WHITE PAPER

NEVERFAIL
CONTINUOUS AVAILABILITY

# CONTINUITY ENGINE: TECHNICAL WHITE PAPER

## Proactively Protect Your Business From Downtime

## Purpose

This technical white paper is designed to help IT Infrastructure and Operations (I&O) personnel as well as application administrators better understand how Neverfail Continuity Engine (CE) keeps critical applications and business services in the datacenter continuously running and operational 24/7, while protecting them against down-time from any type of threat or failure conditions.

This paper includes coverage of both key enabling technologies essential for keeping critical applications continuously available, highlights the limitations of other alternative backup and replication solutions, and provides technical details of CE product features and architecture.

## Protecting the Most Critical Applications

Modern business operations are underpinned by a high degree of IT automation, and with many of these systems exposed directly to customers via the web, the business increasingly demands 24/7 availability for the most critical systems. In a highly competitive landscape, consequences arising from system outages can be severe, including reputational damage, lost revenues and, in extreme cases, a threat to the organization's ability to survive. While IT managers in most businesses have some form of a business continuity or recovery plan, a large majority of such plans are not regularly exercised and there may be little confidence that deployed resiliency infrastructure will be up to the job of protecting their most critical IT assets.

As virtualization infrastructure has become commonplace within the datacenter, IT professionals have leveraged virtual clustering and disaster recovery (DR) services to extend protection to the majority of IT workloads for the first time with greatly reduced cost and complexity compared to older approaches. However, the performance of protection provided by native virtual infrastructure, while "good enough" for the majority of applications, will inevitably fall short of keeping the most critical applications continuously available. For example, fully restoring service for the more complex multi-tiered applications deployed across both virtual and physical platforms following an outage can still take hours. The same is equally true for those systems protected solely by commercial backup and replication solutions because these approaches focus primarily on system recovery rather detecting problems and fixing them before infrastructure or application failure causes downtime. Critical systems demand continuous availability because this is what is required to ensure 24/7 commercial operations. The key technologies to look for within any infrastructure designed to keep critical applications continuously available are as follows:

- **Application awareness.** The majority of backup and replication solutions cannot detect application service degradation or application failures. As a result, these solutions will always be reactive to outages that will interrupt normal business operations. A true continuous availability solution will, by contrast, proactively detect and remediate any impending failure rather than reactively reboot a server to restore application service. Application awareness is also critical to ensure the successful failover of complex multi-tiered application components seamlessly.

> Commercial backup, replication and recovery infrastructures will not protect critical applications

- **Support for Recovery Point Objective (RPO) targets measured in seconds.** Replicated or backup copies of a protected application or server image are only as good as the most recent data synchronization. Most backup and replication solutions cannot handle RPOs lesser than 15 minutes, which means that critical data will almost certainly be lost during an outage. By contrast, continuous availability demands minimal or zero data loss which often drives the deployment of expensive storage array replication and/or WAN optimization infrastructure. Dedicated continuous availability infrastructure, on the other hand, will deliver the required RPO targets at a fraction of the cost.
- **Support for Recovery Time Objectives (RTO) measured in minutes to seconds.** Recovering from a disaster, especially a site-wide outage, can often take hours to reconfigure the network and connect multiple components to each other then verify that everything is working. Most backup and replication solutions have to recover a physical or virtual machine image from the replicated copy before individual servers can be powered back up, which can take hours. Compound this with the need to recover groups of applications and servers in a designated priority order across critical business services, and almost all of these solutions fall flat in terms of timely recovery. Critical applications cannot afford this level of downtime, so look for orchestration and automation to address these key issues within your failover infrastructure.
- **Low Total Cost of Ownership.** The level of protection offered by advanced continuous availability infrastructure should be easy to justify given the need to keep your most critical applications always-on. At least in theory, until the total cost of ownership figures associated with storage-based backup and replication systems are presented to senior management. Many IT professionals assume that only those with the largest enterprise-class budgets can afford to deploy this level of protection. Ironically, this is not the case, with specialized vendors such as Neverfail offering the best performing, most capable protection infrastructure, at an SMB-friendly total cost of ownership.

In the next section, we'll explain how CE tackles these problems while delivering on its promise of "no application downtime," irrespective of any threats or failure across any type of infrastructure.

## Continuity Engine Overview

CE provides total protection for critical business applications, ensuring 24/7 availability regard-less of any threats to uptime. CE prevents application failure by proactively detecting application failure signatures and switching degrading applications to a hot standby server before application failure causes user downtime. With its built-in replication, WAN optimization, continuous availability, disaster recovery and data protection capabilities, CE provides the most comprehensive protection for your most critical business services.

The key features and associated benefits of CE include:

- **Unified Continuous Availability and Disaster Recovery.** CE provides complete protection for your critical business services against application, server, network, storage, or site failures.
- **Built-in Replication.** CE's built-in replication greatly lowers the threshold for data loss by delivering near-zero RPOs. By working across heterogeneous storage or server hardware, not requiring shared storage or depending on other storage replication software and being able to replicate application, registry and file system data, it significantly reduces overall costs.
- **Proactive Application Health Monitoring.** CE prevents application failures by proactively monitoring application health in real-time and detecting patterns of degradation before a failure can occur. If such patterns are detected, automated remediation mechanisms are triggered to maintain application continuity.
- **Multi-tier Application Groups Protection.** CE can coordinate expedited failover of any arbitrary collection of application components spread across multiple servers, especially in the context of site failovers, and recover them in designated priority order.
- **Built-in WAN Acceleration.** CE's built-in data compression and data de-duplication capabilities significantly reduce disaster recovery operational costs by reducing WAN replication traffic and associated network bandwidth requirements by up to 80%.
- **VMware vSphere (HA) and vMotion integration.** CE extends application intelligence and proactive application health monitoring to workloads running on VMware vSphere infra-structure by enabling administrators to configure automated remediation triggers associated with VMware vSphere HA, vMotion, Storage vMotion or enhanced vMotion actions, when application faults or degraded health conditions are detected.
- **VMware SRM integration.** CE adds disaster recovery capabilities for physical machines within VMware Site Recovery Manager (SRM) by enabling synchronized virtual machine replicas of protected physical machines to be recovered during execution of VMware SRM recovery plans.
- **Integrated Data Protection.** CE provides a Data Rollback Module (DRM) that integrates with Windows Volume Shadow Copy Service (VSS) to prevent data corruption and data loss by creating shadow copies of application data that can be leveraged to roll back the application state during recovery to a previous point in time.
- **Tertiary Node Support.** CE provides flexible topology options for extended redundancy combining local HA and remote DR failover, as well as for multi-site DR. CE can also transition from an existing HA or DR pair to tertiary server quickly and easily.

- **Flexible Network Configuration Options.** CE allows administrators to deploy multiple network configurations whether it's a single NIC configuration or dual NIC configurations. These can be on the same subnet or different subnets.
- **CE Management Service Web Client.** CE's management web client is designed to facilitate point-and-click deployment of CE clusters by providing centralized administration over all CE deployments. See Figure 1.

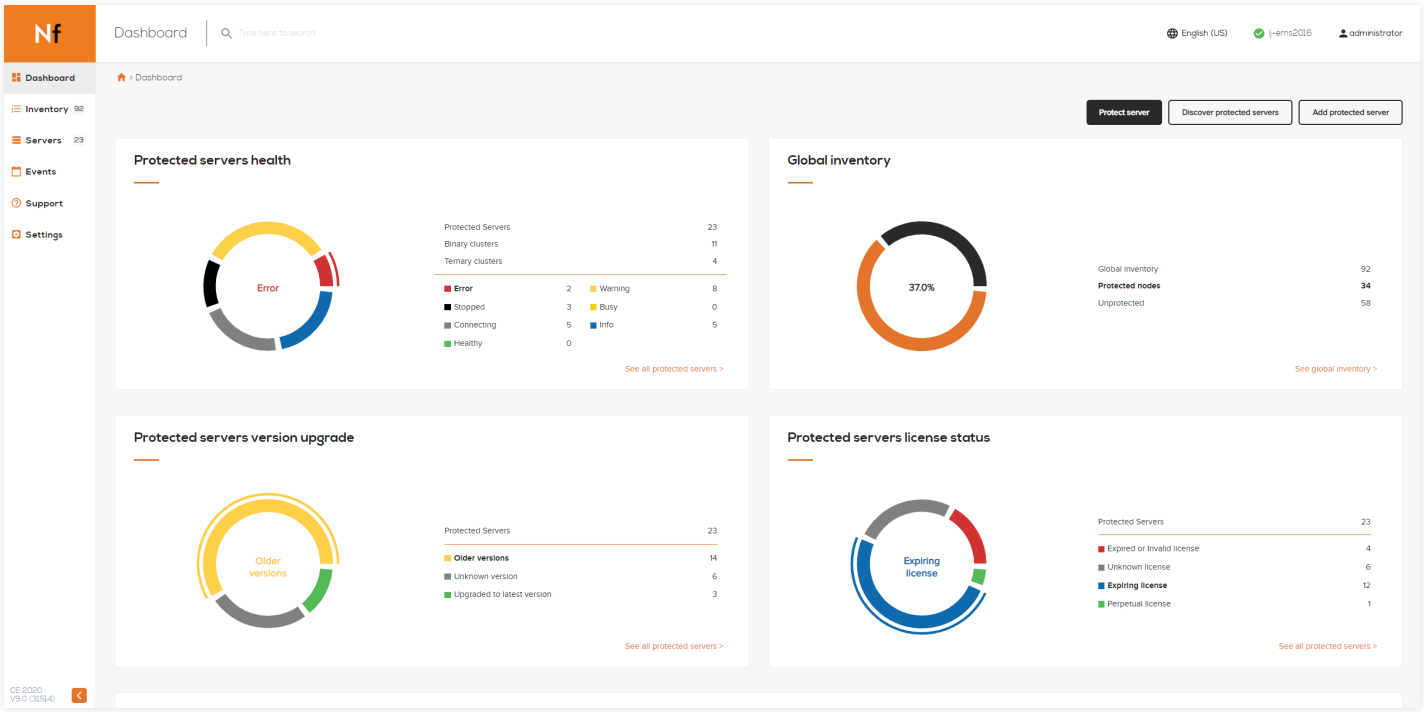


Figure 1: Engine Management Service Web Console

## Continuity Engine in Detail

The following sections focus on CE architecture as well as the core continuity operations available within CE.

### Solution Architecture

CE keeps critical applications continuously available by failing over to hot standby instances configured in either a protection pair or trio topology, determined by the number of server nodes participating in the CE "cluster." An CE protection pair refers to an application cluster across two servers, while an CE protection trio (or tertiary) refers to an application cluster spread across three servers. Deployment of CE onto servers associated with an CE cluster results in assignment of an Identity (Primary, Secondary or Tertiary) to each of the servers, resulting in either an CE Pair or an CE Tertiary configuration. See Figure 2 on the next page.

In addition to an Identity, each server within the CE cluster is also assigned a Role (Active or Passive). While the server Identity describes the membership order within the CE cluster, the server role describes the status of the protected application on that server node. Although the server Identity will typically not change between the cluster nodes, the server Role will change depending on

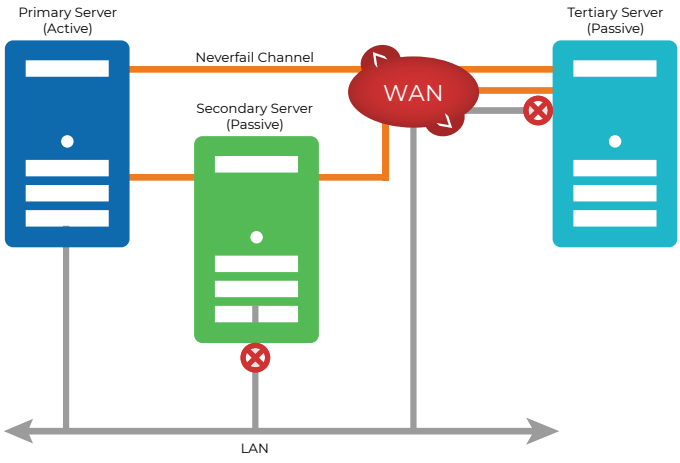


Figure 2: Engine Protection Pair

the status of the application instance on each cluster node as determined by CE. In its simplest form, CE operates as a Protection 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 end-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 as a pair, CE can be deployed for either HA using a high speed Local Area Network (LAN) connection, or for disaster recovery (DR) using a lower bandwidth Wide Area Network (WAN) connection where bandwidth optimization (using CE's built-in WAN Acceleration feature) may be required. When deployed in the Tertiary configuration, CE provides HA of protected applications over the LAN, and simultaneously provides DR with a third server located at a remote site, over a WAN connection. CE can also transition an existing HA or DR pair into tertiary quickly and easily. All that is required is an updated license key.

***Take Note:***

An Engine cluster can include physical and virtual machines as well as servers already participating in a VMware vSphere Cluster.

## Data Replication

CE keeps servers within its cluster synchronized by means of its built-in replication capabilities. See Figure 3 on the next page. Once configured, CE tracks disk I/O changes pertaining to the protected application data and places any updates to protected files in a queue on the Active server (that is, the send queue), ready to be sent to the Passive server with each request numbered to maintain its order in the queue.

Once the "send queue" reaches a specific configured size, or the configured time duration is reached, the update is sent to the Passive server, which places all the requests in an array of log files termed the "receive queue." The Passive server then confirms that the changes have been logged by sending the Active server an acknowledgment. The Passive server's "receive queue" is then read in numerical order and a duplicate set of file operations are applied to the disk of the passive server.

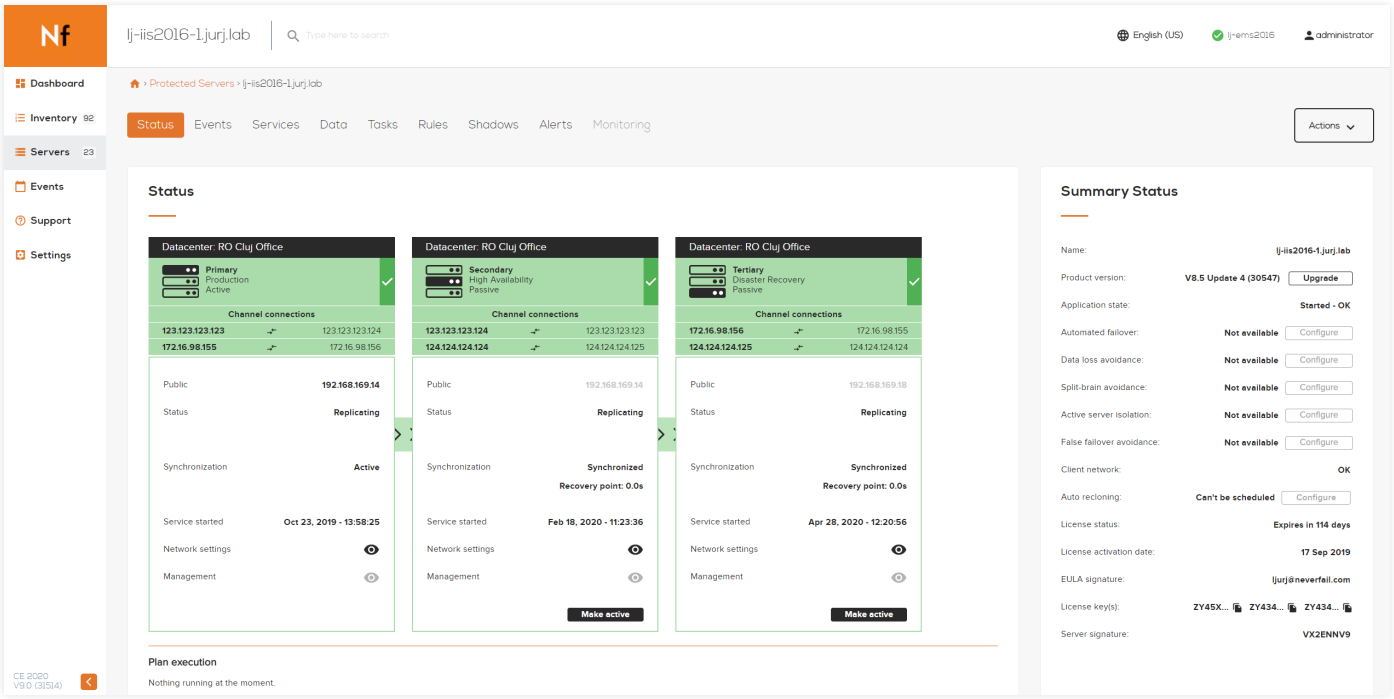


Figure 3: Engine Desktop Management Console

## Network Configuration

CE protected applications and servers rely on some Neverfail specific networking constructs for ensuring continuous availability across both high availability and disaster recovery deployments of the solution. These include:

- **Public IP or Public Name.** This is the IP address or fully-qualified domain name (FQDN) associated with a production application service that is critical for business operations and is registered with the organization's global DNS servers. This must be a static IP; DHCP is not supported.
- **Public NIC.** This is the NIC that is configured with a Public IP/Name for clients to communicate with the protected application or server.
- **Neverfail Channel IP.** This is the private (typically internal) IP address assigned to each server node within an CE cluster that corresponds to the Neverfail Channel network over which data replication occurs.
- **Neverfail Channel NIC.** This is the NIC that is configured with a Neverfail Channel IP on each server node within an CE cluster for communicating with other nodes over a private Neverfail Channel network. This can either be a dedicated NIC separate from the Public NIC (dual NIC configuration) OR the channel and public can occupy the same NIC when CE is deployed in a (single NIC configuration). See Figures 4 and 5.

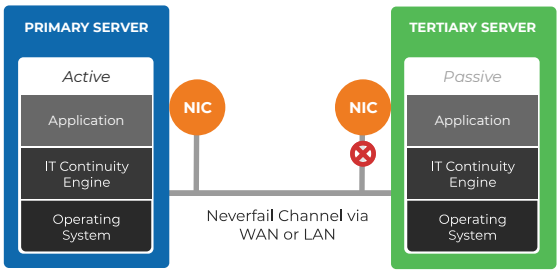


Figure 4: Single NIC Configuration

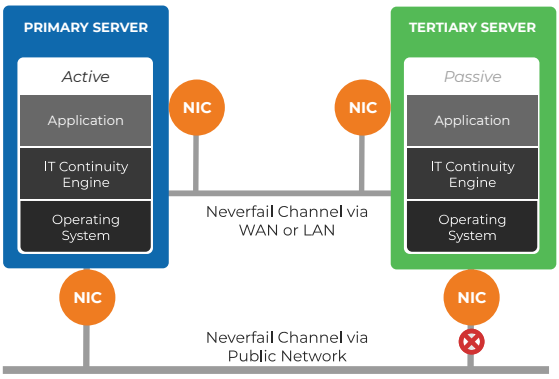


Figure 5: Dual NIC Configuration

When an CE cluster is configured for HA on the Local Area Network (LAN), the Public NIC on the passive server uses the same public IP address as the active server. However, the passive server is hidden and prevented from communicating with the live network by the means of an IP packet filtering system that is installed with CE. This packet filter prevents traffic from using the Public IP address from being committed to the wire. It also prevents NetBIOS traffic utilizing other IP addresses on the NIC from being sent to prevent NetBIOS name resolution conflicts.

When the CE cluster is configured for Disaster Recovery (DR) to a remote site with a different IP subnet, CE must be configured to use a different Public IP address for each of the Primary and Secondary servers. When a failover or switchover is performed, the DNS server will be automatically updated to redirect client connections to the new Active server instance at the DR site. These DNS updates are not required when the same subnet is in use at the DR site.

### *Take Note:*

Neverfail has implemented an improved version of its packet filtering system. The Windows Filtering Platform (WFP) is a set of API and system services that provide a platform for creating network filtering applications. Engine's implementation of this is based on its network packet filtering capability which prevents passive nodes from communicating on the network. This new integration allows Engine to control traffic in and out of a Windows 2008 R2, Windows 2012 or Windows 2012 R2 server dynamically without the need for complex device drivers designed for specific network interfaces.

## Scope of Protection

CE provides comprehensive protection against single points of failure to ensure that applications are continually running and operational for serving critical business services. These include:

- **Server Protection.** CE provides application continuity in the event of a hardware failure or an operating system crash. IT administrators have complete visibility using the CE management console into the health of the servers participating in the CE cluster. The Passive server actively monitors the Active server by sending frequent "I'm alive" heartbeat messages to it and expecting an acknowledgement in return over a network connection referred to as the Neverfail Channel. If the Passive server detects that the Active server is no longer responding, either due to a hardware failure or loss of network connectivity on the Active server, it can initiate a "failover" task. During a failover action, the Passive server is enabled to immediately take on the role of the Active server. The mechanics of failover are discussed in detail below.
- **Application Protection.** CE continually monitors protected applications and associated services that are running on the Active server. If any protected application should show failure signatures that indicate impending failure, or should the application suddenly fail, CE first attempts to restart the application on the Active server. If the application restart fails to remediate the application, CE can initiate a "switchover" task. A switchover action gracefully closes down any protected applications that are running on the Active server and restarts them on the Passive server along with any dependent application services required for application availability. The mechanics of switchover are discussed in detail below.
- **Network Protection.** CE proactively monitors the ability of the Active server to communicate with the rest of the network by polling up to three defined nodes around the network, including devices such as the default network gateway, primary DNS server, and the Global Catalog server at regular intervals. If all three nodes fail to respond, for example, due to a NIC or network switch failure, CE can gracefully switch the roles of the Active and Passive servers (referred to as a switchover) allowing the previously Passive server to assume an identical network identity to that of the previously Active server. After the switchover, the newly Active server then continues to service the clients.
- **Performance Protection.** CE proactively monitors system performance indicators to en-sure that your protected applications are truly operational and are performing adequately to deliver on expected quality of service to its end-users. This is made possible by CE's patented Application Management Framework (AMF) that is modular and extensible to include any Windows-based application, thereby providing comprehensive application protection coverage across thousands of applications being deployed in IT organizations. Additional details pertaining to AMF are described in the sections below.

The ability to express an application's key performance indicators are captured across pre-de-fined rules and adjustable thresholds in the form of a business service or application "plug-in". These plug-ins allow CE to monitor specific application attributes to ensure that they remain within normal operating ranges. Rules can be enabled or disabled as desired, and can be set to trigger specific preemptive corrective actions whenever these attributes fall outside of their respective ranges.

- **Data Protection.** CE ensures availability of application data as well as other file system data across all nodes within the CE cluster. CE can be configured to protect files, folders, and even specific registry settings of the Active server by mirroring them in real-time to the Passive servers. This means that if a failover or switchover were to occur, all files that were protected on the failed server would continue to remain available on the newly Active server after the failover or switchover event.
- **Site Protection.** CE provides application and business services availability even in the event of site-wide outages. By deploying the secondary or tertiary server within an CE cluster at a remote datacenter site, CE enables push-button disaster recovery capabilities along with WAN Acceleration support to ensure orchestrated recovery of individual applications or entire business services in designated priority recovery order at similar availability service levels as a local HA configuration.

## Failover and Switchover

Neverfail Continuity Engine offers different procedures to change the role of Active and Passive servers. Switchovers are associated with protected applications and failovers are associated with protected servers.

### Switchover Process

The switchover process is instigated either manually or automatically. In both cases the result is application control passes from an active to a passive server in a cluster. There are seven steps that document the switchover process. See Figure 6 for an illustration of the steps.

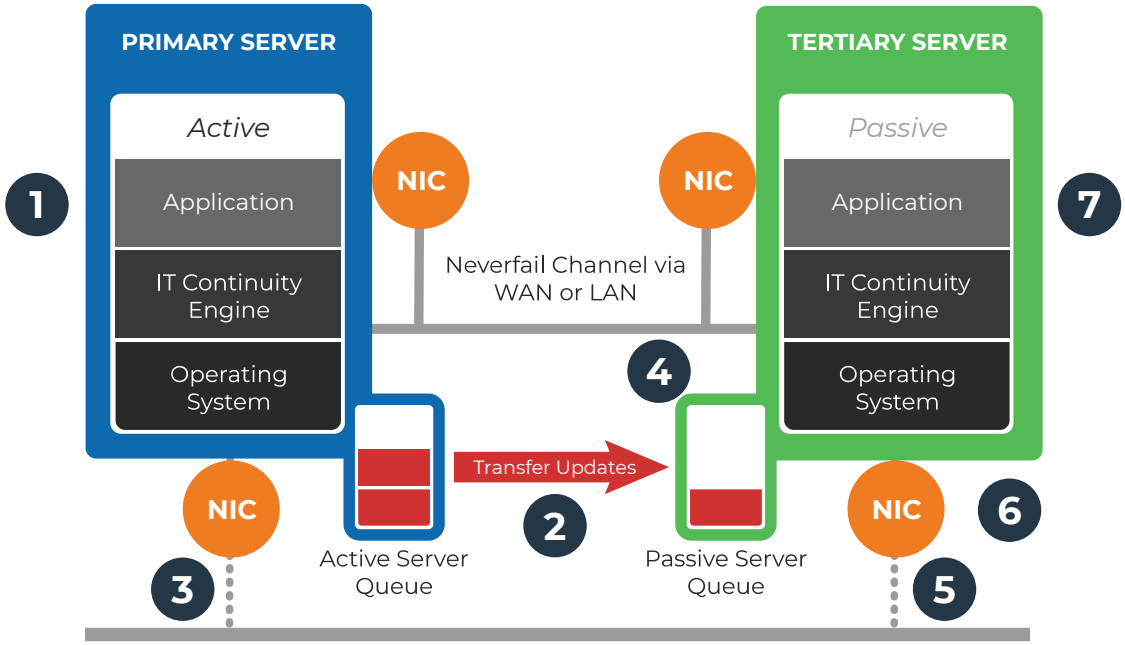


1. Stop applications.
2. Transfer updates.
3. Blocked from network. Server becomes passive.
4. Applied queued updates.
5. Expose to network.
6. Start interceping updates. Server becomes active.
7. Start applications.

Figure 6: Engine Switchover Process

**Managed Switchover.** A managed switchover can be initiated manually from the Neverfail CE management interface by selecting the appropriate “Make Active” button on the thick client or the Actions menu and selecting “Make Secondary Active” function. When a managed switchover is initiated, the running of protected applications is transferred from the Active server to a Passive server in the Cluster, that is, the server roles are reversed.

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

An auto-switchover is different from a managed switchover in that although the server roles are changed, the CE service is stopped on the previously Active server to allow the administrator to verify the integrity of the data on the newly Passive server and to investigate the cause of the auto-switchover. Auto-switchovers are similar to failover (discussed next) but initiated upon the failure of a monitored application. Once the cause for the auto-switchover is determined and corrected, the administrator revert the server roles to their original state.

## Failover Process

The failover process can also be instigated either manually or automatically. There are four steps that document the failover process. See Figure 7 for an illustration of the steps.

**Automatic Failover.** When a Passive server detects that the Active server is no longer running or responding to its frequent "I'm alive" heartbeat messages, it assumes the role of the Active server.

**Managed Failover.** A managed failover is similar to an automatic-failover in that the Passive server automatically determines that the Active server has failed, and can warn the IT administrator about the failure; but no failover occurs until the IT administrator chooses to trigger this operation manually.

1. Process replication queue updates.
2. Expose network.
3. Start intercepting updates. Server becomes active.
4. Start applications.

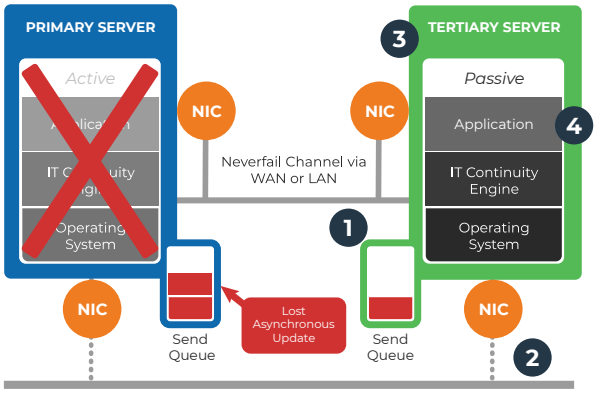


Figure 7: Engine Failover Process

## Extended Remediation Mechanisms

CE provides expanded remediation options for recovery starting with the CE 7.0 release that allow VMware administrators to leverage VMware vSphere technologies such as VMware HA, vMotion, Storage vMotion and enhanced vMotion for recovering from or avoiding any application continuity issues.

### VMware HA integration

- CE can be configured to trigger a VM Restart action utilizing VMware HA functions that execute a clean shutdown and restart of the Guest OS on the virtual machine experience an application issue or outage.
- Furthermore, CE can also be configured to reset the virtual machine using the VMware vSphere HA application monitoring mechanism, instead of using CE's native application remediation mechanisms.

### VMware vMotion integration

CE can be configured to trigger VMware vMotion or storage vMotion in response to any detected application degradation issues. Two examples would be:

- An application might be performing badly due to inadequate compute resources for the associated virtual machine (VM) on the existing ESX host, and the best remediation strategy would be to vMotion it off to another ESX host within the VMware vSphere Cluster that has adequate compute capacity.
- Another reason for why an application might be degraded is because of possible disk I/O (storage) latencies on the existing datastore and the best means to recover would be to relocate the impacted virtual machine disk(s) to an alternate datastore.

The above vMotion associated actions can be further configured with an application restart task to be executed by CE for a more aggressive remediation action (if necessary).

Continuity Engine delivers comprehensive breadth of protection simultaneously across all levels to ensure that all facets of the operating IT environment are maintained at all times, no matter what failure scenarios present themselves.

## Expanded Disaster Recovery for VMware Environments

CE provides expands disaster recovery capabilities of VMware's Site Recovery Manager (SRM) solution to include physical machine recovery along with its native support for virtual machine recovery. This is extremely beneficial for IT datacenter administrators who struggle continuously with deploying and managing multiple disparate disaster recovery strategies across their physical and virtual environments.

CE allows VMware administrators to create an "SRM Step" that can be added to a SRM Recovery Plan for allowing any physical and virtual machines protected by CE to participate in the failover and recovery of systems orchestrated by SRM. The "SRM Step" creates a script that designates which virtual machine replica of the physical server will be the Active server upon execution of the SRM Recovery Plan. You'd typically need to create two recovery scripts within the "SRM Step", one for failover and one for failback. These scripts are to be co-located on the same server as the SRM server(s) and are inserted into the SRM Recovery Plan(s) as a "Command Step" entry.

Once configured, VMware administrators can exercise all available SRM functionality, including recovery plan tests and actual recovery plan runs, in the context of physical and virtual machines protected by CE.

### *Take Note:*

- The standalone Neverfail Engine installer is available with Engine v6.7 (formerly known as Heartbeat) and prior versions
- P2P protection pairs are currently not supported with the Engine Management Service Enterprise or Desktop integrated installer package.

## Packaging & Deployment

CE is available in two flavors of deployment packages, depending on the desired mode of configuration. These include:

- CE Management Service integrated CE installer package
- Standalone CE installer package

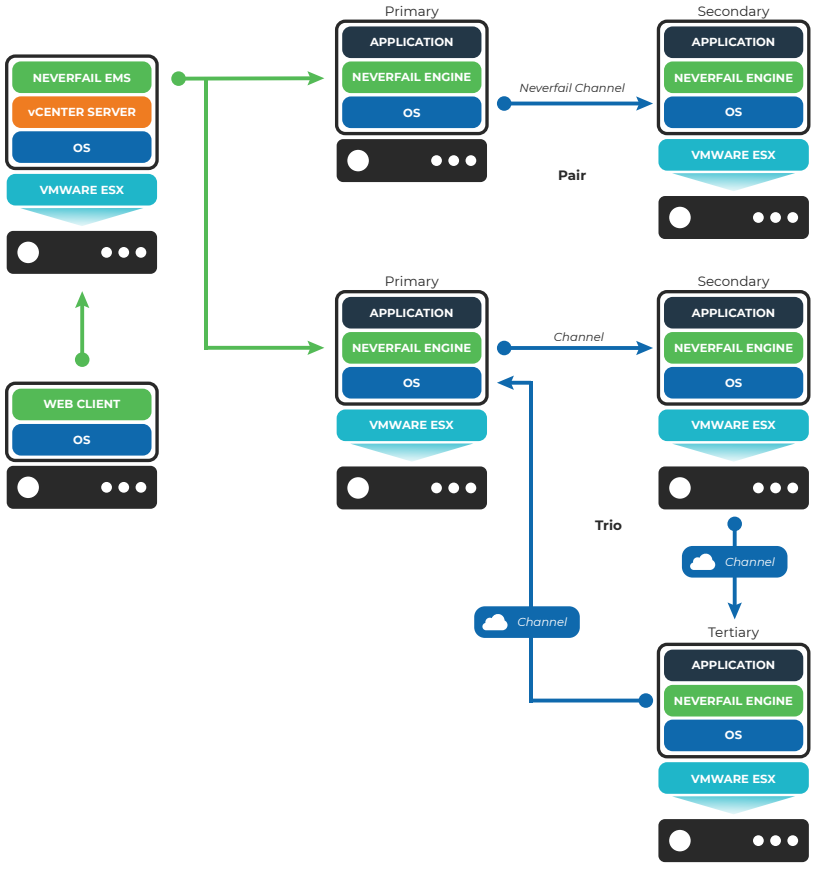


Figure 8: Engine Management Service Framework

## VMware CE Management Service Enterprise or Desktop Integrated Installer

This deployment method is best recommended for VMware virtualized environments when protecting applications running on either physical or virtual machines and when the standby secondary server is a VMware virtual machine, that is, a physical to virtual (P2V) or virtual to virtual (V2V) protection pair.

With this option, the Engine v7.1 installer package deploys an Engine Management Service (EMS). The EMS controls Engine services and configuration deployed across multiple remote protection pairs (or Engine clusters) to provide application aware continuous availability. As part of the installation process, EMS runs as a standalone web service so it can be managed by administrators.

The EMS web services can be accessed with any standard web browser. The EMS also provides a hook into the VMware vSphere Web Client for quick access to vSphere services. The EMS Enterprise Edition can be installed as part of a platform for deploying and managing all Engine deployments. In addition, the EMS Desktop Edition is designed for small implementations — typically five pairs or less that will not grow or need enterprise management. It should be installed on a separate Windows server with network connectivity to a remote instance of VMware vCenter Server.

With the EMS, administrators can deploy, configure and manage end-to-end Neverfail Engine protection, while carrying out most routine availability operations.

When deploying a new Engine protection pair using this installation method, Engine leverages its integration with VMware Converter to automatically create a clone of a target physical ma-chine towards deployment of the secondary server. See Figure 9.

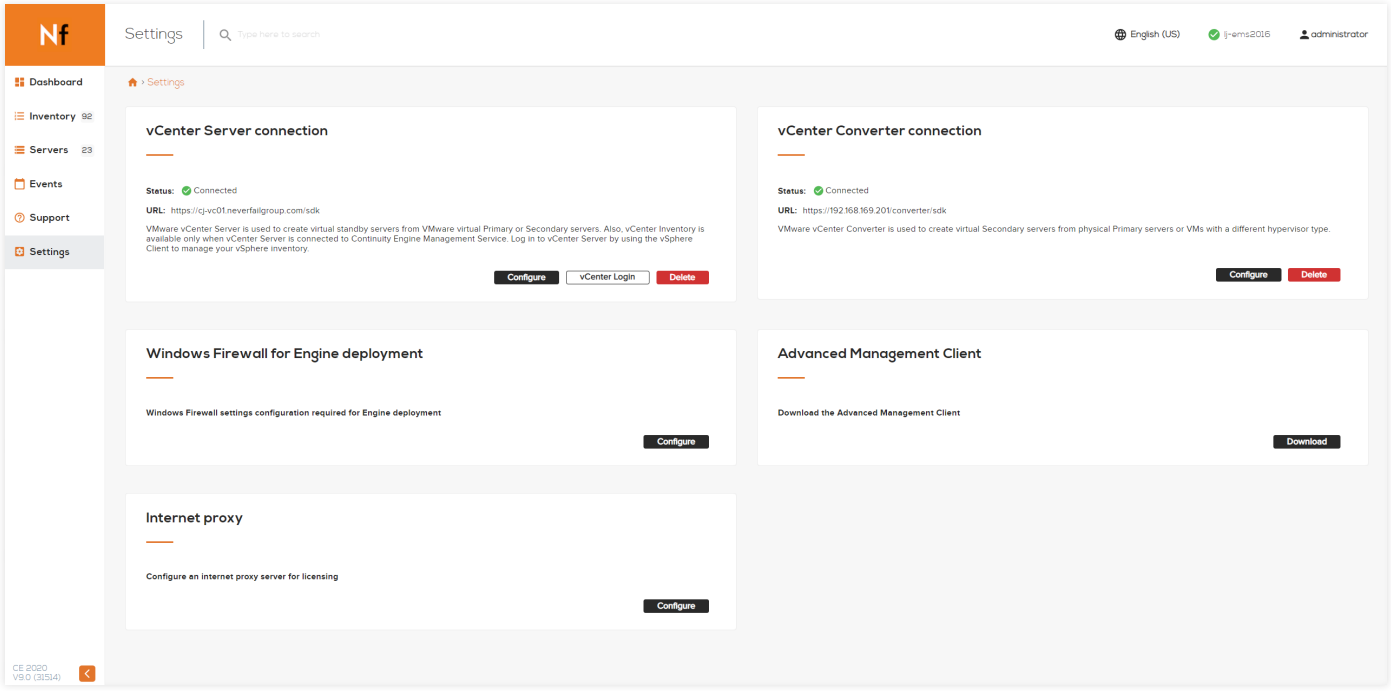


Figure 9: VMware Converter Integration

Likewise, Engine also automates secondary server creation for a virtual machine (V2V) protection pair by leveraging VMware vCenter Server's built-in V2V cloning mechanism, thereby significantly reducing the overall deployment process of the Engine protection pair. One can also discover and manage existing deployments of Engine clusters deployed using Engine (Heart-beat) 6.7 EMS Web Client interface for Engine. See Figures 10 and 11.

For more advanced configurations associated with an existing Engine cluster, a standalone CE Desktop Management Client is also available.

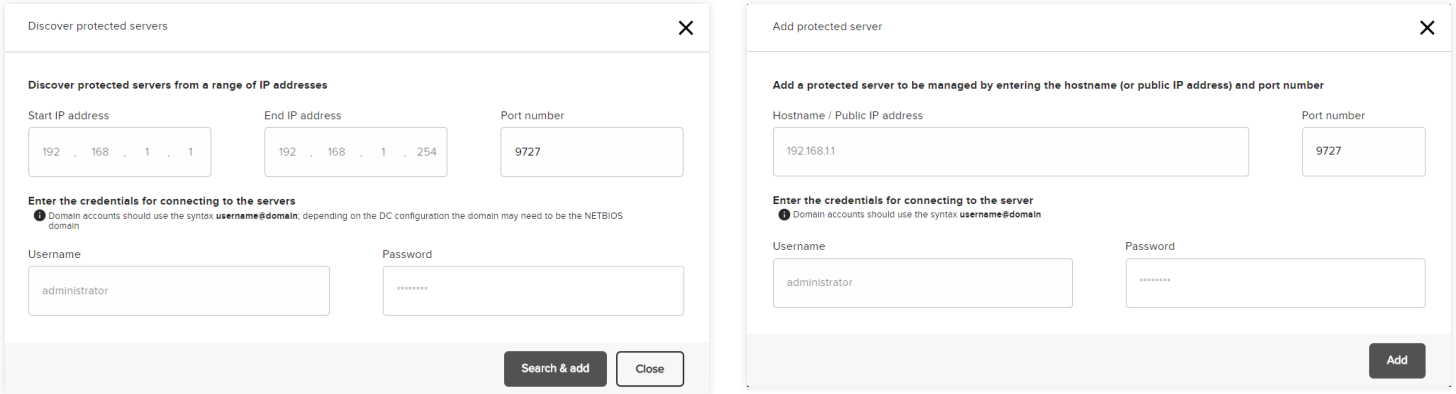


Figure 10: Manage Existing Engine Clusters

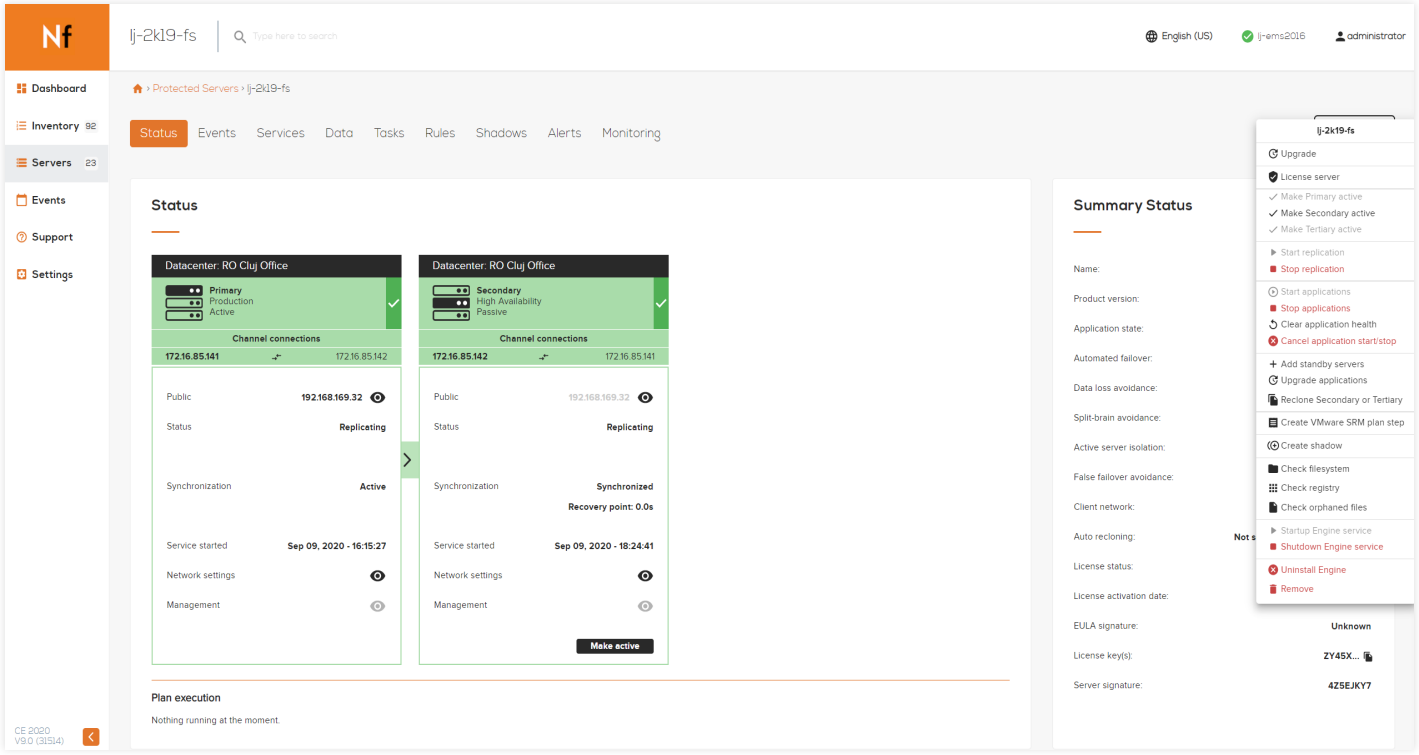


Figure 11: Basic Administration Using Engine Management Service

## Component Architecture

Neverfail Continuity Engine is comprised of several interrelated concepts and components that work in harmony to provide supreme levels of continuous availability and protection from a broad range of failures.

The architecture (Figure 12) outlines the key logical components of a single instance of Neverfail Continuity Engine.

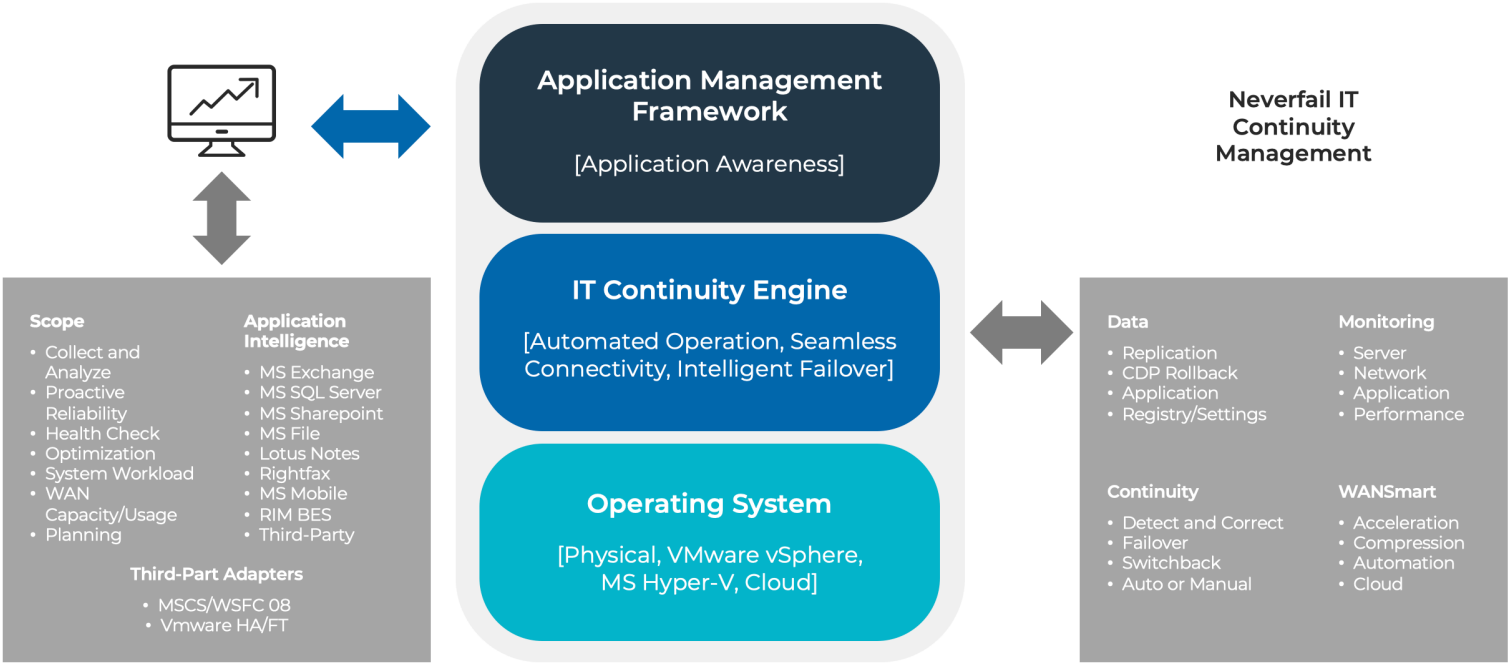


Figure 12: Engine Component Architecture

### Server Check, Optimization, and Performance Evaluation (SCOPE)

The first component to consider is SCOPE. It ensures the success of a Neverfail Engine deployment by providing current, accurate, and complete information about the server environment. It also provides detailed information about the current running state of your server environment and recommendations for optimizing your servers before installing Neverfail Engine.

It is embedded into the Engine 7.x product and automatically executed as part of the Engine deployment workflow. When utilizing the standalone Neverfail Engine 6.7 installer as the deployment option, SCOPE is available as a standalone software utility that can be manually executed on each server node that is intended to be configured as part of an Engine cluster.

### Engine

At the core, Continuity Engine component orchestrates operations and manages communi-cations between servers. It performs the complex core product functions of replicating data to and from other Engine protected servers at the Windows kernel level while live applications are running in the operating system layer above. Engine also manages coordinated failovers, switchovers and switchbacks between the various servers in a Continuity Engine cluster, synchronizing activity as required between active and passive instances.

### Application Management Framework

The other major component of Neverfail Engine is the Application Management Framework (AMF). The AMF is responsible for real time detection of faults, discovery of changes in the state of any protected application, managing interdependencies and live registration/de-registration of new protected applications through specific business application modules (plug-ins) or third party adapters.

Neverfail application plug-ins allow information about how best to protect a specific application to be encapsulated including inter-dependencies between related services and registry entries. And since AMF knows the state of each business application module plug-in, including the state of any associated resources such as services or registry entries it can be configured to manage interdependencies between applications as well.

For example, the Microsoft SQL Server application plug-in may detect that while SQL queries (read operations) on a critical production SQL database server are being executed correctly, no SQL write transactions have been recorded within the expected operational threshold period. Therefore expected services levels are not being met.

In this event Engine can be configured to automatically raise an alert through any of several notification methods, or even switchover to another instance of SQL Server. However, if we switchover SQL Server to another instance, we may affect the communications between SQL Server and a related application service such as a SharePoint.

As such, even though the SharePoint instance is operating correctly we may need to switch this application over at the same time in a coordinated fashion.

Finally, the AMF can be customized to protect any crash-consistent Windows application. Using a generic business application module ("plug-in"), Neverfail Engine can monitor and manage the state of any application's related Windows services. Custom tasks can also be implemented to provide application-specific monitoring of key performance indicators. This means we can extend the protection afforded by Neverfail Engine beyond that supported by standard business application modules (plug-ins) as long as the applications involved meet certain "restartability" conditions. See Figure 13.

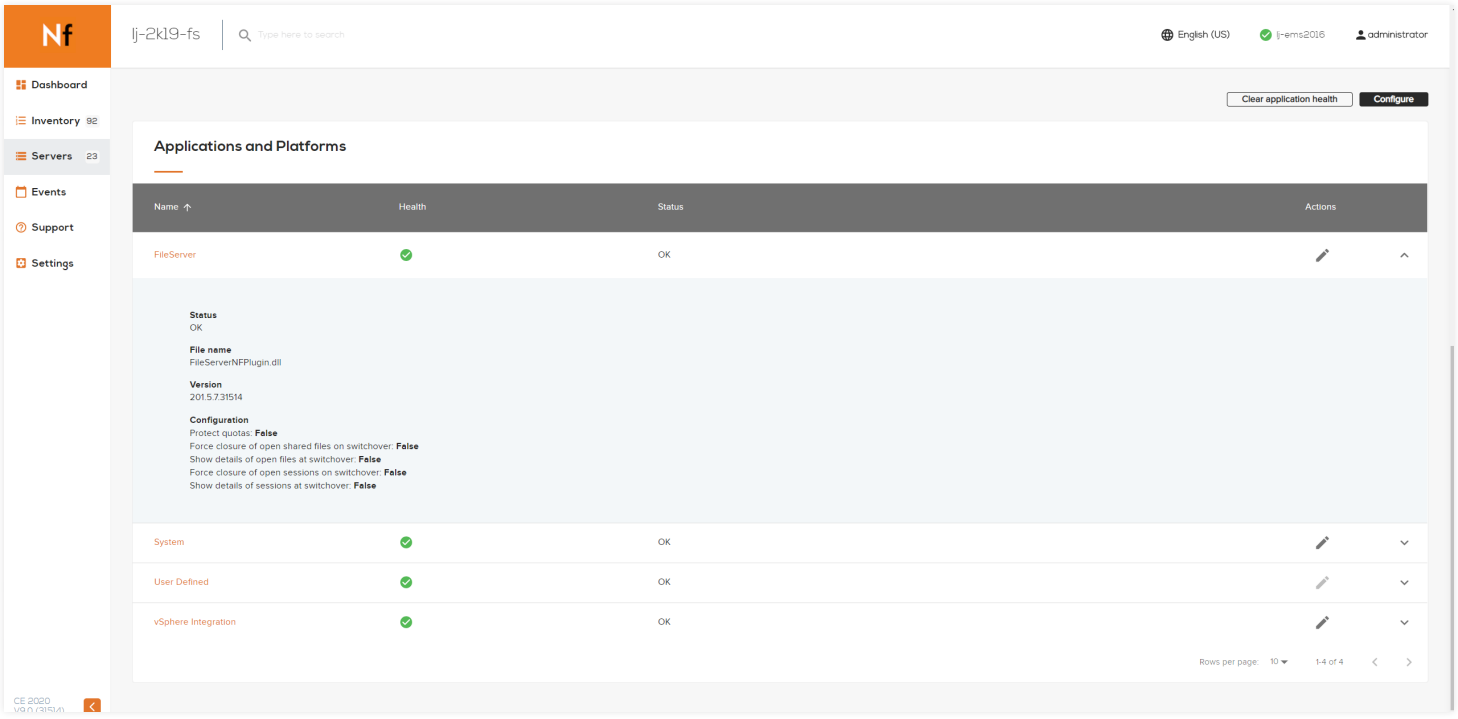


Figure 13: Application Plug-In

## Unified Availability Management Console

The Neverfail Engine Management interface allows multiple instances of Engine clusters to be monitored and managed from a single location. This presentation layer is currently available in the following formats:

- Standalone (Java-based) CE Management Client
- CE Management Service Web Client

By associating with the Public Name (FQDN) or Public IP of the protected server, the Neverfail Engine Management interface connects to the Engine instance on that server and provides visibility into the primary and secondary (and tertiary, if configured) servers within the Neverfail Engine cluster. The role, status, and state of each server is readily available on the overview page. Depending on the type of management client being utilized, additional navigational menu options or tabs provide visibility into details regarding the state of the application, the network, the data, and replication processes. Configuration changes for existing protected Engine clusters can be performed using the standalone Neverfail Engine Management Client without disrupting the replication process.

Furthermore, unique application groups may be defined in the Engine management console for the purpose of both monitoring and managing more complex applications. When application groups are defined within the Neverfail Engine Management Client, warnings and alerts from any member server will roll up to the group level. The status of all groups are visible simultaneously within the Neverfail Engine Management Client, and any issues within the protected infrastructure are quickly identified. Additionally, application groups can optionally be configured to switch over to a remote DR location as a collective logical container. By coordinating the switchover of the entire application group, Neverfail Engine reduces the complexity and the dependencies related to moving an application to a DR site.

## WANSmart

Continuity Engine offers WAN Acceleration capabilities (WANSmart) that, in addition to the default data compression, provide an on- the-fly data deduplication algorithm to dramatically reduce the amount of bandwidth required to support replication over a WAN. Similar to other hardware-based WAN optimization solutions that are typically very expensive to procure and implement, Continuity Engine's WANSmart is a software implementation across servers participating in the Engine cluster, reducing the amount of data that needs to be sent to the remote site by up to 80%.

When servers in the Engine cluster have a high change-rate on their protected data, WANSmart helps to minimize the cost of the connection between the datacenters as well as ensures that changes arrive at the DR site faster than if data was sent in an uncompressed state.

## Summary

For more than a decade, Neverfail has helped companies implement best-in-class business continuity strategies and continuous availability solutions to over 2,500 companies from the Fortune 50 to SMB businesses. By leveraging years of experience and a mature set of proven technologies, Neverfail Continuity Engine applies a broad range of availability concepts to delivering continuous availability and disaster recovery for all types of applications and workloads in any IT organization.

Combining modular elements of data replication, server clustering, network management, and application-centric monitoring of key performance indicators, Continuity Engine provides total protection for your critical business services and associated applications, ensuring 24/7 avail-ability regardless of any threats to uptime. It is the only application-aware IT continuity solution focused on eliminating the risk of downtime before it can impact any critical business service or application that the business depends on.

## About Neverfail

Neverfail enables businesses to achieve 100% uptime through the world's most resilient business continuity and secondary storage solutions. Made for mission-critical businesses, Neverfail solutions mitigate the risk of downtime in the face of any potential outage. By delivering seamless business continuity, we empower our partners and clients to realize their full potential without the risk of downtime.

Honeywell vmware Mitel McKESSON citi Bank of America