

Eliminating End User and Application Downtime

Architecting the “Right” Continuous Availability and Disaster Recovery Environment

March 2010

Table of Contents

Introduction	3
Where to Start	3
Moving to Continuous Availability	4
Data Centric (Backup/Restore & Disaster Recovery)	5
Application Centric (Clustering)	5
Business Centric (Continuous Availability)	5
Putting it all Together	6
Finishing Up	9

Introduction

Architecting the “right” continuous availability (CA) and disaster recover (DR) environment can be a difficult, if not daunting task. However, by understanding which applications in your environment are critical to the successful operation of your business, the task becomes a bit less troublesome. Short of performing a complete business impact analysis (BIA), which is a time consuming and costly exercise itself, there are ways to determine what your availability needs are and in which direction you should turn for a solution to solve the problem.

Neverfail provides a CA and DR solution that ***ensures your users stay connected to their applications and data, regardless of the type of failure or the failover location.*** With Neverfail’s unique, patented, proven technology and comprehensive management capabilities, Neverfail can solve even the most challenging CA and DR problems.

This white paper will discuss how Neverfail can provide the foundation for a solid CA and DR environment, allowing you to architect the “right” environment.

Where to Start

The process of designing and architecting a reliable, comprehensive, cost effective CA and DR solution can only begin by understanding your application uptime and data recovery requirements. In other words, what is your stated service level agreement (SLA) for application and data recovery time and recovery point objectives (RTO and RPO)?

RTO deals with the how quickly an application has to be back up and running after a failure, while RPO deals with the amount of data that can be lost. RTO and RPO can also be defined in terms of money. That is, ***how long can you afford to be down*** and ***how much data can you afford to lose.*** When defined in monetary terms, it becomes quite clear as to where the focus needs to be.

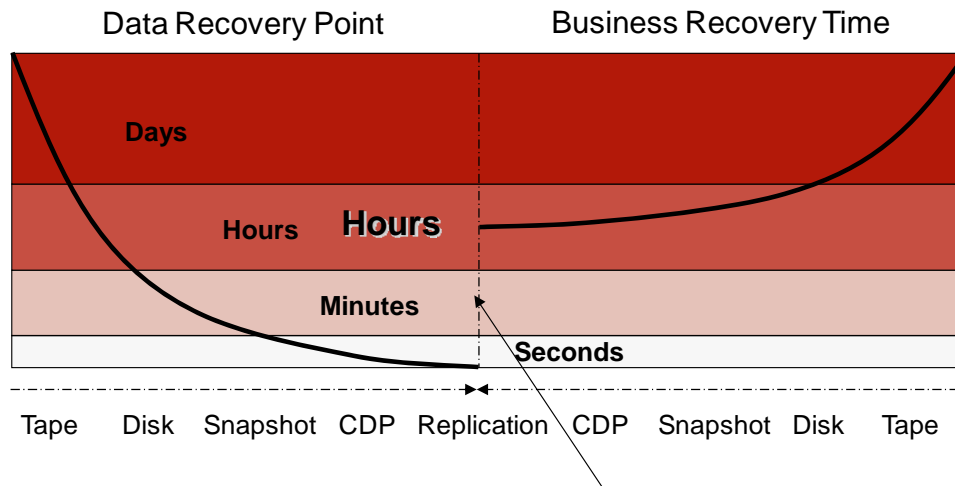
There are many solutions that provide data RPO in the seconds range. However, by comparison, there are very few solutions that can also provide the same, for application RTO. Neverfail is one of the select few, to have developed the necessary technologies required to deliver a solution that provides application RTO and RPO in the “near-zero” range, effectively eliminating application downtime.

So, how do you know which applications are the most critical?

Which require RTO/RPO in the “near zero” timeframe?

The answers to such questions are actually pretty simple. Which applications create the most help desk calls when they are unavailable? Which applications, when unavailable, cause the company to lose money, or productivity, or customer confidence?

Making these decisions and being aware of the impact of application downtime will help you to determine which applications require the most attention.



Continuous Availability - Eliminates the Availability Gap

Figure 1 - The Availability Gap

Establishing which applications require the most stringent RTO and RPO is the first step in determining what type of solution and architecture will work best for your organization.

Moving to Continuous Availability

Over the years, a vast array of point products has been developed to solve the specific availability problems of the day. Whilst there has always been the need for a solution to bring applications back online quickly after a failure, with as little loss as possible, finding a complete solution to the problem has remained elusive. However, more recent technology advances (in replication, clustering, de-duplication etc.) mean that it is now possible to move away from the conventional Recovery Centric approach to solving the problem, towards a more Availability Centric approach.

The following diagram illustrates the evolution away from recovery/data centric solutions to availability/business centric solutions, such as those provided by Neverfail.



Figure 2 - The Evolution of Continuous Availability

Data Centric (Backup/Restore & Disaster Recovery)

Solutions falling in this category range from traditional “data to tape” style backup/restore to advanced CDP and “disk-to-disk” DR solutions. Though these solutions may have some use for some applications, they are really only focused on ensuring that the **data** can be recovered in a timely fashion. Replication and snapshot solutions provide additional data recovery capability and many times are the basis for other availability solutions, but in most solutions using these technologies, there is no focus on or awareness of the application to ensure it is brought back online in a timely fashion. Trying to design a solution for continuous availability, using this type of solution, will leave a big gap in the application RTO requirement.

Application Centric (Clustering)

Most clustering solutions, whether active/passive, active/active or multi-node provide a hardware, fault tolerant and high availability capability that focuses on ensuring the application is available **if some hardware components should fail**. In many cases, this may be adequate to reach the RTO requirement. However; most clustering solutions have a number of limitations, not the least of which is a shared storage architecture. Shared storage architectures introduce a “single point of failure” which can be detrimental to application uptime. If the shared storage should fail, no amount of switching to other nodes in the cluster will bring the application back on line. Furthermore, most clustering solutions do not provide any data protection. In order to ensure RPO requirements are met, another solution would be required, creating additional administrative overhead.

Business Centric (Continuous Availability)

In today’s highly competitive business climate, traditional data centric and availability centric solutions do not deliver against the ever increasing RTO/RPO requirements that most organizations now face, let alone meet regulatory or compliance obligations.

Continuous Availability solutions must **eliminate end user and business application downtime completely**. Neverfail’s Continuous Availability suite achieves this, by providing a unique blend of continuous data protection, replication, de-duplication, performance monitoring, application awareness and user transparent failover/failback capabilities, all from within a single package. In addition, a centralized management system provides a “single pane of glass” to view of all of the underlying business applications, servers and services within the protected environment.

By combining all of the above technologies, Neverfail delivers a comprehensive solution that enables organizations to meet their RTO/RPO objectives, as well as meet ever increasing regulatory and compliance obligations.

In order to have application RTO approach zero, all of the elements that make up the application must be managed in a way that ensures their continued availability. The diagram below illustrates the various architectural elements, which are required to deliver continuous availability.

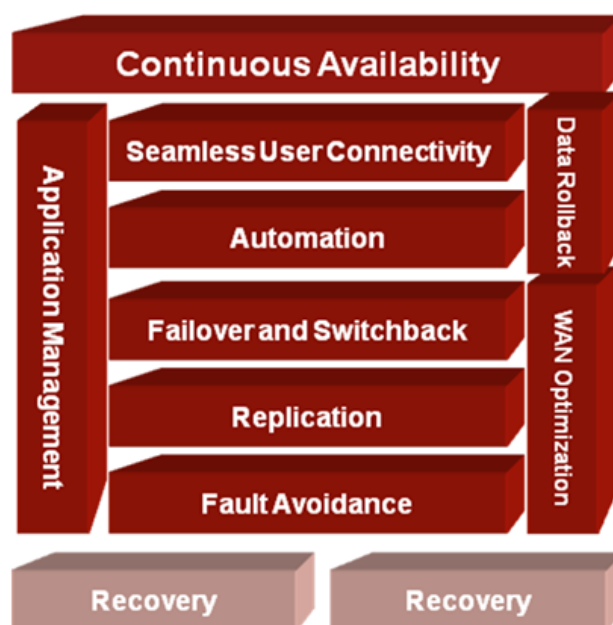


Figure 3 – Architectural Elements for Continuous Availability

Any solution that does not provide all of these capabilities, will lack the ability to provide a “no down time, near-zero RTO”.

Putting it all Together

Even though decisions have been made regarding the criticality of applications in an environment, there is still an overriding fact that cannot be ignored – application downtime is intolerable for most application owners, users, customers and management. Given this, it seems like ensuring the application is well protected and monitored, so that a pre-emptive action can be taken to solve a problem before it becomes an issue, is key.

Neverfail has built its solution so that the application user is the focal point for the operation. As shown below, the core of the Neverfail solution is based around Data, Monitoring, Automation, Continuity and WANSmart functional areas. These core operations ensure that data is replicated and protected. It ensures that the applications are monitored for, not only proper operation, but proper performance as well. Continuity provides the detect and correct mechanisms – remediating problem before they become major issues and WANSmart for acceleration, compression and de-duplication over a wide area network.

Built on top of the core operation is the Application-Aware Management Framework (AMF). The AMF is the “brains” of the operation. It provides the foundation for application specific plug-ins. The plug-ins, tailored for many different applications, understand how the application should run, what services make up the application and what performance metrics are key to the application running successfully. By monitoring these items, Neverfail can make decisions on how to handle everything from a minor application “glitch”, resource limitation through to a complete site outage.

The entire operation is managed by the Continuous Availability Director, which provides a single point of management and control as well as a global health view across the entire environment.

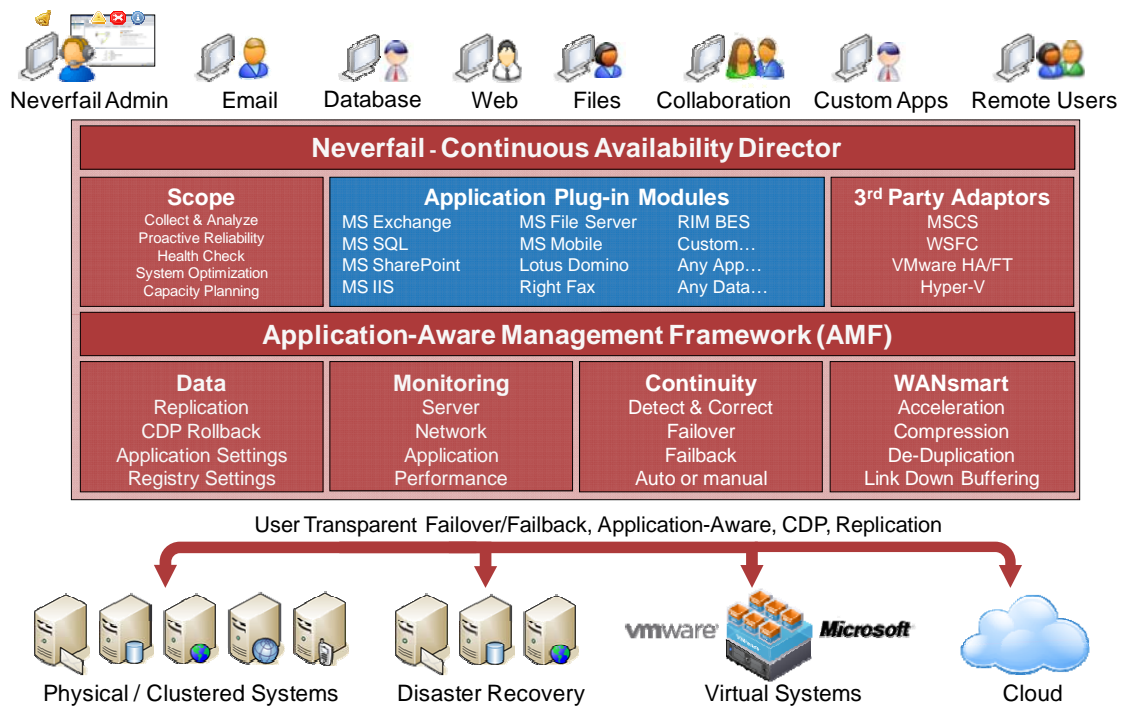


Figure 4 - Neverfail Block Diagram

Often, there are many outside factors that also play a part in determining what architecture will satisfy the continuous availability and/or disaster recovery requirements of your organization.

Continuous availability solutions may be chosen to protect against legacy hardware failures that can cause more frequent application down time. Other factors like being in a hurricane zone for instance, may lead to a disaster recovery or remote decision. Migrating to a virtualized environment or moving data centers can also factor.

Whatever the scenario, Neverfail provides the ability and flexibility to implement an environment that fits your requirements. If continuous availability is required and the production servers are physical machines, Neverfail gives you the option of choosing physical or virtual servers as secondary machines. Virtualization is a very popular choice for secondary servers because of the improvement in the technology and the cost savings associated with this technology.

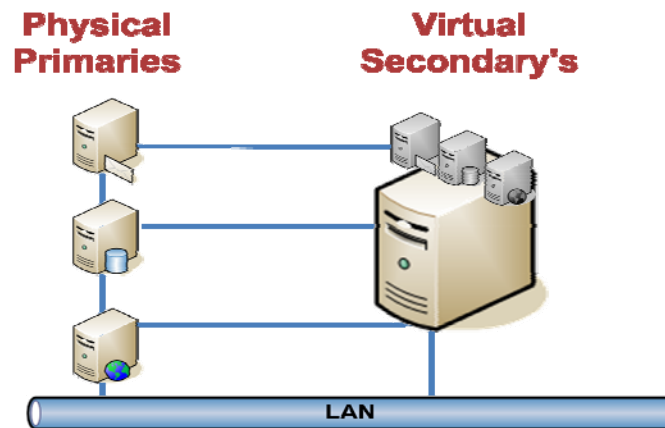


Figure 5 - Continuous Availability Implementation

Adding DR capability then becomes simple. It's just a matter of implementing the appropriate environment at the remote DR site, using Neverfail to synchronize the data and application configuration.

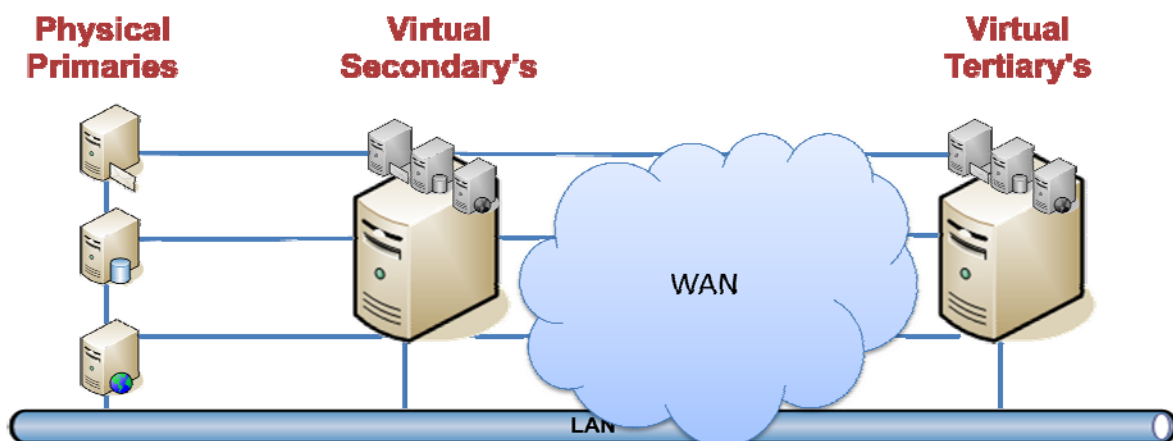


Figure 6 - Continuous Availability and Disaster Recovery Implementation

Neverfail also eases the effort associated with migrating from a physical server environment to a virtual server environment as a by-product of the continuous availability and disaster recovery capability inherent within the solution. 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 re-purposed as needed.

Finishing Up

There are many solutions available within the market today, that claim both high availability and disaster recovery capabilities, and so it's easy to understand why making a decision and architecting the "Right" solution can be difficult. However, in the end, there are really only a few questions that need careful consideration. Does the solution you're considering: -

- Meet the *application* RTO and RPO requirements for your business?
- Support the "Right" architecture for your current environment? If not, what is the most cost effective way to get there?
- Provide complete, application-aware protection against downtime for any Microsoft based Business Applications such as Exchange, SharePoint, BizTalk, Communication Server, Commerce Server, SQL etc.?
- Provides the freedom of choice to deploy Microsoft based Business Applications in local high availability, stretched farm and/or remote site disaster recovery configurations
- Provide easy, affordable, user transparent WAN Failover for Microsoft Business Applications
- Significantly reduce the ongoing management burden of protecting your Microsoft Business Applications, by providing centralized, monitoring, management and protection
- Focus on eliminating user downtime

With Neverfail, you can have confidence that the solution will satisfy your requirements by providing the "right" solution for your environment, *ensuring that your end users and business applications are managed and fully protected against downtime – whatever the cause.*

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