

Eliminating End User and Application Downtime

Protecting Mission-Critical Multi-Tier Business Applications with Neverfail and Virtualization

July 2010





Table of Contents

Introduction	3
Protecting Mission-Critical Applications in a Mixed Physical, Virtual World	3
Neverfail Continuous Availability for Mixed Physical, Virtual Environments	5
Neverfail Mixed Physical, Virtual Deployment Scenarios	6
Neverfail All Virtual Deployment Scenarios	8
Neverfail Local and Remote Site Availability Deployment Scenarios	8
Neverfail Physical to Virtual Migration Scenarios	9
The Importance of Fast, User Transparent, Automatic Failover & Failback	9
The Importance of Heterogeneous Replication	10
The Importance of WAN Optimization for Remote Site Availability and DR	10
Neverfail Continuous Availability - Built for Mission Critical Applications	12
Application-Aware, Centralized Management and Control	13
About Neverfail & Virtualization	14

Introduction

The rapid adoption of server virtualization has enabled many organizations to significantly reduce their IT capital expense, operating costs as well as adapt to changing business requirements more easily and improve overall server availability.

However, when it comes to extending the benefits of server virtualization to their mission-critical business applications, many organizations often struggle with how to re-architect their environment and strike the right balance between performance, scalability, manageability and most importantly, high availability and disaster recovery.

This white paper outlines many of the common deployment scenarios that such organizations face and describes how Neverfail Continuous Availability and virtualization can be used to deliver a more complete, consistent and cost effective availability solution for mission-critical, business applications.

Protecting Mission-Critical Applications in a Mixed Physical, Virtual World

Most modern business applications incorporate multiple tiers of interdependent servers (Web Front End, Application, Database/Storage etc.) each of which have differing roles and therefore different performance, management and availability requirements. However, to ensure the continuous availability of the entire business application, each of these systems must be managed and protected in a consistent manner in order to minimize the risk of business application downtime.

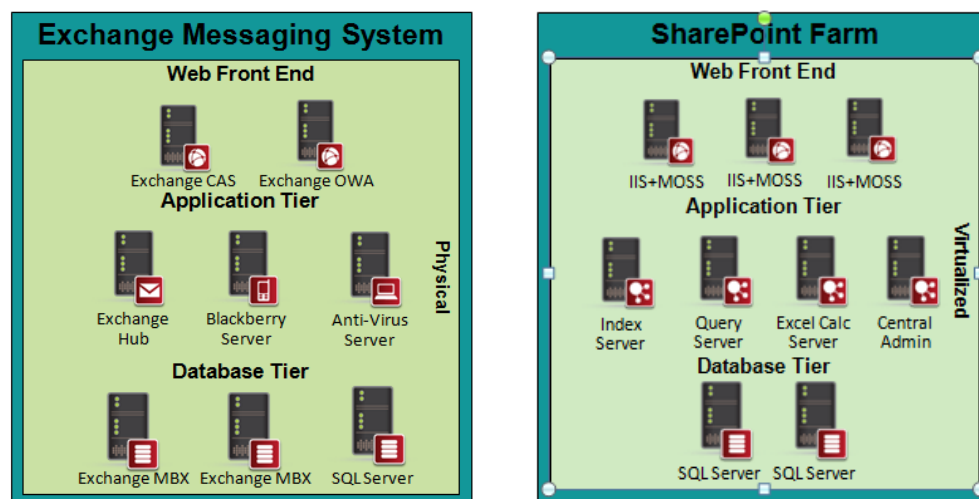


Figure 1 – Exchange & SharePoint: Mission-critical, multi-tier business applications

When organizations look to virtualize their mission-critical business applications, careful consideration must be given to ensure that nothing impacts performance, scalability and manageability and that nothing compromises the continuous availability of the entire business application. Such concerns have led many organizations to part-virtualize their server farms, often leaving the Database tier completely untouched as well as many other high performance application servers.

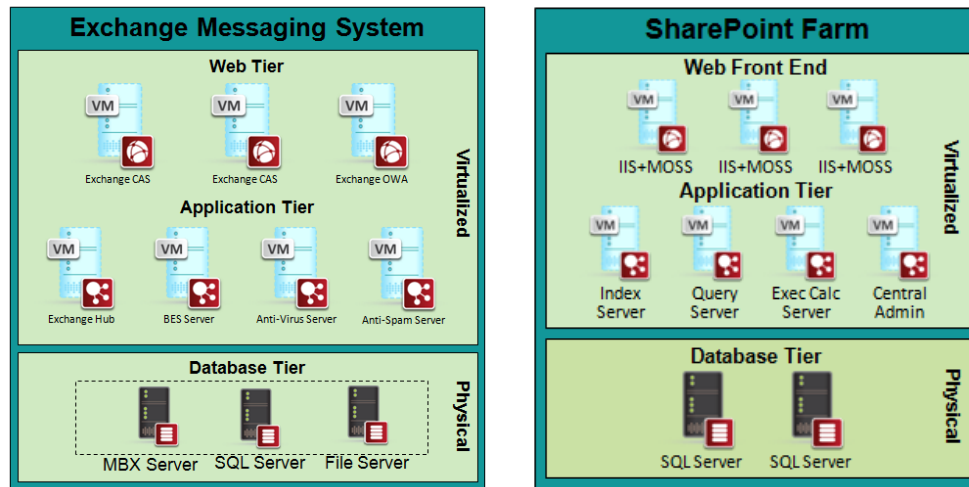


Figure 2 – Typical deployment scenarios: Part-virtualized, multi-tier, business applications

However, while part-virtualizing the server farm ensures performance and scalability, it does require that the availability of virtualized and non-virtualized servers be handled in different ways. This is due to the fact that the various server virtualization platform availability features, such as VMotion, HA/FT, CSV, Live Migration etc. cannot protect both virtual and physical servers. This situation is compounded, when looking to design for remote site availability or disaster recovery, as neither VMware SRM nor Microsoft DPM support mixed physical, virtual, local and remote site availability configurations either.

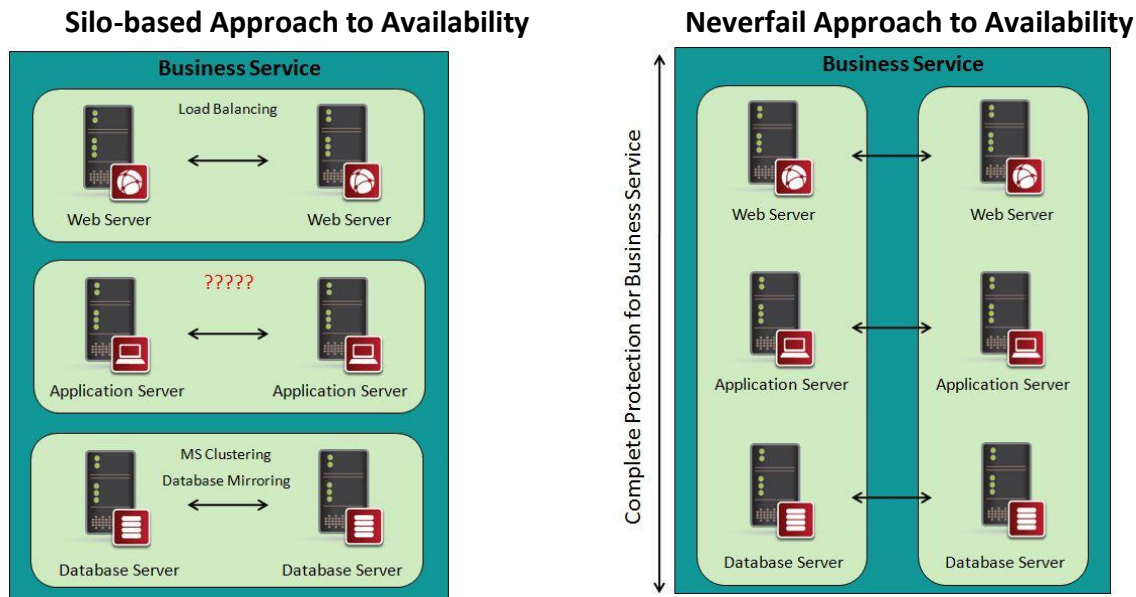


Figure 3 – Conventional, silo-based approach to availability vs. Neverfail Continuous Availability

As illustrated in the above graphic, taking a “Silo-based” approach to availability and disaster recovery creates significant additional complexity, as it requires that each solution to be installed, configured, managed and maintained in different ways. Unfortunately, this complexity will invariably result in an increased risk of downtime for the mission-critical business application that you are trying to protect!

Neverfail Continuous Availability for Mixed Physical, Virtual Environments

The VMware vSphere and Microsoft Hyper-V platforms incorporate many platform-level technologies (VMware HA, FT, VMotion, Storage vMotion, SRM, Live/Quick Migration etc.) that can significantly improve overall system resilience and server availability.

However, when it comes to protecting your mission-critical business applications, in a mixed physical, virtual environment, across both local and remote sites (for disaster recovery) these platform-level technologies do not provide *complete* protection for your mission-critical business applications.

In addition, unlike Neverfail, these platform-level availability technologies do not protect your mission-critical business applications from the most common causes of downtime, such as application failure, service failure, performance degradation, configuration drift, user error etc.



Product Availability/Protection	Microsoft Hyper-V	Microsoft DPM	VMware HA/FT	VMware SRM	Neverfail v6.2
Local High Availability	✓	X	✓	X	✓
Remote Site Protection	X	✓	X	✓	✓
Local & Remote Site Availability	X	X	X	X	✓
Virtual Machine Level Protection	X	✓	X	✓	✓
Physical & Virtual Server Protection	X	✓	X	X	✓
Storage System Failure Protection	X	✓	X	✓	✓
Network Failure Protection	X	X	X	✓	✓
Operating System Failure Protection	X	X	Basic	X	✓
Application Aware Failure Protection	X	X	X	X	✓
Configuration Drift Protection	X	✓	X	✓	✓
Automatic & Manual Failover/Failback	X	X	X	X	✓
Embedded Replication (for DR)	X	✓	X	X	✓
Embedded WAN Optimization	X	✓	X	X	✓

Figure 4 –Neverfail Continuous Availability vs. VMware and Microsoft Hyper-V availability features

As can be seen from the above table, when it comes to providing *complete* protection against downtime for mission-critical business applications in a mixed physical, virtual environment, Neverfail’s Continuous Availability Suite is the only solution that covers all of the common high availability and disaster recovery deployment scenarios.

Neverfail Mixed Physical, Virtual Deployment Scenarios

Neverfail provides the flexibility and freedom of choice to implement both local and remote site high availability solutions that meet your technical and business requirements, whatever the scenario. If local high availability is required and the primary production servers are physical machines, Neverfail gives you the option to choose physical or virtual servers as the secondary server for the high availability pair.




Figure 5 – Neverfail protecting a physical server with a secondary physical server (P2P local HA)

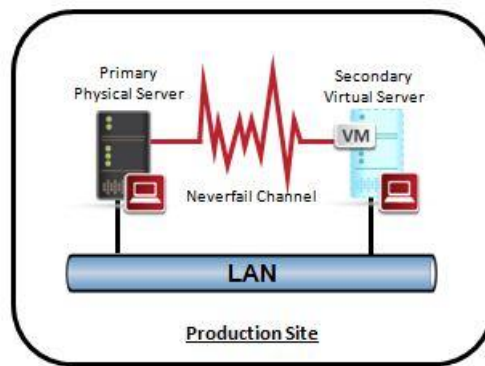


Figure 6 – Neverfail protecting multiple servers in a mixed physical, virtual environment (P2V HA)

The Virtualization of secondary servers (as detailed above) is a popular choice for local high availability, for the following reasons: -

- There is no application disruption as the primary applications remains untouched
- The consolidation of the secondary HA server (via virtualization) reduces the total cost of HA
- The secondary servers facilitate the testing of virtualized mission-critical business applications
- The Neverfail vXtender licensing lowers the cost of protecting the virtual secondary servers

Neverfail All Virtual Deployment Scenarios

Where virtualization is used to implement a local virtual cluster for server high availability, using Neverfail to replicate and maintain mission critical applications at a remote location, can bring immense benefits in terms of disaster availability. It is likely that only a few mission critical applications require Continuous Availability. By focusing on these applications, the business can be protected without requiring complex infrastructure upgrades such as stretched or replicated SANs which would be more suited to full site protection.




Figure 7 – Neverfail protecting multiple servers in a purely virtual environment (V2V HA)

Neverfail Local and Remote Site Availability Deployment Scenarios

With Neverfail, once a local high availability pair has been established, adding a remote site availability or disaster recover site becomes simple. It's just a matter of implementing the appropriate environment at the remote DR site, using Neverfail Continuous Availability to synchronize the data and application configuration.




Figure 8 – Protecting mixed physical, virtual servers with local & remote site availability (P2V2V)



Neverfail Physical to Virtual Migration Scenarios

Neverfail also reduces the effort associated with migrating from a physical server environment to a virtual server environment as a by-product of the local high availability and replication capabilities, inherent within the Neverfail product. After installing the secondary virtual servers and synchronizing the data, switching over to the secondary machines and making them the production servers is as easy as pressing a button. This process can be performed in the local environment, the remote environment or both.

Once it is determined that there is no application performance degradation, the physical machine can be shut down and removed altogether, or re-purposed as needed.

The Importance of Fast, User Transparent, Automatic Failover & Failback

In order to meet the aggressive Recovery Time Objectives (RTO) that mission-critical business applications require, Neverfail provides fast, user transparent, automatic failover and failback for the entire business application.

Neverfail's ability to perform fast, user transparent, automatic failover of mission-critical business applications, in mixed physical and virtual environments, is what sets the Neverfail solution apart from other Data Protection and Recovery solutions such as VMware Site Recovery Manager (SRM) & Microsoft Data Protection Manager (DPM). For example, The VMware SRM product is primarily designed to automate the site disaster recovery process ("Run book automation") for up to 1,000 virtual machines, a job which it does extremely well. However, when it comes to providing Continuous Availability for mission-critical business applications when a disaster strikes, the process of bringing the business application back up and running on the remote site could take anywhere between 20 minutes to several hours. By comparison, with Neverfail's Continuous Availability solution the same mission-critical business application (taking the same environment and configuration) would be up and running again in a matter of seconds or at most, a few minutes.

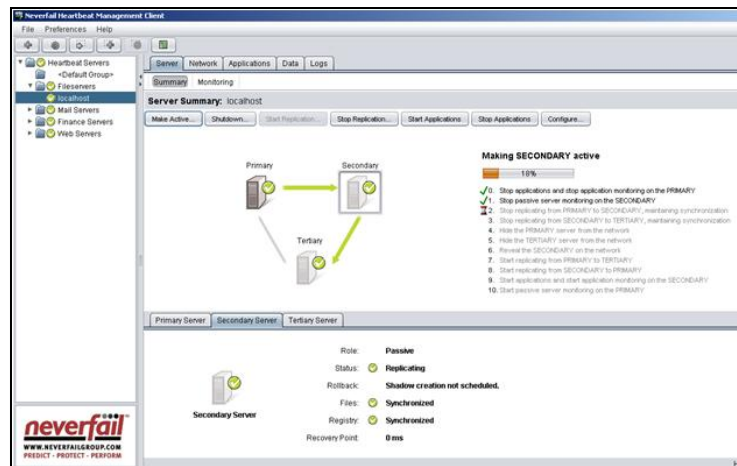


Figure 9 – Server failover, showing the failover progress, replication state and actual Recovery Point

The Importance of Heterogeneous Replication

Neverfail Continuous Availability incorporates heterogeneous replication as standard. Unlike VMware SRM, which requires SAN based replication in order to provide site protection, Neverfail’s replication technology can work with Direct Attach Storage (DAS) or Storage Area Network (SAN) attached storage and enables any system to be replicated to any other system, regardless of the make or model of the underlying storage. This “open” approach to system replication provides organizations with the flexibility and freedom of choice to configure system replication in the most cost effective manner (e.g. high end EMC to low end NetApp) which can deliver tremendous cost savings.

The Importance of WAN Optimization for Remote Site Availability and DR

When it comes to designing a remote site availability or disaster recovery solution, the number of options available will often be constrained by network bandwidth and/or the physical distance between sites, servers or virtualization hosts. After all, not everyone can afford the high costs associated with deploying a low latency network (required for SAN based synchronous replication) or justify moving large amounts of traffic over wide area network (WAN) links.

Neverfail Continuous Availability incorporates WAN optimization which makes implementing a remote site availability or disaster recovery solution for business applications affordable and easy to implement by overcoming the common issues of replicating over long distances and/or limited bandwidth.

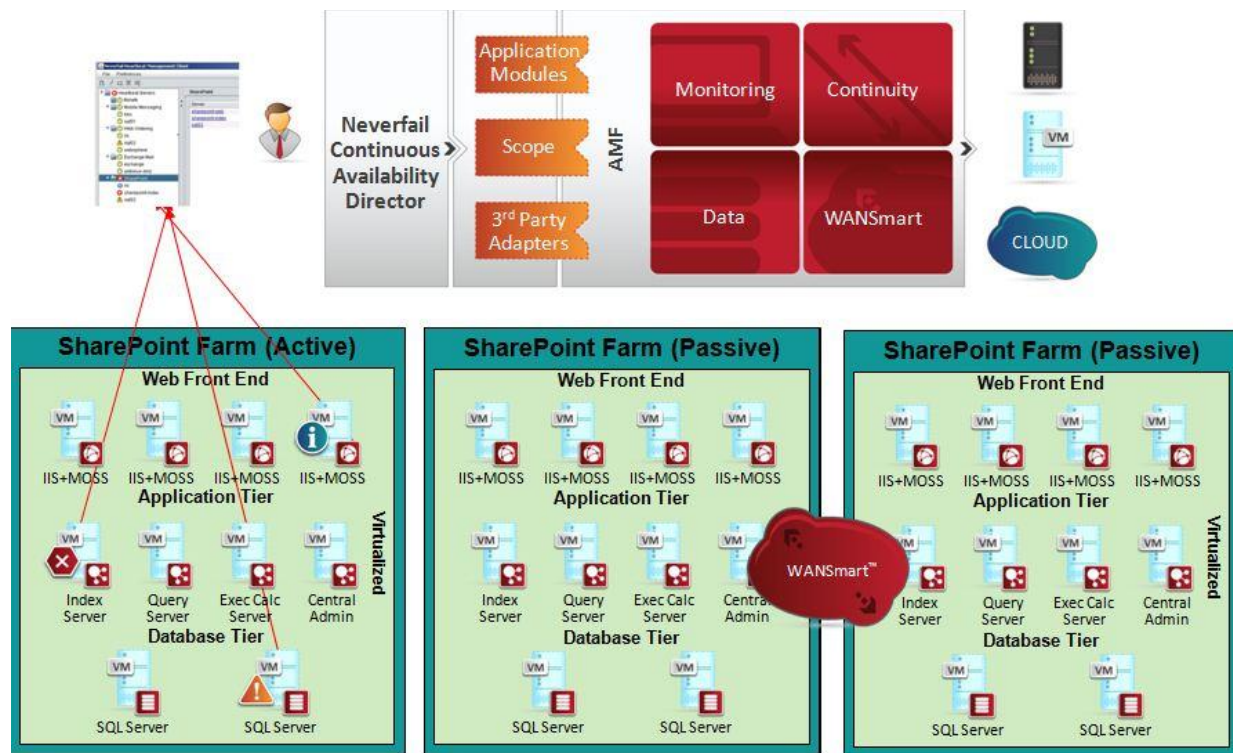


Figure 10 –Using Neverfail WAN optimization (WANSmart) to improve RTO & RPO

WANSmart replication combines the proven benefits of data compression with Neverfail's advanced de-duplication technology. When combined, these two technologies accelerate the replication of business application(s) across the WAN, significantly reducing the bandwidth requirements (by a factor of up to 30) as well as enabling significantly better Recovery Time (RTO) and Recovery Point Objectives (RPO) to be achieved.

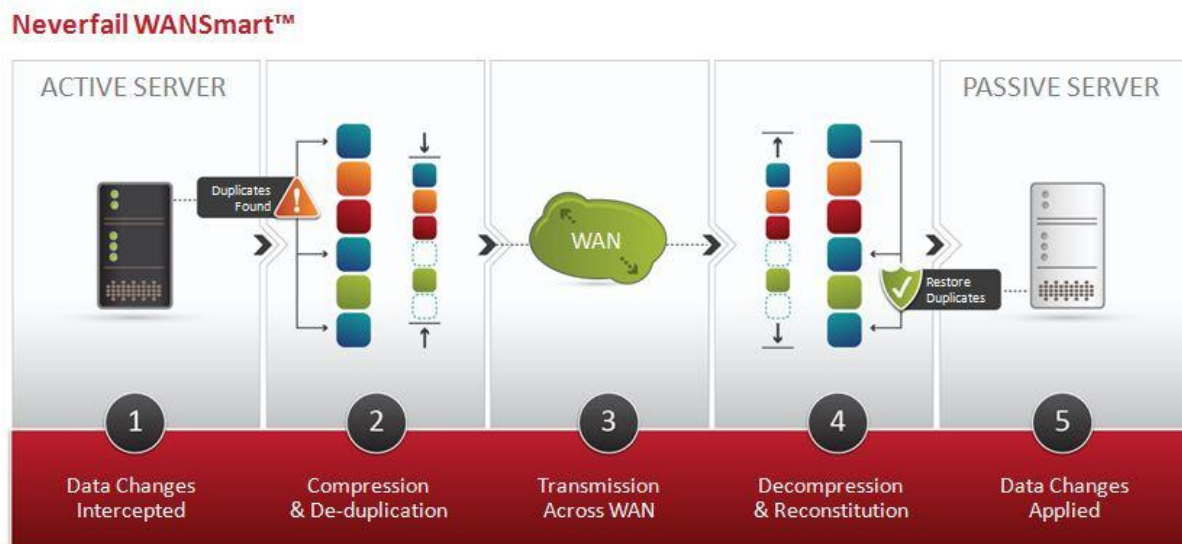


Figure 11 –WANSmart de-duplication and compression reduce bandwidth requirements by up to 30x

To reduce RTO's further, WANSmart also performs the necessary DNS updates automatically, which ensures that the mission-critical business application WAN failover is as seamless as possible and that the highest level of continuous availability is achieved for the business users.

Neverfail Continuous Availability - Built for Mission Critical Applications

Neverfail's Continuous Availability Suite is unique in that it proactively manages and protects the entire business application ecosystem whether built upon physical servers, virtual servers or a combination of both. It can protect mission-critical business applications against any type of failure, such as the physical server hardware, storage, network, operating system, service, application, configuration drift, human error etc. as well as a total disaster, such as a complete site failure. If a problem occurs, Neverfail ensures your applications are protected.

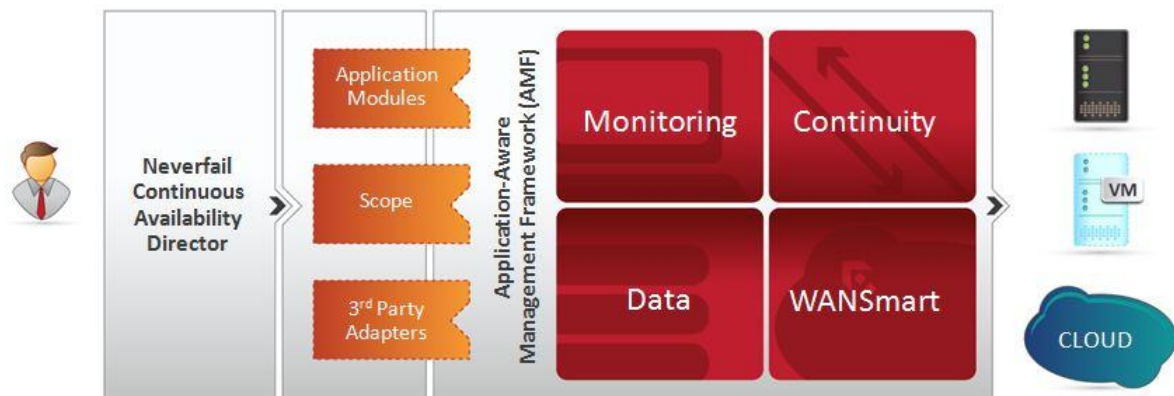


Figure 12 – The Neverfail Continuous Availability Suite

Application-Aware, Centralized Management and Control

The Continuous Availability Director management console provides organizations with an enterprise-wide, business-centric view of mission-critical business applications and IT services. The view spans both physical and virtual environments, as well as providing a consolidated view of local and remote systems. It has a flexible approach, which allows logical grouping of application, database, messaging and other servers. Groupings provide a way of visualizing interdependencies across servers to provide a central console where important events, alerts and overall health of the business application can be viewed.

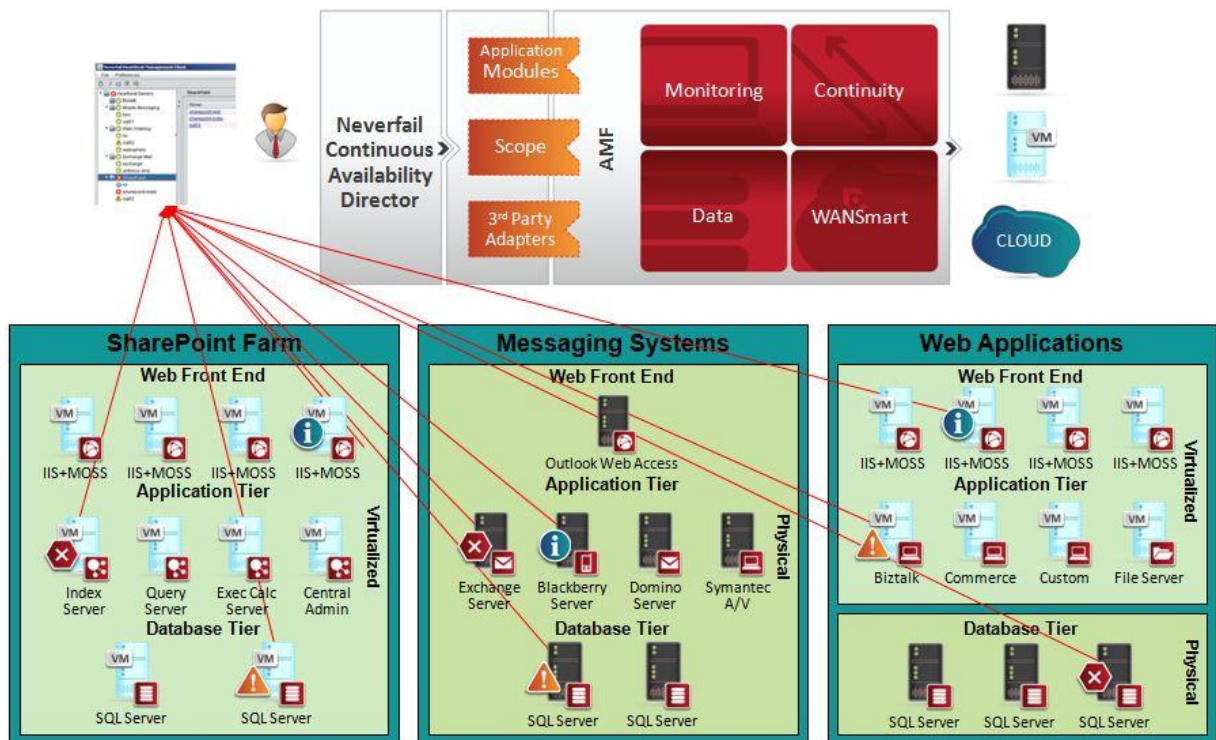


Figure 13– Continuous Availability Director: Protecting multiple, mission-critical business applications

About Neverfail & Virtualization

Neverfail's mission is to provide the most sophisticated and comprehensive Continuous Availability solutions within the marketplace which encompass both physical, virtual and cloud environments. To this end, Neverfail has forged strong OEM and technology partnerships with leading virtualization and cloud companies, such as VMware and Microsoft.

One example of this is VMware, who have licensed Neverfail software to build VMware vCenter Server Heartbeat which provides Continuous Availability for the vCenter Server. The vCenter Server is the central control point for the entire vSphere environment; if the vCenter Server goes down, then many of critical components in vSphere, such as vMotion, SRM and VIEW, no longer operate.

Look out for more exciting virtualization announcements as Neverfail continues to work towards a world without downtime.

Neverfail...Eliminating User Downtime



All rights reserved.

Neverfail® is a trademark of Neverfail Group Limited. All other trademarks are trademarks of their respective companies. No part of this publication may be reproduced, transmitted, transcribed, or translated into any language or computer language, in any form or by any means without prior express, written consent of Neverfail Group Limited.

www.neverfailgroup.com