SOLUTION BRIEF:
TRANSFORM YOUR WINDOWS® DATA CENTER:
SOLVE KEY STORAGE AND HIGH AVAILABILITY

Transform Your Windows® Data Center: Solve Key Storage and High Availability Challenges

Who should read this paper

Windows® IT architect and IT director for Windows® Server

Content

Introduction

Today's data centers have become highly evolved, with virtualization being widely deployed and many organizations starting to move toward the cloud. These technologies offer tremendous benefits such as increased flexibility, reduced capital and operational expenditures, and reduced hardware footprints in the data center. However, as with any change, challenges, both anticipated and unanticipated, can arise. These challenges, if not met head on, have the potential to slow down and, in some cases, derail progress. To survive and remain competitive in today's evolving technology landscape, businesses have to adapt, not only by adopting these new technologies, but also by successfully meeting the challenges that come with them.

This paper will identify some of those challenges and discuss how Symantec can help address them to allow administrators to provide more dynamic and resilient IT environments.

Challenges in evolving Windows® data centers

Among the many challenges that arise with the introduction of virtualized workloads and the accelerating rate of cloud adoption are increased demand for storage capacity, maintaining highly available for applications in those environments and providing true disaster recovery. Administrators need tools that help them to better meet these challenges if they are to keep their virtual and cloud environments highly available and scalable, while keeping performance at a premium and operating costs low.

With organizations picking up new technologies, the amount of data generated by applications is growing exponentially. Not only does this data need to be retained for longer periods due to regulatory and other requirements, but live migration of data associated with critical applications places additional demands on storage and creates problems for administrators with respect to both increasing storage cost and management of complex storage infrastructure. According to Gartner, data will grow 800 percent in the next 5 years.[1] Some companies stand up more and more storage to accommodate this massive data explosion. Sadly, organizations rarely know what storage is actually used or available before they purchase more capacity.

Virtualization and cloud adoption are on the rise, with businesses and government agencies starting to deploy their own private clouds and/or moving towards leveraging some aspects of public cloud. Forrester estimates that the amount spent on the global market for cloud will grow from $40.7 billion in 2011 to more than $241 billion in 2020.[2] While VMware® remains the dominant server virtualization platform, other players are making strides of their own. Microsoft Hyper-V® is considered by many to be “good enough” and a cost-effective alternative to VMware in some areas and is seeing increasing levels of adoption. Other platforms such as Kernel-based Virtual Machines (KVMs), Sun® Logical Domains (LDOMs), and Logical Partitions (LPARs) are also on the horizon.

Application requirements don’t change when they move to a virtual environment, but providing high availability can become increasingly challenging as an ever growing number of applications are hosted within virtual machines. While clustering applications addresses some of these challenges, protecting the application within the virtual machine presents its own unique set of challenges as virtual solutions typically do not provide visibility to the applications. Leaving you, as an administrator, unsure how well an application is running. Additionally, virtual solutions very or lack the ability to stop or restart the application.

When unpredictable disaster strikes, companies want and need to be best prepared. Disaster recovery can prove prohibitive due to its complexity and cost. Without an affordable, easy-to-manage, testable, bullet-proof disaster recovery solution, some companies decide to forgo a disaster recovery solution. For non-heterogeneous data centers customers will not have singular visibility to remedy disaster recovery challenges before they impact the bottom line. Companies need comprehensive disaster recovery solutions.

1-Gartner’s webcast “Technology trends you can’t afford to ignore.” April 11, 2012, by Raymond Paquet
2-Forrester Research, Inc “Sizing the Cloud,” April 21, 2011, by Stefan Ried and Holger Kisker

In addition to the new challenges that come with evolving data centers, more traditional challenges in physical environments continue to preset themselves and have to be addressed. An example of this is customers who are deploying larger configurations, which have the potential to negatively impact the time taken to complete a cluster failover. In environments with large storage configurations, some customers reduce their logical unit number (LUN) and volume count to provide faster failover at the expense of granularity or they might deploy Veritas™ Clustered File Systems from Symantec and Veritas™ Clustered Volume Manager from Symantec to provide multinode access and speedy failover.

Transform your Windows data center: Solve key storage and high availability challenges

Customers need solutions that are flexible, built for virtual and physical environments and can be deployed in small and large environments. Veritas Storage Foundation™ from Symantec and High Availability solutions for Windows is a suite of products that provide advanced volume management for maximum flexibility, optimal storage utilization and performance, and state-of-the-art clustering for the highest levels of application availability and disaster recovery. It is comprised of four distinct products—Veritas Storage Foundation™ for Windows® from Symantec, Veritas Storage Foundation™ HA for Windows® from Symantec, Veritas™ Cluster Server for Windows from Symantec, and Veritas™ Dynamic Multi-Pathing for Windows from Symantec—and a rich feature set that provides a comprehensive solution set to make the data center more resilient and cost-effective.

Minimize storage spend

Without a healthy storage infrastructure, data access can be impeded or blocked, or data can be lost. Deploying the correct storage is essential to maintaining that health in the most cost-effective manner. Having the ability and flexibility to select appropriate storage for specific needs frees a business from the confines of vendor lock, allowing freedom of choice, so that it can more effectively negotiate the best and most cost effective storage solutions that meet its needs. However, having different types of storage from different vendors can present significant management overhead, as each vendor has its own set of management tools, with some vendors even requiring different tools for their various array families.

To overcome this challenge, businesses need solutions that allow them to easily manage these mixed environments. Storage Foundation for Windows is a host-based, advanced volume manager that allows storage that's seen by a host to be virtualized, creating a single pool an administrator can create virtual spaces for different purposes. These virtual spaces are comprised of disk groups, which are created with a subset of disks or LUNs available to the host, and volumes. These volumes are created on those disks. As management is done at the host-level, operations can be performed across disks selected from the same array or from different arrays. Practical applications of this include mirroring across arrays to provide array fault tolerance or allow a cluster to be stretched across data centers to provide site-level fault tolerance, creating snapshots to less expensive tiers of storage for cost savings, and replicating data to remote locations without having to deploy the same type of storage at those locations.

In virtual and cloud environments, storage resources are typically integrated into a single pool that serves the needs of multiple business units and applications, and is required to handle several different workloads. The storage has to be highly and easily scalable, and should be able to take advantage of thin provisioning for optimal utilization. It should also provide a means to optimize performance by allowing easy movement across storage tiers so that the right storage can be dedicated to appropriate workloads.

Identify waste and optimize storage usage

To maximize utilization of their storage, many data centers are moving away from traditionally thick storage, where all space is allocated in the array as soon as it's provisioned, and are adopting thin provisioned storage where space is virtually provisioned up front, but dynamically allocated in the hardware as data is written to it.

For all of the benefits of thin provisioned storage, data centers face a new challenge with the potential to run out of disk space due to over-subscription. Storage administrators have to keep an eye on actual storage usage and proactively add additional capacity as necessary. It would be beneficial if they could not only have insight into actual storage utilization, but real insight into actual needs; not all allocated storage is actually utilized. Users also routinely delete files, freeing up space on the allocated storage. Reclaiming that free space would delay the need to purchase additional storage. Thin reclamation in Storage Foundation for Windows provides a mechanism to easily reclaim freed up space on thinly provisioned storage. Storage Foundation for Windows interfaces with the file system to determine available free space and communicates that to the array so that storage can be reclaimed. Reclamation can be immediate or scheduled and can be done at the volume level, the disk group level or both.

Space can also be reclaimed by shrinking volumes while they remain online. The volume shrink feature in Storage Foundation for Windows allows for dynamic volumes that are either formatted with New Technology File System (NTFS) or Raw Volume (RAW) to be shrunk by up to 50 percent so that the space can be repurposed, allowing new volumes to be created or existing volumes to be relocated there.

Storage Foundation for Windows also allows administrators increased flexibility through its ability to monitor volume capacity and take various actions when specified thresholds are reached. Administrators can enable automatic volume growth for those volumes so that, in the event that a critical threshold for the amount of used space is reached, the volumes grow automatically by pre-set amounts so the applications that depend on them are not in danger of going offline due to a lack of space. This functionality is very important in cloud environments where automation plays a major role.

Increase visibility into entire storage infrastructure

For organizations who want even more visibility from the host to storage, Storage Foundation for Windows has a new feature called, Extended Attributes. Veritas™ Operations Manager further extends that visibility across the entire IT enterprise.

Administrators generally have to trust that they've been provisioning the correct type of storage for their application, as they have no visibility into the underlying storage. This may waste valuable time troubleshooting poor performance of a database, for example, only to discover that the root of the problem lies with the underlying storage. Having visibility into their storage devices could save several troubleshooting steps and valuable hours that could otherwise be spent tracking down the problem.

With Storage Foundation for Windows, information is discovered and displayed for arrays and LUNs, giving administrators the visibility they need to maximize efficiency and reduce potential downtime. The following is a full list of attributes:

- Vendor ID
- Product ID
- Revision ID
- Array replicaiton LUN
- Array media type
- Array transport protocol

- Cabinet serial number
- Array volume ID
- Array LUN type
- Array Redundant Array of Inexpensive Disk (RAID) level
- Array snapshot LUN
- Array port World Wide Name (WWN)
- Array port serial number
- Array controller ID
- Array hardware mirror

Ease of storage migration in virtual and physical environments

Storage Foundation for Windows adds intelligence to storage migration operations and provides a level of automation and orchestration so that multiple volumes can be easily moved to new storage locations while online. The storage migration wizard allows you to set up one time migration operations. You can schedule them from the disk group perspective or in Hyper-V environments, from the perspective of a virtual machine. This allows administrators to easily accomplish the task of ensuring that their virtual machines and application data are located on appropriate tiers of storage for performance and/or cost optimization. In virtual and physical environments, coupled with Storage Foundation for Windows, the SmartMove feature allows administrators to easily migrate from capacity-inefficient, thick storage to capacity-efficient, thin provisioned storage without incurring any downtime. Volumes migrated to thin storage will cause just enough allocation in the array to host the data they contain. As the data set grows, additional allocation will occur as dictated by the thin array.

Keep storage highly available for critical business application

From solely a storage focus to include server-level benefits that provide protection against server and application failure to reduce both planned an unplanned downtime. Storage Foundation HA for Windows scales from simple two-node clusters to metropolitan-sized clusters and full wide-area disaster recovery solutions that protect applications at the site level.

To keep continuous access and increase bandwidth to storage and the data that lives there, data centers deploy multiple paths between servers and the storage they connect to. Once multiple paths have been established, a robust and flexible multi-pathing solution is required. At its most basic, it should protect data from being multiple instances at the host so that corruption does not occur. Also, it should fail over I/O to an alternate path in the event of a failure of the active path, and provide some level of intelligence so that all available paths are optimally utilized for data transfer. Most array vendors provide one or more solutions that are specific to their hardware. In multivendor environments, organizations can find themselves having to deploy and manage multiple solutions. In single vendor environments, organizations are faced with higher than necessary costs due to vendor lock-in and may still have to deploy multiple solutions for different array families from the same vendor.

Dynamic Multi-Pathing provides a consistent multi-pathing solution across a huge list of storage array families across multiple vendors enabling customers to reduce both capital and operational expenses. Unlike vendor provided solutions, which are very proprietary with support for a limited set of that vendor's arrays, and the generic native solution that ignores array differences, Dynamic Multi-Pathing exploits array specific capabilities through a library of device specific modules (DSMs), which provide array specific support for a wide variety of the leading array families, with several load-balancing policies that allow for performance optimization. In addition to providing array specific support, Dynamic Multi-Pathing also provides Array Vendor ID (AVID) naming. Devices are named based on an attribute in the array,

making these names consistent across servers that have access to these devices and making it easier for server and storage administrators to communicate. Dynamic Multi-Pathing provides I/O statistics at the disk level, for all disks in an array or for disks across all arrays connected to a host. It also provides aggregate statistics, which allow statistics to be grouped by device to show statistics per device, or by array to show statistics at the array level. Aggregating statistics are useful to storage administrators when trying to determine where and why performance bottlenecks exist. Dynamic Multi-Pathing for Windows fully integrates with the Microsoft® Multipath I/O (MPIO) architecture for a fully Windows logo-certified solution.

Dynamic Multi-Pathing is available both as a feature in Storage Foundation for Windows and as a standalone product. Users, who want a heterogeneous multi-pathing solution without the advanced volume management provided by Storage Foundation, can install the standalone version of Dynamic Multi-Pathing and use the operating system's native disk management to manage their storage space.

Recover from disasters and reduce application impact with full-mirror snapshots

With the FlashSnap feature in Storage Foundation for Windows, you can create independently addressable, full-mirror snapshots of data volumes. It integrates with Microsoft Volume Shadow Copy Services for application consistent snapshots of Microsoft® Exchange, SQL®, and SharePoint® servers, and Symantec Enterprise Vault™. These full-mirror snapshots can be used on-host as backups to quickly recover from disaster, or moved off-host for operations such as backup, test and development, and reporting. It also integrates with the native Volume Shadow Copy-on-Write provider in Windows to provide space saving snapshots, which can be used on-host to go back to previous versions of a volume, or to recover from data corruption or other disaster.

The Fast File Resync feature is also available to allow quick recovery of individual files, such as those used by a database or a virtual machine from snapshots of volumes that host multiple instances of those files. This is important as data centers trend toward large LUNs hosting multiple files rather than dedicating LUNs to individual databases or virtual machines.

Prevent data lost using data replication across any distance

Veritas™ Replicator from Symantec replicates data between geographically dispersed sites over an IP network. The sites can be any distance apart. An additional product to Storage Foundation HA for Windows, Veritas Replicator supports three modes of replication: synchronous, asynchronous, and synchronous override. Synchronous replication keeps the target (secondary) up-to-date and is supported within distances that conform to latency limitations. Synchronous replication impacts application performance if the secondary's acknowledgement is slow or it fails to respond. Asynchronous replication has no distance limitations, does not impact the application, but can be behind the source (primary), which can lead to data loss if the source fails.

Synchronous override leverages the other two modes so that the secondary is up to date as long as it s available, but the application is not impacted if the secondary becomes disconnected. With synchronous override, replication is synchronous as long as the secondary is connected, but changes to asynchronous if it becomes disconnected.

The replication process maintains write-order fidelity for the volumes in a Replicated Volume Group (RVG) via a circular log that is referred to as the Replicator Log or SRL. Consistency of the data in the Replicated Volume Group is guaranteed by the replication process, which sends the data from the SRL at the primary site to the secondary site in the same order that it is written to the volumes. Under normal operation conditions, applications, such as databases and their logs, will always be consistent at the secondary i.e. the database can be mounted against its logs even if the secondary site is behind the primary site.

Bunker replication is also available to ensure a zero Recovery Point Objective (RPO) when asynchronous replication is used. With bunker replication, a secondary (bunker) site is setup close enough to the production site to allow its replication logs to be replicated synchronously to it. In the event of a failure at the production site, recovery at the secondary site is accomplished by Veritas Replicator copying any outstanding writes from the bunker site to the secondary site and then putting the secondary site into production.

Veritas Replicator also offers data compression for optimal bandwidth utilization, memory tuning to monitor and control the amount of Non-Paged Pool (NPP) memory it uses so that it does not contribute to servers freezing due to memory running out, and graphs that show bandwidth and memory usage for better visibility into how resources are being consumed. Bandwidth usage can be displayed in real time or historically.

With Veritas Replicator, administrators have an easy route to the cloud via a mechanism that allows them to easily migrate application data to the cloud or to setup what could be the basis for a disaster recovery solution to or from the cloud. It can also be used to migrate data centers by replicating the data and then migrating operations to the new data center. At the end of the migration, data is in a form that's ready to take advantage of the other benefits that Storage Foundation for Windows offers, and as a host-level solution, Veritas Replicator gives control of what will be replicated and when it will be replicated to the system administrator.

Extend fault tolerance to your site level using campus clusters

Single clusters can be stretched beyond the confines of a data center to extend fault tolerance to the site level. Host-based, volume mirroring across arrays located at different data centers provide campus clusters for site-level fault tolerance. If inter-data center connections aren't available to support mirroring, data can be synchronously replicated across an IP network, allowing replicated data clusters to be setup. Veritas™ Cluster Server for Windows, included in Storage Foundation HA for Windows, monitors the replication and reverses its direction when failover occurs to the other site. These stretched clusters are limited to metropolitan distances so that latency requirements for mirroring and synchronous replication are not violated.

Without Storage Foundation for Windows, core functionality growing or mirroring volumes are quite challenging to achieve in a Microsoft cluster. Storage Foundation for Windows provides the Cluster Option for Microsoft Failover Clusters. With Storage Foundation for Windows, Microsoft clusters can be stretched to create campus clusters by mirroring volumes between arrays located at different data centers, volumes can be dynamically resized, migrated to new storage or replicated, and snapshots can be created for on-host or off-host operations, making those environments more highly available and resilient.

Clusters can have multiple nodes at each site. To better control where failover occurs i.e. locally instead of remotely, Veritas Cluster Server supports the concept of a system zone. A system zone is a subset of systems that are the initial targets for failover. Systems outside of the zone or in another zone will be failover targets if none of the systems in the active zone are available for failover. This allows control of failover so that it occurs at a particular site before moving to another site. Upon site failure, failover then occurs to the other site.

Achieve wide-area disaster recovery with global clusters

The need for resilient disaster recover typically requires a site that's located at a distance beyond those that support synchronous replication and can exist on different network segments. Veritas Cluster Server Global Cluster Option allows separate clusters to be setup and uses a wide-area connector to establish communication between the clusters. As with replicated data clusters, data is replicated between the clusters with either Veritas Replicator or hardware replication. In the event of a disaster, a remote switch can be performed and Veritas Cluster Server handles all updates, such as Domain Name System (DNS) and replication reversal, and brings the application(s) online.

Setting up wide-area disaster recovery clusters can be a complex task, involving several steps that include configuring storage, setting up the cluster, installing applications, clustering the applications, setting up replication, and configuring disaster recovery. Configuring disaster recovery with Veritas Cluster Server is simplified with the Solutions Configuration Center, which provides workflows and configuration wizards.

Test your disaster recovery solution without disruption

Establishing a disaster recovery solution that works for your organization will provide the highest levels of protection against disasters, whether man made or natural. However what good is a disaster recovery site if it fails to recover your applications when an actual disaster occurs. A standard practice for any disaster recovery solutions should be to regularly test the disaster recovery site to ensure that it functions as expected. In keeping with reducing downtime for the highest availability, testing procedures should be completely non-disruptive.

Veritas Cluster Server also includes the Fire Drill feature, which allows you to proactively and non-disruptively test your disaster recovery sites. Snapshots or clones of data at the disaster recovery site are created and mounted in a clone of application service group to allow testing by the application. Successful testing gives peace of mind that, in the event of a disaster, the application can be successfully recovered at the disaster recovery site. If testing fails, it allows for troubleshooting and correcting the problem before disaster strikes.

Fire Drill can be set up with the Veritas Cluster Server Solution Configuration Center. Currently supported replication solutions include Veritas Replicator, EMC® SRDF, and Hitachi True Copy.

Reduce failover time to seconds

As seen earlier in this paper, the Storage Foundation for Windows multi-access disk group architecture removes the need to import an entire disk group when ownership changes to another host. As storage configurations grow in size this can be beneficial to cluster failover times when that storage is a resource in the cluster.

Veritas Cluster Server adds the Fast Failover attribute to its Volume Manager Diskgroup (VMDg) agent so that clusters with large storage configurations can take advantage of this new disk group architecture. Failover is fast, regardless of storage configuration size, as disk groups are already in an imported state on the failover target. On failover, read access is changed from read-write to read-only on the node being failed over from, and from read-only to read-write on the node being failed over to. As a result, failover time can be reduced from tens of minutes to seconds. Figure 1 shows failover times for a SQL Server 2008 R2 cluster configuration with 150 LUNs, 25 disk groups, and 125 volumes. With this configuration, traditional failover took an average time of 27 minutes. With fast failover enabled, average failover time dropped to 55 seconds.[3]

3-Symantec Performance Engineering Group

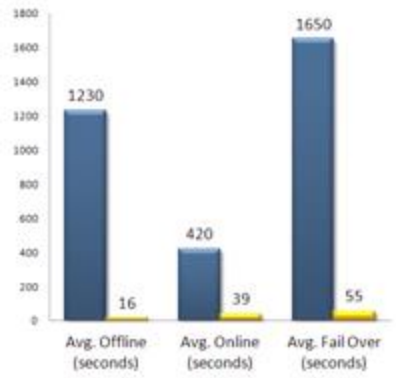

Figure 1. Average failover time drops from 27 minutes to 55 seconds with Fast Failover

Storage Foundation for Windows introduced a new architecture, which allows cluster dynamic disk groups containing shared storage to be simultaneously imported on more than one host. In previous versions, to protect against potential data corruption when storage is shared by more than one host, disk groups could be safely imported on only one host at a time. Traditional cluster and private disk groups place SCSI reservations on the disks in the disk group so that only the node with a majority of the disks in the disk group could import and own it.

With the new multi-access architecture, disk groups can be imported on all nodes that have access to them. Protection against data corruption is provided by controlling write access; owning nodes have read-write access to the disk group, while non-owning nodes have read-only access (Figure 2). When ownership changes, write access is changed so that the new owner has read-write access and the previous owner's access changes to read-only. As the disk group does not have to go through the process of being deported and imported during a failover, the time taken to change ownership is reduced. This is a huge benefit for clusters that use shared storage, as failover time is no longer dependent on the size of the storage configuration. With large storage configurations, the process of deporting a disk group and importing it on another node could take several minutes. With multi-access disk groups, the entire process can occur in a minute or less.

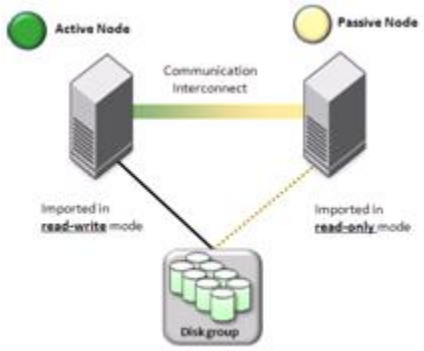

Figure 2. Disk group imported on multiple nodes

The new disk group architecture also allows for live migration of virtual machines that are located on dynamic disks in Hyper-V environments. These virtualized environments can fully benefit from the advantages of Storage Foundation for Windows both in guest and at the host level. Volumes can be mirrored for redundancy and grown dynamically, snapshots can be created for quick recovery from corruption and virtual machines or their data can be migrated to new storage locations, new data centers or to the cloud, all while remaining online and accessible.

Monitor framework intelligently

Cluster monitoring is traditionally done via a polling mechanism, in which the cluster intermittently polls resources to check their state. If faults are detected, corrective action is initiated based on policy. However, depending on the monitor interval, time to detect a state change can be extensive, resulting in delayed reaction. This can be shortened with a lower monitor interval, but this takes additional system resources, resulting in increased memory and CPU utilization. Administrators are forced to choose between cluster and system performance.

Veritas Cluster Server for Windows solves this with the Intelligent Monitoring Framework (IMF) feature, which replaces traditional polling with an asynchronous, event driven method of monitoring resources. When a cluster comes online, its resources that are configured for Intelligent Monitoring Framework register themselves with the agent framework. If a resource's state changes, it's immediately detected by the Veritas Cluster Server agent and communicated to the Veritas Cluster Server engine for corrective action. With IMF, the cluster helps detect failures instantly and, due to the fact that it is not constantly polling resources, is more efficient in its use of system resources. CPU utilization can be seen to peak at over 88 percent and averages over 22 percent without the Intelligent Monitoring Framework.[4] With the Intelligent Monitoring Framework enabled, peak and average utilization falls to zero. It will rise momentarily if a fault is detected, and then immediately returns to zero.

Ensure continued high availability with campus clusters

Mirroring provides fault tolerance so that if a disk or LUN fails or is removed there are surviving plexes of a volume from which data can continue to be accessed. This works fine when mirroring to other LUNs in an array as long as the array itself is available, but is defeated if there is a failure at the array level. Mirroring across arrays solves this by providing fault tolerance at the array level. It also provides an opportunity for clusters to scale beyond a single data center.

As a host-based volume manager, Storage Foundation HA for Windows allows data volumes to be mirrored across storage arrays, whether the arrays are from a single vendor or from different vendors. If there are interconnects between data centers that allow access to the arrays from servers at each location, campus clusters can be setup by mirroring across the arrays at each site so that each cluster node has access to volume plexes at each site. Storage Foundation HA for Windows is architected so that, if the non-active data center fails, the disk group and volumes remain online at the surviving and active data center. If the active data center fails, the surviving data center can be brought online as it has mirrored copies of all volumes in the cluster.

In these configurations, consideration also has to be given to changing landscapes; volumes may have to be grown and new storage may have to be added to accommodate expanding data sets. It's important that, as configurations grow, they do so in a way that takes into account the requirement that they be evenly dispersed between data centers. Storage Foundation for Windows provides a mechanism to ensure this with site-aware allocation, which allows administrators to tag storage resources with a site name so that all growth takes place at the appropriate location to keep the configuration consistent. If the volume is resized, relocated or its layout is changed, site boundaries are maintained so that a complete plex of the volume is always located at each site. This makes it easier to automate reconfiguration without having to be concerned that the site to site configuration will be violated.

Tie multi-tiered applications together

With the continued growth in virtualization, we're seeing more mixed platform environments. A vast majority of Tier-3 applications are virtualized, but most Tier-2 and Tier-1 applications are still running on physical servers. Operating systems are also varied, with multi-tiered applications typically being deployed on different operating system (OS) platforms across the different tiers. This increases the challenges

4-Symantec Performance Engineering Group

that business face managing their business services in these complex environments. Challenges such as starting and stopping a service, providing high availability and disaster recovery, security, and determining status of the business service across the various tiers and platforms need to be addressed.

As a feature of Veritas Cluster Server for Windows and Veritas Operations Manager, Virtual Business Service provides a solution to these challenges by allowing the various tiers to be tied together as a single business entity that can be controlled as a single application, while allowing each layer to make independent operational decisions. With Virtual Business Service, clusters running at each tier can be configured with their own policies for handling faults, whether occurring at that tier or at another tier. Virtual Business Service propagates faults across the business service so that, depending on how they're configured, the other tiers can decide to go offline, remain online or restart in response to a fault at another tier. If the business entity has to be brought online or taken offline, Virtual Business Service will coordinate the order of operations at the various tiers to honor dependencies.

Virtual Business Service also supports disaster recovery so that, if the business entity has to failover to a disaster recovery site, the tiers are brought down in the correct order at the faulted site and brought up in the correct order at the disaster recovery site. As businesses make the move toward private clouds, Virtual Business Service allows them the flexibility to utilize their existing infrastructure to get there.

Manage disaster recovery in Hyper-V environments

In Hyper-V environments, Storage Foundation HA for Windows coordinates with Microsoft Failover Cluster to recover from all failure scenarios. Failover Cluster recovers virtual machines in response to local failures. Storage Foundation HA for Windows recovers virtual machines in response to a disaster. Veritas Cluster Server includes the Disaster Recovery Manager feature that provides disaster recovery automation in Hyper-V environments.

Disaster Recovery Manager can be easily deployed into existing Hyper-V environments to extend their capabilities beyond a local data center to remote data centers hundreds of miles away for true disaster recovery. Veritas Cluster Server for Windows is installed in a separate virtual running in a Microsoft Failover Cluster at each site. The Hyper-V DR Manager Virtual Machine Configuration Wizard is used to configure the connection between the sites and virtual machines are easily configured for disaster recovery via PowerShell. Hardware replication is setup between the production and disaster recovery sites with either EMC SRDF or Hitachi True Copy. For high availability, multiple disaster recovery Manager virtual machines can be configured as resources in the Failover Cluster.

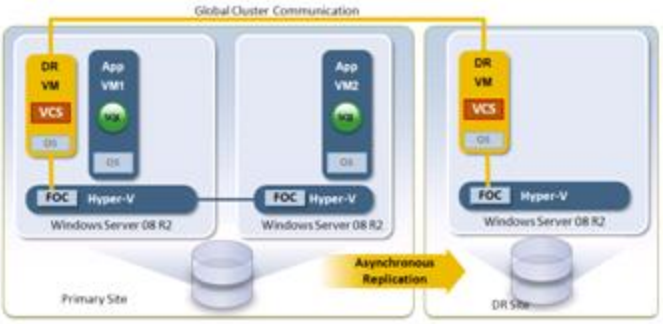

Figure 3. Veritas Cluster Server for Windows Disaster Recovery Manager installed in a Hyper-V virtual machine

On initial setup, virtual machine configurations are exported to shared storage at the production site and imported at the disaster recovery site on initial failover. Subsequent failovers do not require additional imports of the virtual machine configurations as long as the configuration remains unchanged. If new virtual machines are added, their configurations have to be exported and will be imported at the disaster recovery site on the next failover.

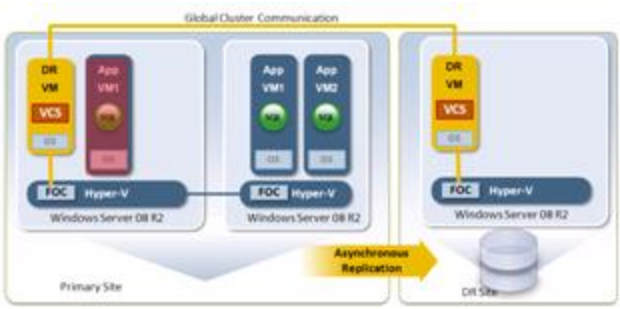

Figure 4. Microsoft Failover Cluster performs local failover

Hyper-V disaster recovery works at the site level; all virtual machines configured for disaster recovery are moved to the other site during a failover. Microsoft's Failover Cluster handles all local failover (Figure 4), while the Disaster Recovery Manager handles site failover (Figure 5). If a virtual machine fails, Failover Cluster will attempt to start it on another Hyper-V host in the cluster. If failover is required to the disaster recovery site, Veritas Cluster Server for Windows (aposturefy) disaster recovery for Hyper-V.

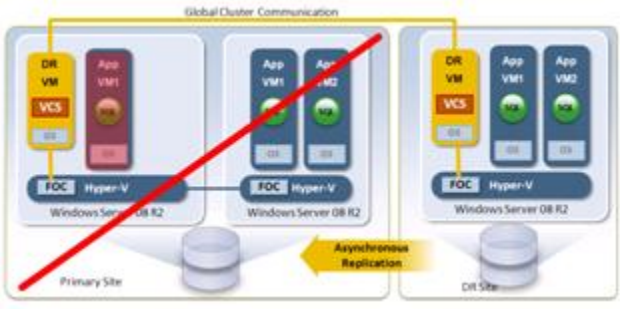

Figure 5. Disaster Recovery Manager performs site failover

Provide advanced clustering that scales

Veritas Cluster Server for Windows provides cost-effective high availability with configurations that support single roaming spare servers. Where a single server is dedicated as a failover target for other servers in the cluster, and has no spare configurations. Additionally, all servers are active and can act as failover targets for other servers in the cluster, hosting multiple application instances when necessary. Service group workload management provides intelligence that allows Veritas Cluster Server to determine where an application fails over to based on available system resources and application requirements, making it possible to more cost effectively deploy much larger clusters.

Veritas Cluster Server supports various replication technologies that include Veritas Replicator and several hardware replication solutions, including EMC SRDF, Hitachi True Copy, and HP® Continuous Access. Replication extends Veritas Cluster Server capabilities to distances beyond a single data center, either with single stretched or replicated data clusters within limited distances, or across clusters, subnets, and unlimited distances for wide-area disaster recovery. This makes Veritas Cluster Server highly scalable solution that supports from one to 32-node clusters, in a single data center or across data centers.

Veritas Cluster Server also adds a framework for intelligently monitoring resources so that faults are instantaneously detected and cluster reaction time minimized for faster failover. It also supports Storage Foundation for Windows's multi-access disk groups for fast failover, independent of storage configuration size.

How customers are currently using these solutions

Customers see value in using Storage Foundation HA for Windows to implement and maintain their high availability, disaster recovery, and storage management environments. Below are two customers who have successfully solved their business challenges with advanced solutions from Storage Foundation HA for Windows:

- SwapDrive, which explains why they've never had an outage with Storage Foundation HA/DR for Windows' flexible architecture
- MTR Corporation Ltd who saved 33 percent of their hardware costs by using Storage Foundation HA for Windows products

Conclusion

Data centers are fast adopting new workloads like virtualization and cloud. This transition is stressing existing storage infrastructure management. Additionally, with business critical applications running in virtual and cloud environment, high availability, storage management, and automation of disaster recovery become important.

Storage Foundation HA for Windows enables wide-area disaster recovery across multiple sites over any distance. It supports the widest array of data replication technologies, including multi-site disaster recovery configurations with Veritas Replicator. Coupled with Veritas Operations Manager, it provides cloud-based solutions by answering several of the challenges, including deep visibility, comprehensive reporting, high availability, disaster recovery, and centralized control and management.

Storage Foundation HA for Windows also provides advanced volume management for maximum flexibility, optimal storage utilization, and performance, IT administrators have a powerful set of tools from Symantec for managing across a broad landscape that spans physical, virtual, and cloud environments.

About Symantec

Symantec is a global leader in providing security, storage, and systems management solutions to help consumers and organizations secure and manage their information-driven world. Our software and services protect against more risks at more points, more completely and efficiently, enabling confidence wherever information is used or stored. Headquartered in Mountain View, Calif., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our website.

Symantec World Headquarters
350 Ellis St.
Mountain View, CA 94043 USA
+1 (650) 527 8000
1 (800) 721 3934
www.symantec.com

Symantec helps organizations secure and manage their information-driven world with **endpoint virtualization**, **server virtualization**, and **application virtualization**.

7/2012 21257818