



100% Storage Uptime: Overcoming vSphere's Single Point of Failure

Chris Gannon is a senior systems administrator with Toys 'n More, a national distributor for children's toys and merchandise. While his company doesn't subscribe to IT's most bleeding edge technology, Toys 'n More's philosophy has long supported incorporating what's new when it makes business sense.

As a result of that philosophy, Toys 'n More was an early virtualization adopter. Their move to virtualization cost them dearly in hardware during the early days, but that investment paid dividends over the years in rapid deployment of services and improved optimization of resources. Thanks to virtualization, they escaped needing to expand their physical datacenter infrastructure with added space, power, and cooling. That cost elimination alone was worth virtualization's early price of entry.

Yet while their virtual practices are noticeably mature, their now-aging storage hardware isn't. Purchased years ago during the early up-swell of interest in SAN technology, Toys 'n More's monolithic storage infrastructure doesn't enjoy many of the features seen in today's hardware.

On this day, Chris and Toys 'n More are about to experience a failure of such magnitude that their entire virtual infrastructure – and indeed their entire computing infrastructure – will cease to function. Creating an unplanned disaster of such scale that no recovery plan was even considered, this failure will almost put the company out of business.

It all starts with a single mouse click.

With that click, Chris' application of a minor but flawed SAN firmware update will inadvertently take down their centralized storage. All of it. As the SAN fails, so too does every single virtual machine. At the click of a button, Toys 'n More loses its entire operations. Business grinds to a halt. Users lose access. Orders cease processing.

At the keyboard, Chris can only blink at the chaos he's wrought.

WHAT MAKES 100%

Virtualization's promise led us to believe that virtual servers were in every way superior to their physical analogs. Virtual machines can be snapshot. They're easier to back up. The infrastructure they lie on comes equipped with high availability features not possible with physical servers.

For nearly all these promises virtualization's products have delivered. Yet they've done so by introducing what could be IT's single greatest point of failure, one that would have never passed the requirements test had virtualization not been such a game changer. *That single greatest point of failure is your SAN storage.*

Think for a minute about the advanced features available today in virtual platforms like VMware vSphere. With the right components in place, this platform includes built-in support for failing over virtual machines when hosts stop functioning. Its DRS feature also adds load balancing functionality to protect hosts from resource overutilization.

Higher availability requirement mandate advanced technology to create a Fault Tolerant virtual machine. With VMware FT enabled, a virtual machine whose host goes down needn't experience downtime at all. Replicating in lockstep with that failed VM is a partner. Located atop some other host, that partner is able to begin servicing clients almost instantaneously after its primary fails. Downtime is measured in units of time so small that clients never see the failover occur.




Figure 1: High Availability Features Still Require Shared Storage

While these technologies are undeniably exciting and great for increasing server availability, they all rely on one insidious requirement: *Shared SAN storage*. As you can see in Figure 1, VMs that failover after a host failure indeed move their processing to alternate hosts. It's their disk files that don't. They continue their operation in the same volume, LUN, and SAN storage throughout the entirety of a vMotion event.

This centralization of disk files makes sense in the typical failover situation. Relocating the contents of a VM's memory state requires the transfer of only a few gigabytes of data. Conversely, that VM's disk file might be measured in tens or hundreds of gigabytes. Even terabyte and sub-terabyte disks are no longer the rare exception. With the sheer size of disk files being this large, transferring that disk to new storage requires an effort bordering on herculean. That effort takes time, and time is exactly what isn't in great supply when a host failure occurs.

As a result, one can't necessarily fault the virtualization vendors for building products that inadvertently elevate storage as single point of failure. Adding features like HA and FT to virtual platform software was necessary for their enterprise viability. They protect against a specific set of failures. Without them, virtualization would never have evolved the industry in the ways it has.

It is critical, however, to recognize that your shared SAN storage is potentially your riskiest single point of failure. As such, proper due diligence mandates that mitigations are put into place to protect against a failure of your entire storage.

In essence, 100% uptime must be your goal.

100% UPTIME ISN'T DISASTER RECOVERY

Most companies have recognized for a long time that this risk exists. To mitigate it, an entire ecosystem of products and services has become available in recent years. Some exist as features within the storage layer itself. Enterprise-grade storage from today's biggest storage vendors now supports redundancy inside nearly every subsystem: From power supplies and storage processors, to network connectivity, disk redundancy, and even entire-node redundancy.

All of these are common in shoring up the storage device

against the range of potential failure states. Yet, by themselves, none truly eliminates the storage subsystem in whole as a single point of failure. Removing the "single" in "single point of failure" requires a horizontal expansion.

Attempting to solve this problem is an entire category of disaster recovery solutions. The central tenet of such solutions centers around the occurrence of some kind of disaster. That incident causes large-scale impact to datacenter equipment. While the problem's cause might not necessarily be the proverbial tornado, hurricane, or other natural disaster, invoking the use of these products tends to focus on one important decision: *You must declare a disaster.*

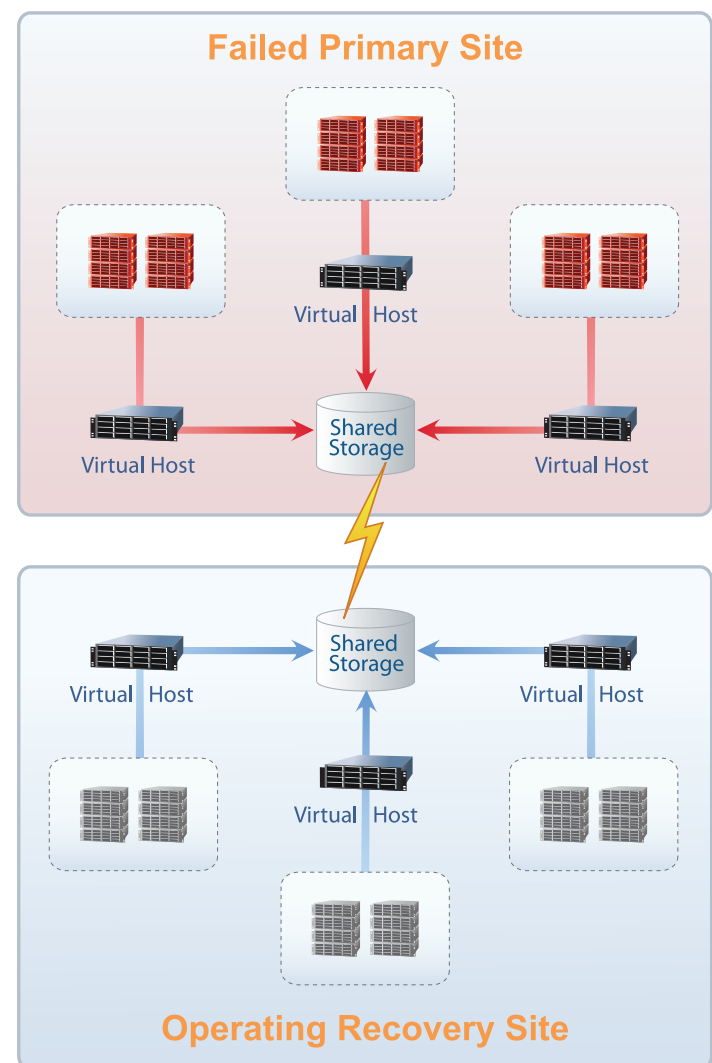


Figure 2: Disaster Recovery Products aren't Necessarily the 100% Uptime Fix

It's this declaration that sets disaster recovery products apart from what solutions might fix the 100% uptime problem. At issue is the scope of the problem these products attempt to resolve. Figure 2 shows an example of that scope, along with the hardware commonly considered part a disaster recovery architecture.

In that picture you can see that disaster recovery products are there to protect against a failure of the primary site itself (or, in some cases, large portions of the primary site). During a primary site failure, its entire scope of equipment is also considered to have failed as well. It's the level of failure that's critically important towards determining whether to declare a disaster or not.

Declaring a disaster with such products can be a business-impacting event as services reposition to the backup site. Positioned there, services can operate with fewer resources or slower network performance. In the interests of time, cost, or complexity, many non-priority services might not have been considered part of the disaster recovery plan's scope.

Further, invoking a disaster recovery plan is one thing. Getting those services back is yet another. Extensive planning can be required to get services back onto their primary site. Often, failback is a more risky event than failover.

CONSTRUCTING THE 100% UPTIME SAN

At the end of the day, constructing a SAN that can assuredly support 100% uptime requires a combination of tactics. Incorporating redundancy features into hardware components is absolutely a start. Expanding that SAN horizontally to create redundancy is yet another. Yet implementing everything else that creates disaster recovery environment above might be going a step too far.

100% uptime desires a highly-available SAN infrastructure that can support the loss of even a SAN itself. This situation might occur even as other hardware in the site remains operational. Accomplishing this goal can be done using software as well as hardware solutions.




Figure 3: Replicated Shared Storage for 100% Uptime

A 100% uptime SAN isn't a complete replication of your entire datacenter, but instead a copy of the SAN itself. Figure 3 shows how this might look from an architectural perspective. In this picture, the shared storage that before represented the single point of failure has been augmented with an additional storage device. Replication in this architecture doesn't occur over WAN links to a backup site. Instead, it occurs over existing high-speed LAN connections.

Implementing a storage infrastructure in this way adds availability over and above what you get with hardware redundancy alone. Virtual hosts are connected to both instances simultaneously, with each instance configured to support bidirectional replication on demand. One SAN may act as primary with the other in a secondary role, or both may be deployed in an active-active configuration.

Building a replicated environment using hardware SANs can get expensive. However, the constituent components that comprise what we think of as "a SAN" could also be constructed

out of lower-cost commodity components. “A SAN” is a set of storage processors, not unlike a server’s processors. It comes equipped with disks and network interfaces. It also runs any of a range of specialty operating systems that connect incoming requests to appropriate data.

For many hardware SANs in operation today, that specialty operating system is an edition of Microsoft Windows called Windows Storage Server. With Windows Server as its core, Windows Storage Server adds a range of additional functionality for serving data to connected servers.

Similar architectures are also possible using Windows Server with a class of software commonly referred to as *software SANs*. These software SANs expose a level of features very close to those found in traditional hardware SANs. *However, because software SANs can run atop existing commodity hardware, their cost basis can be much lower than traditional hardware SANs.* Environments that use them also tend not to suffer the experience gap with IT professionals unfamiliar with hardware SAN technologies.

One can’t talk about a software SAN atop Microsoft Windows without pointing out the Windows OS’ regular need for updates and reboots. Reboots in a SAN environment go counter to the requirement for 100% uptime. Thus, the software SAN you select must support high availability features like replicated shared storage. The right solution is also trivial to manage, and capable of repairing itself as a failed server comes back online.

SOFTWARE SANs ENABLE COST-EFFECTIVE 100%

SAN hardware and ongoing maintenance represents a primary cost for many datacenters. While the cost of each gigabyte of online storage continues to drop, that reduction drives business towards a greater need for availability. Getting to 100% uptime, whether your environment is virtualized or not, is a worthy goal. With businesses today placing ever-greater demands on IT services, ensuring their consistent availability has become one of IT’s most important tasks.

The components with which to construct that 100% uptime environment are many in form. Hardware can accomplish the task, with its associated price point. Alternatives to hardware SAN solutions such as software SANs just might give you the cost-effective path you need to elevating your storage infrastructure to 100%.

Don’t put yourself in Chris’ shoes, blinking at the storage chaos you’ve created through a single mouse click. Get your storage uptime to 100%, or be prepared for the consequences.



Turn Any Server Into a SAN.
Enterprise Features.
SMB Price.™

www.starwindsoftware.com

ABOUT STARWIND

Since 2003, StarWind has been the storage solution of choice for thousands of global customers in over 50 countries, from SMBs, to governments, and to Fortune 1000 clients. StarWind has pioneered the iSCSI / IP SAN industry with its storage virtualization software that converts any Windows Server into a reliable and scalable shared storage.

©2011, StarWind Software Inc. All rights reserved. Reproduction in any manner whatsoever without the express written permission of StarWind Software, Inc. is strictly forbidden. For more information, contact StarWind. Information in this document is subject to change without notice. StarWind Enterprise Server is a registered trademark of StarWind Software.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.