Neverfail Solutions for VMware:

Continuous Availability for Mission-Critical Applications throughout the Virtual Lifecycle

Table of Contents

Virtualization

One of the hottest topics in today's IT corridors is the uses and benefits of virtualization technologies. IT professionals everywhere are implementing virtualization for a variety of business needs, driven by opportunities to improve server flexibility and reduce operational costs.

With each and every server that becomes virtualized, new challenges are manifesting themselves – most notably risk management concerns involving the physical-to-virtual migration of critical applications and the ongoing protection of application service delivery within the virtualized environment. The latter of these, continuous application availability and disaster recovery, are perhaps the most mission-critical to organizations focused on business performance.

Benefits of Virtualization

The fundamental aim of virtualization is to enable server consolidation and containment by running software applications on fewer servers. The virtualization market leader, VMware, Inc., provides exactly this capability not only for the data center but also for development and QA test labs, and virtualized desktops. This allows IT administrators to make the most efficient use of their hardware resources and development operations to quickly provision test systems for a variety of combined OS and software platforms.

The clear benefits of virtualized technologies are that they allow companies to increase their hardware utilization (since most servers are typically underutilized) and consequently decrease their capital and operating costs. Resource optimization tools such as VMware's® VMotion improve overall service levels by dynamically and intelligently moving workloads to the host system that can handle it the best.

Of course, the fundamental design of this approach, running multiple virtual servers within a single physical host system, inherently introduces new IT risks. By consolidating multiple servers that perform a variety of business critical functions onto a single host system, availability of that system becomes a significant risk point. VMware has taken some huge steps to mitigate these risks with high availability and business continuity products to protect the *physical* platforms on which virtualized applications run.

As a VMware Technical Alliance Partner, Neverfail completes the solution by layering *application-level* monitoring and protection onto VMware's virtual infrastructure platform. Additionally Neverfail provides options for high availability and disaster recovery that do not rely on shared storage, thus providing greater flexibility and lower cost options, particularly for remote disaster protection. This combination offers customers complete protection to ensure the continuous delivery of services for critical business applications.

Continuous Availability Defined

The best way for IT to ensure consistent business performance is through the implementation of a continuous availability solution that focuses on the business need and end user experience; rather than protecting data, hardware and applications as isolated components. True continuous availability of business critical applications requires a solution which ensures that business is not disrupted by IT

outages and keeps users seamlessly and transparently connected to critical applications. Continuous Availability solutions address continuity across all possible forms of outage.

The reasons for outages vary, and may include everything from data loss, server failure, application failure, or network failure to planned downtime, application performance degradation and corruption or a complete site outage (disaster). It's a fact of life that IT outages will happen; resultantly, a critical goal should be that when an outage occurs, it should not result in business disruption and downtime – end users should simply continue operating as if nothing has happened, thus delivering on the promise of consistent business performance.

Virtualization and Potential Risks to Continuous Availability

Continuous availability technologies provide uninterrupted uptime for mission-critical applications Continuous availability demands there is no interruption to services provided by mission-critical applications. This means email must keep flowing, websites must keep working, and a user's BlackBerry® and other mobile devices must continue to provide access to information. This must happen in all situations from component failure through user error to natural disaster. The critical component in all of the above is the application itself. It's not possible to deliver continuous availability without understanding the performance and availability of individual applications.

Risk #1 – P2V Migrations

The first risk to availability of mission-critical applications is the physical-to-virtual (P2V) migration process. The process itself will depend on the criticality of the application and whether downtime can be considered acceptable. As it involves cloning systems and data, extensive testing is required to ensure a seamless, uninterrupted switchover. A staged approach allows time for evaluating the overall benefits and impact, and of course, continuous availability of the virtualized environment must apply from day one. From the standpoint of risk management, it would be ideal to adopt a migration strategy that also accommodates an immediate and automatic failback of applications onto the original physical platform should unexpected problems manifest themselves within in the virtual environment. In fact many Neverfail customers will carry out P2V migrations in stages. Firstly, they will use Neverfail to clone the application to a virtual machine giving immediate high availability options. They will then carry out a series of "live-tests," for example, switchovers during planned maintenance are often implemented in order to gain experience on the virtual machine, before switching back seamlessly to the physical server and analyzing the test results before a final switchover.

Risk #2 – Application Failures

VMware's Virtual Infrastructure contributes class-leading levels of high availability focused primarily on maintaining hardware availability. If a host system critically fails, all virtual machines can be restarted on alternate hardware using VMware HA. If hardware resource levels drop below acceptable thresholds, VMotion can dynamically reassign the load to higher capacity servers automatically, before the resource issue degrades service. While this may overcome the risk associated with the *hardware* and operating system that supports a virtualized application, it does not account for the largest risk category with the highest probability of failure – problems that occur to the *applications* running within the virtual environment.

For example, if the Windows® operating system running within a VM exhibits any form of failure (such as a blue screen); or if the virtualized network fails or is misconfigured; or if any of the mission-critical applications running within the VM itself exhibit problems, VMware will be completely unaware of the situation. Given that these types of issues are so much more common than physical hardware failures, it is clear that virtualization on its own still does not address all the requirements for continuous availability – keeping your business running and end users connected no matter what the cause of an outage.

Risk #3 – Single Points of Failure

Continuous availability also demands that disaster scenarios are covered. In a crisis, keeping communication lines open is a high priority. This means email systems, such as Microsoft® Exchange, and mobile platforms, such as RIM BlackBerry, have to be protected off-site. VMware alone cannot provide this level of protection but instead must rely on third-party solutions to replicate data across multiple locations. These technologies are complex and expensive to deploy and operate, and require significant time to restore service after a site failure, not to mention a SAN itself is a single point of failure.

Resource optimization using VMware's VMotion technology can be a great aid in delivering well-configured, optimized systems. VMotion allows VM's to be provisioned to other host servers to provide load-balancing, and also provide a level of protection should indications show a potential hardware failure. However, VMotion's capability depends upon a persistent connection to VMware VirtualCenter. Should VirtualCenter fail, the ability to manage the VMware infrastructure is severely diminished. Therefore, VirtualCenter itself is a single point of failure with no effective mechanism for high availability or disaster recovery.

Neverfail and VMware: The Best of Both Worlds

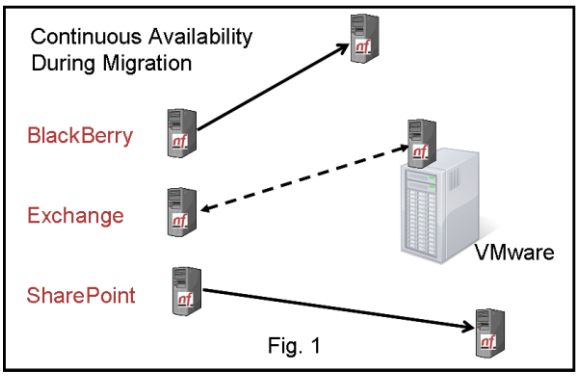

Fig. 1

Neverfail is designed to provide continuous availability of mission-critical applications. This means keeping users connected and productive through potential service interruptions that may result from hardware failure, application failure, natural disasters and many other situations. Neverfail's solutions are designed to allow seamless switchover *and* switchback should maintenance or other work be required on the primary servers. Neverfail does not rely on shared storage architectures; instead, it uses advanced replication technology to maintain cloned system and data copies locally or remotely. Therefore, it can be relied upon for protection against disasters. It also comes into its own as part of a physical-to-virtual (P2V) process as highlighted previously.

The ability of Neverfail to clone an application and keep data synchronized with a replica copy means multiple critical applications can be migrated onto a VMware platform, and run in parallel, without interrupting service delivery from the physical platform during the migration. Once the application has been cloned onto the virtual platform, Neverfail then provides automated switchover capability by moving control to the virtual clone. The key benefit of this approach to P2V migration is that the

service can also be switched back from virtual-to- physical without service interruption, should anything go wrong. Neverfail's ability to synchronize data back as part of this switchback means no data will be lost. Unlike other single-direction P2V migration strategies, Neverfail gives the customer complete confidence that the migration is reversible without risk. This approach also provides an immediate opportunity to implement a high availability solution as the first step in the virtualization lifecycle (see Figure 1).

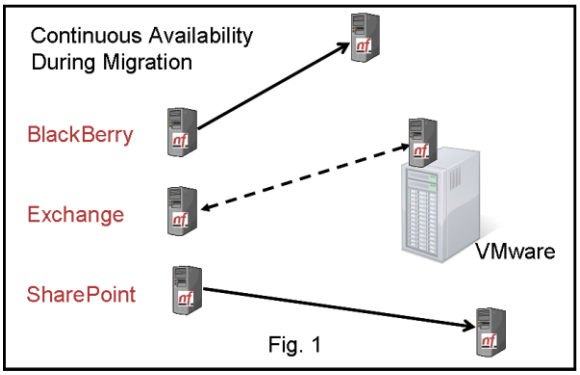

Fig. 1

Adding Application Awareness

The Neverfail technology offers full monitoring of all critical components of the application stack within the virtual machine itself. If a primary server (physical or virtual) fails for any reason, Neverfail will detect the problem and failover seamlessly to the secondary system. If the primary server exhibits networking issues and is unable to connect to critical network resources throughout the rest of the IT environment, Neverfail will detect this condition as well and seamlessly switch application workload to the secondary system. Even if the applications running within a virtual machine exhibit problems or any form of performance degradation, Neverfail can identify these problems, dynamically attempt to resolve the problem, and then ultimately failover to the redundant secondary system when all automatic preemptive solutions have been exhausted. **All of these failover and switchover operations are completely seamless to the end user, allowing them to remain productive without the need to reconfigure or even restart their client applications.** These same processes provide a perfect environment for migration testing as well.

Virtual-to-Virtual (V2V) Protection

Now that the organization is benefiting from a fully monitored and automated high availability solution that protects both individual applications and VM's, it doesn't make sense to sacrifice these facilities when moving from primary servers to a virtual VMware implementation. A major benefit of Neverfail is its ability to support failover from virtual-to-virtual environments (see Figure 2). This is a natural extension that delivers the benefits of virtualization without compromising the need for continuous availability. It also makes no demands on particular storage requirements allowing the re-use of direct attached storage, if appropriate.

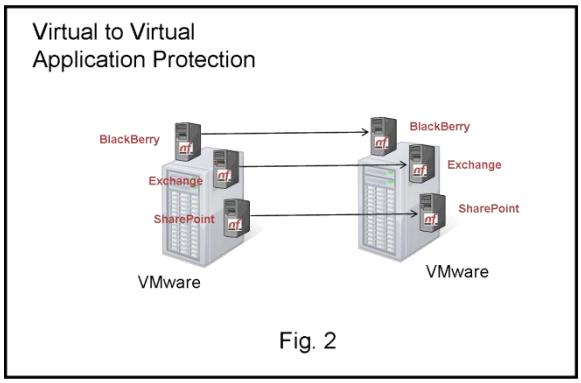

Fig. 2

Taking this one step further, Figure 3 demonstrates a typical architecture with mixed virtual and physical environments being protected at an application level. If desired, VMware HA can be used as an additional level of protection to shield the physical server from catastrophic hardware level failure. This is not really necessary because even if the primary server fails, Neverfail is managing the applications in the secondary VM's and will detect the failure and continue running without interruption to users.

Although catastrophic failures may be in the class of a component failure, it is also advisable to protect against disaster scenarios. To guard against this situation, VMware requires third-party solutions that are often costly and complex. Only then is it possible to provide a remote site provisioning implementation.

Fortunately, Neverfail does not rely on a shared storage infrastructure. Replication is carried out by Neverfail itself, and optimized over WAN connectivity. As a result, it is simply a configuration decision to provide the ability for offsite replication (see Figure 3).

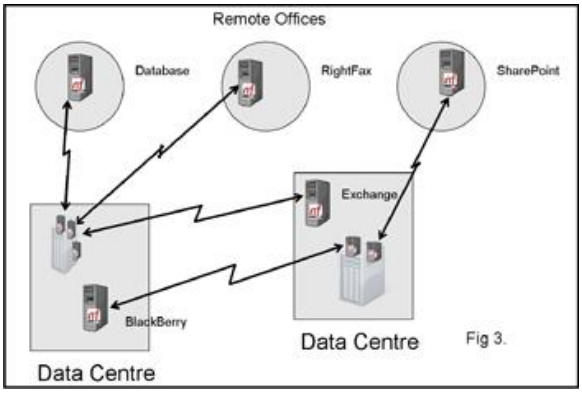

Fig 3.

Neverfail for VirtualCenter

VMware VirtualCenter is an essential tool in the management of complex virtual infrastructures. Everything from VMware VMotion provisioning to VMware Distributed Resource Scheduler (DRS)

optimization depends on VMware VirtualCenter being available. Centralized management, monitoring and higher levels of security also depend on VirtualCenter.

As a single point of control it is essential that VirtualCenter is continuously available and protected against hardware failures, planned maintenance, configuration issues and disaster situations. In fact, keeping VirtualCenter available through any type of outage should be a high priority to ensure maximum return on investment for your virtual infrastructure.

Neverfail for VMware VirtualCenter is the only continuous availability solution designed to automatically protect your VirtualCenter management environment, including protection of the License Server. It clones, then continuously replicates, configuration settings, data and the SQL Server databases to ensure that a consistent, up-to- date copy of VMware VirtualCenter is maintained on a secondary server.

Neverfail then proactively monitors VirtualCenter itself to predict threats to availability. Threats include services becoming unavailable, performance degradation, operating system failures and total machine failures. When an outage is unavoidable, Neverfail automatically switches VirtualCenter control to a secondary machine, there is no interruption to the management and control operations offered by VirtualCenter, and management policies simply continue to run while end users experience no downtime whatsoever.

VMotion Protection

A big advantage of the VMware approach to virtualization comes from VMotion. The ability to provision virtual machines across physical servers to maximize resource utilization, and to do this dynamically, gives great flexibility and opportunity to deliver a greater return on investment.

However, as previously mentioned, provisioning VMotion requires the installation and ongoing availability of the VMware VirtualCenter application. If VirtualCenter is not available, the value of VMotion is lost. Neverfail for VirtualCenter is designed to provide the confidence that VirtualCenter, and components utilizing it such as VMotion will remain continuously available (Figure 4).

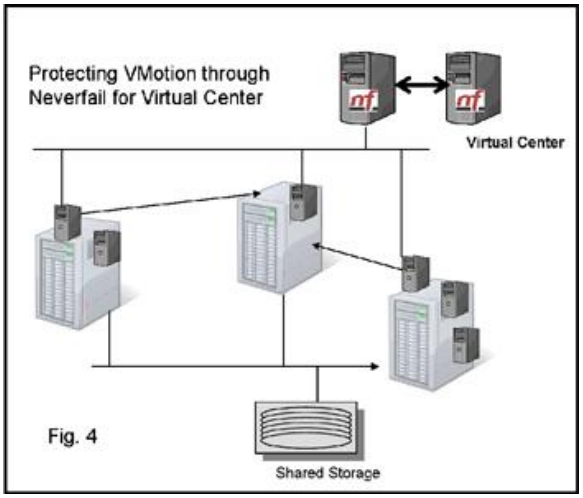

Fig. 4

It is also important to note that while VMotion enables quick reprovisioning of a virtual machine to a different physical host, there is still only one running virtual machine (a single VM file stored on the shared SAN) – by definition a single point of failure. Therefore, failures within that VM won't be solved simply by "VMotioning" to another ESX server, and planned maintenance encounters the same issue.

With implementation of Neverfail's Continuous Availability Suite, full redundancy (even of the VM, since Neverfail for VirtualCenter sets up a secondary VM) is assured. Therefore, failures within the VM environment and/or the need to perform planned maintenance don't result in business downtime as users are kept connected to their mission-critical applications continuously.

Final Thoughts

By combining virtualization technology with Neverfail, one can achieve the best of both worlds – an optimally-provisioned, consolidated server environment that, through a combination of Neverfail's focus on the business application and VMware's focus on the platform, is always available to end users, allowing them to continuously perform their job functions without perceiving any outages. Only through achieving both of these goals simultaneously can the objective of maximized business performance and continuous availability be realized within the VMware environment.

- Maximizes investment in virtual infrastructure with P2V, V2V and V2P Protection
- Secures Application Level High Availability in LAN and Disaster Recovery in WAN environments
- Highly affordable, rapidly installed and easy-to-use product expedites return on investment
- Provides seamless automated or manual Failover *and* Failback, Switchover *and* Switchback
- Keeps VirtualCenter continuously available by protecting VirtualCenter configuration settings and Microsoft SQL Server database
- Ensures continuous VMware VMotion and DRS functionality
- Guards against License Server Failures
- Extensive monitoring of entire ecosystem: application, network, software, hardware (hardware-

About Neverfail

Neverfail is a leading global software company providing affordable data protection, high availability, and disaster recovery solutions focused on keeping users productive. Neverfail's software solutions enable users to remain continuously connected to the live software application irrespective of hardware, software, operating system, or network failures.

Neverfail's mission of eliminating application downtime for the end user delivers the assurance of business continuity, removes the commercial and IT management costs associated with system downtime and enables the more productive use of IT resources. More information can be found at www.neverfailgroup.com

Neverfail's software solutions enable users to remain continuously connected to VMWare® VirtualCenter. Microsoft® Exchange, IBM® Lotus® Domino®, RIM BlackBerry®, Microsoft SharePoint®, IIS, File Server and other Windows®-based applications irrespective of hardware, software, operating system, or network failures.

www.neverfailgroup.com