are <i>Synchronized</i> .	
Navigate to the Actions drop-down and click on Make Primary Server Active .	The Engine Management Service Summary Status 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 Summary Status pane shows the applications starting on the newly active server. Both the File System and Registry status are Synchronized.
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 Continuity Engine Trio

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.

In this document, the term "Trio" refers to a Continuity Engine trio. Refer to the **Glossary** for more information about Continuity 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, Continuity 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 pre-configured 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, Continuity 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 Continuity 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.

Machine ID	Activity	Results
Primary	Open a command prompt.	
	Change directory to <code>C:\Program Files\Neverfail\R2\Bin</code>	
	Execute <code>nfavt.exe</code> When prompted, "Are you sure you wish to continue", click Continue .	Service is switched to the Secondary server and Continuity Engine shuts down on the Primary.
Secondary	Login to the Engine Management Service.	
	In the Servers pane of the Engine Management Service, select the server Cluster.	The <i>Server Details</i> page indicates that the Secondary server is active.

Machine ID	Activity	Results
	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 Continuity Engine and protected services on all servers.
2. Complete the following on all three servers:
 1. Open the *Configure Server* wizard.
 2. Select the *Machine* tab.
 3. Select the *Primary server* as active.
 4. Click **Finish**.
3. On the Secondary and Tertiary servers, right-click the taskbar icon and select **Start Continuity 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 Continuity 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

Continuity Engine provides manual control over switching the active server role to another server in the Cluster. On command, Continuity 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

Continuity Engine is running on the Secondary active server (**S/A**) and Tertiary server (**T/-**). Continuity 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.

Machine ID	Activity	Results
Secondary	Login to the Engine Management Service.	
	Click Snapshots .	The <i>Snapshots</i> screen is displayed.
	Under <i>Snapshots</i> , click Create . In the <i>Create Shadow</i> dialog, select <i>Secondary</i> , and then click OK .	A snapshot / rollback point is created prior to testing Secondary to Tertiary.

Machine ID	Activity	Results
	In the <i>Servers</i> pane of the Engine Management Service, select the <i>server Cluster</i> .	The <i>System Overview</i> screen is displayed.
	In the <i>System Overview</i> page, select the Tertiary server and then click Make Active .	Continuity 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 Continuity Engine and protected applications on the Secondary and Tertiary servers.
2. Complete the following on the Tertiary server:
 1. Open the *Configure Server* wizard.

2. Select the *Machine* tab.
3. Select the Secondary server as active.
4. Click **Finish**.
5. Right-click the taskbar icon and select **Start Continuity Engine** .
6. Verify that the Tertiary server is passive (**TI-**) and then shut down Continuity Engine.
3. On the Secondary, right-click the taskbar icon and select **Start Continuity Engine** .
4. After Continuity Engine starts on the Secondary server, login to the Engine Management Service.
5. Click **Snapshots**.
6. Under Snapshots, select the previously created snapshot (shadow copy) 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 Continuity 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

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

Machine ID	Activity	Results
Primary	Right-click the taskbar icon and select Start Continuity Engine	Continuity Engine successfully starts.
	Login to Engine Management Service.	
	In the Servers 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 server</i> 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 Snapshots .	The <i>Snapshots</i> screen is displayed.
	Under <i>Snapshots</i> , click Create . In the <i>Create Shadow</i> dialog, select <i>Tertiary</i> , and then click OK .	A snapshot / rollback point is created prior to testing Tertiary to Primary switchover.
Primary	In the <i>System Overview</i> screen, select the <i>Primary server</i> and click Make Active .	Continuity 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 Continuity Engine and protected applications on all servers.
2. Complete the following on the Primary and Secondary servers:
 1. Open the *Configure Server* wizard.
 2. Select the *Machine* tab
 3. Select the Tertiary server as active.
 4. Click **Finish**.
 5. Right-click the taskbar icon and select **Start Continuity Engine** .
 6. Verify that the Primary and Secondary servers are passive (**PI-** and **SI-**).

Glossary

Active

The functional state or role of a server when it is visible to clients through the network, running protected applications, and servicing client requests.

Alert

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

Active Directory (AD)

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

Active-Passive

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

Advanced Configuration and Power Interface (ACPI)

A specification that dictates how the operating system can interact with the hardware especially where power saving schemes are used. The Primary, Secondary, and Tertiary servers must have identical ACPI compliance.

Asynchronous

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

Basic Input/Output System (BIOS)

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

Cached Credentials

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

Channel Drop

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

Channel NIC (Network Interface Card)

A dedicated NIC used by the Neverfail Channel.

Checked

The status reported for user account credential (username/password) validation.

Cloned Servers

Servers that have identical configuration settings, names, applications, Security Identifiers (SIDs) and IP addresses, following the installation of Neverfail Engine.

Cloning Process

The Neverfail Continuity Engine process whereby all installed programs, configuration settings, and the machine name, Security Identifier (SID), and IP addresses are copied to another server.

Cluster

A generic term for a Neverfail Engine Pair or Trio and the set of machines (physical or virtual) involved in supporting a single protected server. A Neverfail Engine Cluster can include the machines used in a VMware or Microsoft cluster.

Connection

Also referred to as Cluster Connection. Allows the Engine Management Service to communicate with a Neverfail Engine Cluster, either on the same machine or remotely.

Crossover Cable

A network cable that crosses the transmit and receive lines.

Data Replication

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

Data Rollback Module

A Neverfail Continuity Engine module that allows administrators to rollback the entire state of a protected application, including files and registry settings, to an earlier point-in-time. Typically used after some form of data loss or corruption.

Degraded

The status reported for an application or service that has experienced an issue that triggered a Rule.

Device Driver

A program that controls a hardware device and links it to the operating system.

Disaster Recovery (DR)

A term indicating how you maintain and recover data with Neverfail Engine in event of a disaster such as a hurricane or fire. DR protection can be achieved by placing the Secondary server at an off-site facility, and replicating the data through a WAN link.

DNS (Domain Name System) Server

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

Domain

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

Domain Controller (DC)

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

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 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.