CITRIX

Microsoft

Microsoft | Virtualization

Providing High Availability and Disaster Recovery in a Multi-Site Virtual Environment

End-to-End Solution Enabled by Microsoft Virtualization, HP P4000 iSCSI SAN Solution, and Citrix Essentials for Microsoft Hyper-V Software

Solution Blueprint

Published: March 2010

Abstract: This white paper describes an end-to-end high availability and disaster recovery solution for Windows® virtualized environments. This solution is enabled by Microsoft® Windows Server® 2008 R2 Failover Clustering, Hyper-V™, the HP P4000 iSCSI SAN solution, and Citrix StorageLink™ Site Recovery software and is prepackaged in the HP Disaster Recovery Smart Bundle for Hyper-V. This solution demonstrates automated failover capabilities in a multi-site virtualized Microsoft environment.

Copyright

Contents

Executive Summary

Although organizations would like high availability (HA) and disaster recovery (DR) implemented for all applications, these technologies have traditionally been too expensive and complex. Therefore, companies have only implemented HA and DR for applications when required. This paper describes a solution that breaks through those barriers using virtualization. The solution provides companies with the ability to easily and inexpensively construct an HA environment with DR for all applications including those that are not business critical.

This white paper defines a blueprint to build a robust HA and DR solution for any organization. The blueprint provides a specific approach for structuring an HA and DR solution in a multi-site Microsoft® Windows® Server® 2008 R2 Hyper-V™ environment with the help of failover clustering. The blueprint uses Microsoft's virtualization and failover capabilities complemented with HP's P4000 iSCSI Storage Area Network (SAN) solution and the Citrix® StorageLink™ Site Recovery software. For simplicity, the solution is prepackaged for sale in the HP Disaster Recovery Smart Bundle for Hyper-V and includes all required hardware, Windows Server 2008 R2 operating system with Hyper-V, and provides Citrix StorageLink and Citrix StorageLink Site Recovery software (components of Essentials for Microsoft Hyper-V), as part of the bundle. It is intended for small to mid-size organizations.

The key benefits that an organization using this blueprint will experience include:

- Affordable HA and DR for even non-critical applications.
- HA with zero downtime during single component failure or environmental failure affecting a single site.
- Simplified one touch automated DR failover.
- Dynamic volume support optimized for Hyper-V.
- Quick and easy HA and DR implementation process.

This paper is written for those who have a working knowledge of Microsoft Windows Server 2008 R2 and virtualized Windows Server environments. This paper also assumes that the reader has a basic understanding of storage virtualization, replication, and clustering terminology.

Introduction

For Information Technology (IT), mitigating the risks to critical data, systems, and applications in addition to computing infrastructures in the event of system outages or complete disasters presents an ongoing challenge from both a technological and a business perspective. Organizations must consistently find solutions that not only meet application and data requirements for capacity, performance, and availability, but also have proven return on investment (ROI) and cost reduction capabilities.

Virtualization has been a game changer for many companies. It has enabled companies that previously were unable to afford HA and DR to begin implementing HA and DR solutions. Virtualization has also enabled companies to justify the cost of providing full HA and DR for additional applications. In addition, it has provided more flexibility and more options in providing the HA and DR solution.

The business challenge is having the ability to create a cost effective, highly available, and protected virtual server infrastructure. This infrastructure must ensure that

applications meet business defined service level agreements (SLAs) for HA and DR preparedness. This white paper provides a solution blueprint that meets these requirements.

Microsoft Windows Server 2008 R2 Failover Clustering, Hyper-V, coupled with HP's virtualized P4000 iSCSI SAN solution, and the Citrix StorageLink Site Recovery software can be used to build an end-to-end robust, highly available, and cost effective HA and DR solution. In order to understand this solution, the key concepts and technical components that make up the solution must first be understood.

Key Concepts

The foundation of the blueprint is the virtualization of the server and storage infrastructure which requires a basic understanding of the virtualization concepts. In addition, there are some important HA and DR technologies that need to be defined before discussing the solution.

Virtualization

Server virtualization is a hot trend in the IT world because of the many business and technical benefits it can provide over the near and long term. There are many types of virtualization that can be utilized in the datacenter, but for the purpose of this paper, we will focus on server virtualization, storage for server virtualization, and storage virtualization.

Server virtualization

Server virtualization is increasingly becoming a more important tool for reducing cost, increasing availability, and enhancing business agility. By enabling organizations to run multiple operating systems and applications on a single physical server versus multiple physical machines, server virtualization can help organizations reduce hardware, energy, and management overhead significantly.

From an application perspective, server virtualization enables organizations to consolidate workloads. This increases resource utilization and lowers capital and operational costs. In addition, server virtualization can help organizations improve application performance, availability, management, and agility. The more agile the IT environment is, the more an organization is able to meet changing business and application requirements.

However, server virtualization changes the accompanying storage requirements. Organizations are unable to realize the full benefits of server virtualization without optimizing and virtualizing their accompanying storage environment. For example, many of the additional virtualization benefits, including built-in, cost-effective HA and DR, require external shared storage with a comprehensive feature set to support them.

Storage for server virtualization

The type of storage platform that organizations deploy to support a virtualized server environment is very important. For example, HA with virtualized servers require external shared storage. In addition, the right storage system platform should ensure low latency for maximum performance, provide scalable capacity and bandwidth as application demands change, and automatically balance application workloads as needed. To support the dynamic Hyper-V virtualized environment, the storage platform must make it easy to create, size, and move storage volumes as required. Also,

because server consolidation through virtualization places multiple applications on a single server, HA features become increasingly important for both the servers and the shared storage. Finally, the right storage platform should also give organizations the ability to virtualize their storage environments by pooling the disk resources to optimize and protect the application environment.

Note: *Storage for server virtualization is different than storage virtualization. Storage for server virtualization is externally shared storage optimized to better serve a virtualized server environment. Storage virtualization refers to the process of abstracting logical storage from physical storage. This abstraction provides numerous advantages and capabilities such as thin provisioning or dynamic snapshots that are not available to storage that is not virtualized.*

Storage virtualization

By abstracting the logical storage from the physical storage, the dynamic datacenter gains many advantages including the ability to migrate data while maintaining concurrent user access. Utilization can also be increased using thin provisioning services that eliminate allocated but unused capacity. This enables administrators to allocate any size virtual volume upfront, but only consume actual physical capacity when the data is written by the application. Also with storage virtualization, there are fewer points of management. The multiple independent storage devices that might be scattered over a network appear to be a single monolithic storage device that can be managed centrally.

HA and DR server functionality

Clustering is a key server function in a HA and DR environment. In addition to clustering, virtualization enhances DR capabilities and can improve DR performance**.**

Clustering

Clustering is normally divided into three conceptual types; failover cluster, load balancing cluster, and grid computing. All three clustering capabilities are in the Windows Server System™. In this paper, failover clustering is the focus.

Failover clustering is a very mature technology that can be used for mission critical applications such as file and print servers, application servers, database servers, and so on. A cluster enables two or more servers to work together as a computer group that can provide failover and increase the availability of the application and the data in any situation. In the event of a primary node going down, failover software based on a heartbeat technique triggers an automatic restart of services on the secondary nodes of the cluster. There are two types of failover clustering; local clustering and stretch clustering:

- **Local clustering**: With local clustering, all cluster participant nodes are at the same facility or datacenter and are physically coupled with the heartbeat link. This configuration can provide application failover, but cannot sustain hosting during downtime that affects the entire facility or datacenter. For example, if the whole datacenter is affected by a catastrophic event, the entire facility is subjected to downtime and does not provide maximum uptime. Still, local clustering is a preferred solution for applications that need to be failed over immediately. Authentication domain servers and financial transaction Web servers are some

prime examples of servers that can negatively affect the infrastructure if they have any failover delay.

- **Stretch clustering:** Stretch clustering or geographically dispersed clustering mitigates the issues involved with local clustering. When a primary site goes down due to natural or man-made disasters, local clustering is not enough to achieve the required uptime of mission critical applications. If a specific site has clusters spanning different seismic zones, applications can be failed over to the secondary site that is unaffected by the primary site downtime. Stretch clustering writes data to both the primary storage system and the remote storage system. This extends the capabilities of a single failover cluster solution and guards against downtime with Windows Server Failover Clustering (WSFC).

Virtualization support

WSFC supports physical and virtual environments including physical to physical, virtual to virtual, and physical to virtual configurations.

HA and DR storage functionality

Server virtualization addresses business continuity planning, but it also places requirements on storage that are often overlooked. The HA functionality of server virtualization software demands the continuous availability of shared storage volumes. DR functionality requires consistent, point-in-time remote copies and must be supported by Microsoft Windows Server 2008 R2 Hyper-V. An HP P4000 SAN is scale-out storage that supports both HA and DR with cost-effective solutions that are often superior to those offered by traditional scale-up SAN products.

Synchronous replication for HA

P4000 SANs have integrated, no-cost synchronous replication that supports both single-site and multi-site HA with a level of storage efficiency not available with most SANs. HA is implemented using a combination of storage clustering and Network RAID technologies:

- **Storage clustering**: Storage clustering creates a scalable storage pool by aggregating the resources of a number of storage nodes into a single storage pool. The pool accepts and responds to iSCSI requests as a single unit. All physical capacity and performance are aggregated and made available to the volumes created on the SAN. When more storage is needed, additional storage nodes can be added to the pool. The P4000 SAN seamlessly and non-disruptively reorganizes its storage to accommodate the new systems.
- **Network RAID**: Network RAID synchronously replicates the blocks that make up a logical volume across one or more storage nodes in the storage cluster for HA. A Network RAID level 2 causes every block that makes up a logical volume to be stored on two storage nodes. With Network RAID level 2 implemented, a single component, a storage system, or in the case of a geographically split SAN, an entire location can fail and the data for each volume will remain available. Failover and failback are automatic and transparent to application servers because the pool itself keeps track of which storage blocks are up-to-date, and automatically updates outdated blocks when the failed part of the pool is restored.

HA is cost-effective with the P4000 SAN because synchronous replication is a no-cost option. Companies only pay for the HA they require because Network RAID is

configured on a per-volume basis. In addition, storage ROI increases because storage is used more efficiently. For example, volumes can be thin provisioned. In addition, volumes and their replicated blocks are stored in a single storage system. This increases efficiency by reducing fragmentation.

Note: *It is important to note that as opposed to a HA failover option that might quickly regain functionality, Network RAID provides continuous availability as long as the other nodes remain available. Network RAID provides zero down time during a node failure. This is similar to how disk RAID deals with a drive failure.*

Asynchronous replication for disaster recovery

P4000 SANs support DR through a combination of thin provisioning and an asynchronous replication feature called Remote Copy that maintains both time and space efficient remote copies at a DR site. Remote Copy can be supported in a standard two-site configuration. It also can be supported between any number of remote offices and a central site using a replication client that runs on remote sites and requires no additional hardware.

In addition, Citrix StorageLink™ Site Recovery software is used to provide one touch failover in case of a disaster. It works as the big red button. The new StorageLink technology assures that any environment virtualized with Hyper-V will work seamlessly with all existing storage, backup, and DR systems.

HA and DR Solution Components

A highly efficient HA and DR solution can be created by combining Microsoft, HP storage and Citrix StorageLink Site Recovery technologies. This section discusses the relevant technologies required.

Microsoft technologies

Microsoft virtualization is an end-to-end strategy that can profoundly affect nearly every aspect of the IT infrastructure management lifecycle. It can drive greater efficiencies, flexibility, and cost effectiveness throughout the organization.

A standard Microsoft virtualization implementation is typically structured using the Windows Server 2008 R2 Hyper-V role to enable virtualization and Windows Clustering to handle HA and DR requirements. System Center Virtual Machine Manager (SCVMM) is typically used to simplify virtualization management.

Windows Server 2008 R2 Hyper-V

Hyper-V is the hypervisor-based virtualization technology from Microsoft that is integrated into all Windows Server 2008 R2 x64 Edition and Windows Server 2008 R2 operating systems. As a virtualization solution, Hyper-V enables users to take maximum advantage of the server hardware by providing the capability to run multiple operating systems (on virtual machines) on a single physical server.

The availability of Hyper-V as a role in a mainstream Windows operating system provides several key advantages:

Features	Benefits
Built in technology	Hyper-V enables enterprises to easily leverage the benefits of virtualization without adopting a new technology.
Broad device driver support	The new 64-bit micro-kernelized hypervisor architecture leverages the broad device driver support in the Windows Server 2008 R2 parent partition to extend support to a broad array of servers, storage, and devices.
SMP support	Hyper-V supports Symmetric Multiprocessors (SMP) in virtual machines (VMs).
HA	Windows Server 2008 R2 Failover Clusters provide HA to VMs to minimize unplanned downtime.
Shared storage HA	Microsoft MPIO dynamically routes I/O to the best path and protects against connection failures at any point between a Hyper-V host and shared storage including NICs/adapters, switches, or array ports.
Easy VM migration	Live Migration is available with Windows Server 2008 R2 to support business continuity during planned and unplanned downtime and over distances.
Volume Shadow Copy Support (VSS)	Robust host-based backup of VMs is available by leveraging the existing Windows VSS-based infrastructure.
Easy extensibility	Easy extensibility is available by using the standards-based Windows Management Instrumentation (WMI) interfaces and APIs.
Simplified integrated management	With its tight integration into the Microsoft System Center family of products, customers have end-to-end physical and virtual infrastructure management capability for Hyper-V environments.

Table 1: Hyper-V features

Windows Server Failover Clustering

WSFC helps to ensure that mission-critical applications and services, such as e-mail and line-of-business applications, are available when needed. Beyond the capabilities already mentioned in the previous stretch clusters section, some other important capabilities of WSFC for DR solutions include:

Features	Benefits
Limitless single subnet	Enables cluster nodes to communicate across network routers. The need to connect nodes with VLANs is no longer necessary.
Configurable heartbeat timeouts	Provides the ability to increase timeouts to extend geographically dispersed clusters over greater distances as well as providing the ability to decrease timeouts to detect failures faster and take recovery actions for more efficient failover.
Common toolset	Provides a similar management experience to managing a local cluster to simplify the process.
Automated failover	Provides automatic failover in the event of a complete disaster in one site.
VSS	Provides the ability to back up cluster settings.
Automation support	Automation provided by Windows Server 2008 R2 and PowerShell simplifies system administration.
Cross-site replication tools	Provides mirrored storage between stretched locations in addition to seamless integration with partner hardware or software-based data replication solutions.

Table 2. Windows Server Failover Clustering features

System Center Virtual Machine Manager

Microsoft SCVMM 2008 is enterprise-class management software that enables administrators to easily and effectively manage both the physical and virtual environments from a single management console. This helps administrators avoid the complexity of using multiple consoles typically associated with managing an IT infrastructure. The key capabilities of SCVMM 2008 include:

Features	Benefits
Enterprise-class management suite	Manages both Hyper-V and VMware ESX virtualization environments.
Intelligent VM placement	Provides support for the intelligent placement of VMs.
System Center Operations Manager 2007 integration	Integrates with System Center Operations Manager 2007 to provide proactive management of both virtual and physical environments through a single console by leveraging PRO.
Native P2V/V2V migration	Provides native capability for physical-to-virtual migration and virtual-to-virtual migration.
Failover integration	Provides integration with failover clustering to support HA and the live migration of VMs.
Automation	Provides easy automation capabilities leveraging Windows PowerShell™.

Table 3. System Center Virtual Machine Manager features

HP StorageWorks

HP StorageWorks offers a comprehensive set of storage products spanning environments from the sole proprietorship to the largest enterprise. Within server virtualization environments, the MSA 2000 array provides the very best price performance while meeting the general requirement of a shared storage environment. On the other end of the spectrum, the Enterprise Virtual Array (EVA) product line offers a full feature array family that completely supports all the features of Hyper-V and WSFC. The EVA product line is integrated with the full set of HP data center management tools. The HP P4000 SAN is the perfect choice for the mid-market, offering all the features required for server virtualization plus a high degree of automation, scalability, and manageability.

HP P4000 SAN

The HP P4000 SAN solutions are optimized for database and e-mail applications as well as virtualized servers. With this all-inclusive pricing model and intuitive storage management software that enables companies to pay as they grow, the SAN is perfect for the budget minded. Furthermore, HA and DR no longer seem out of reach. The P4000 eliminates single points of failure across the SAN with an innovative approach to data availability, reducing risk without driving up costs.

In addition, uncertainty about whether or not current storage can meet future needs is no longer an issue. Built on a storage clustering architecture, the HP P4000 SAN enables linear scaling of capacity and performance without incurring downtime or performance bottlenecks, or requiring expensive upgrades. Finally, the HP P4000 SAN

solutions provide creative ways to stretch the IT budget. For example, thin provisioning and SmartClones SAN solutions substantially increase storage efficiency.

The all inclusive feature set provides enterprise functionality at an affordable price including:

Features	Benefits
Storage clustering	Storage clustering enables customers to consolidate multiple storage nodes into pools of storage. All available capacity and performance is aggregated and available to every volume in the cluster. As storage needs increase, the HP P4000 can scale performance and capacity online.
Thin provisioning	Thin provisioning allocates space only as data is actually written without requiring pre-allocation of storage. This raises the overall utilization and efficiency of the HP P4000 SAN, reduces costs, and ultimately increases the ROI.
Snapshots	Snapshots create thinly provisioned, instant point-in-time copies of data on a per-volume basis. Administrators can access snapshots to recover individual files from the volume, or rollback an entire volume. Built-in integration provides automated quiescing for Microsoft Volume Shadow Copy support (VSS) applications.
SmartClone	A SmartClone creates a shared block copy of a volume. SmartClones are very efficient when many volumes that contain similar data are required. An example of this is when virtual desktop images or a set of VMs all do the same thing.
Remote copy	Remote copy replicates snapshots between P4000 SANs at primary and remote locations. Copies are thinly provisioned with no space reservation required. Remote copy enables centralized backup and DR on a per-volume basis and leverages application integrated snapshots for faster recovery.

Table 4. P4000 enterprise functionality

The HA and DR functionality includes:

Features	Benefits
Network RAID	HP P4000 SANs stripe and mirror multiple copies of data across a cluster of storage nodes eliminating any single point of failure in the SAN. Applications have continuous data availability in the event of a power, network, disk, controller, or entire storage node failure.
Multi-site SAN	Distributing storage nodes across multiple local sites isolates and protects against environmental failures that affect one location.
Integrated replication for DR	HP P4000 SANs include integrated replication at no additional cost that simplifies management with simple failover and failback.
No downtime configuration changes	Administrators can add capacity, increase performance, grow and migrate volumes between HP P4000 SAN clusters on the fly with no application downtime.

Table 5. P4000 HA and DR

The solution provides scalable performance including:

Features	Benefits
Minimum purchase requirement	Avoid up-front cost and potential performance constraints. Purchase only what is needed today, and then grow the performance, capacity, and redundancy of the HP P4000 SAN online as storage requirements evolve. Buying storage only when needed simplifies planning and relieves budget pressures.
Simultaneous performance and capacity scaling	Each time a storage node is added to the HP P4000 SAN, the capacity, performance, and redundancy of the entire storage solution increases.
Disruptive upgrade avoidance	Add resources to the HP P4000 SAN without disruption as capacity and performance requirements increase. Applications remain online during maintenance events for best availability.

Table 6. P4000 scalable performance

The solutions provide easy-to-manage SANs for virtualized environments:

Features	Benefits
Centralized management console	Multiple data centers and sites can be managed from an all inclusive single view. All of the HP P4000 SAN features are managed from the Centralized Management Console for simple, easy-to-manage storage.
Automation	Many typical SAN management tasks are fully automated within the P4000 SAN. For example, volume provisioning and block-level load balancing are done entirely by the system. Due to this automation, managing the SAN does not require an understanding of the internal workings of the system and therefore does not require any specialized

	training or expertise.
Non-disruptive Administration	Virtually all administrative tasks within the P4000 SAN can be done while the system is up and operating without affecting any of the online volumes. For example, upgrading firmware, adding a new storage node to the cluster, changing a volume's protection level, and increasing the size of a volume can all be done without disrupting operations.

Table 7. P4000 virtualization capabilities

HP Hyper-V Smart Bundles

To realize the full potential of Microsoft Windows 2008 R2, HP has preconfigured virtualization bundles that deliver an end-to-end virtualized environment based on Microsoft Windows Hyper-V. They are predefined and tested for performance and scaled for different corporate needs. This helps IT departments avoid complicated configuration and sizing decisions and reduce the implementation time. The Hyper-V Smart Bundles are grouped into three families:

- **Entry**: HP ProLiant Internal Storage with HP SmartArray technology.
- **Midrange**: Simple Direct Attached SAS HP StorageWorks MSA 2000 G2.
- **Scalable HA**: Advanced iSCSI HP P4000 SAN with built in HA and DR.

Citrix Essentials 5.5 for Hyper-V

Citrix Essentials™ for Microsoft Hyper-V is a rich set of automation and advanced virtualization management capabilities that extend the capabilities of Microsoft Hyper-V and SCVMM. The Citrix Essentials product line was built from Citrix expertise in server virtualization and a 20 year partnership with Microsoft. With a history of embracing and extending the Microsoft Windows platform, the Citrix Essentials product line enables customers to get more value from their Hyper-V and SCVMM investments. For example, Citrix Essentials provides customers with the ability integrate past infrastructure investments with new storage system purchases while leveraging the advanced storage functionality of the newer storage arrays. For the purpose of this paper, we will focus on the StorageLink™ component of Citrix Essentials and more specifically, its Site Recovery functionality. Site Recovery provides support for the DR capabilities in this blueprint.

Citrix StorageLink Site Recovery

Citrix StorageLink technology enables virtual server infrastructures to fully leverage all the resources and functionality of existing storage systems. StorageLink accelerates VM deployment and simplifies VM storage provisioning through automation and offloading complex storage copy operations to the storage subsystem. StorageLink supports all third party storage architectures and deeply integrates with leading storage platforms to reduce the cost and complexity of managing storage in Hyper-V environments. It offers one click access to native storage devices, provides administrators with easy access to native storage controls to expedite the roll out and to simplify the management of virtual environments. StorageLink also simplifies the setup and management of WSFC by automatically configuring the storage for all the hosts in the cluster and seamlessly integrating incremental hosts or storage into the cluster.

Citrix StorageLink Site Recovery enables organizations to implement flexible, cost effective DR solutions for their Hyper-V infrastructure. Site Recovery leverages underlying storage array replication services giving administrators the toolset they need to enable remote failover strategies for their Hyper-V environments. Citrix Site Recovery includes simple setup, configuration, staging and DR plan testing. Together with workflow orchestration and third party global clustering solutions, Site Recovery enables organizations to implement fully automated DR plans for fast, reliable site recovery of the critical Hyper-V infrastructure.

The StorageLink Site Recovery feature set includes:

Features	Benefits
Simplified configuration	Site Recovery can be configured in a few simple steps using point and click wizards and a single console to simplify setup and configuration. Examples include: • Primary and secondary array designation. • Storage repository creation at both the primary and secondary site. • Storage repository for replication protection and mapping. • Ability to select VMs to protect. • Ability to export configuration data. • Ability to log in to a secondary site and import configuration data.
Automated replication configuration	Site recovery leverages the P4000 native replication capability Remote Copy feature to move data from the primary site to the secondary site. This replication can be established using the StorageLink interface, reducing the time required to set up DR and ensuring fast and accurate data migration.
DR validation	After the data is mirrored and the metadata is transferred to the secondary site, the Validation feature confirms that there are enough hosts with memory and CPUs available at the secondary site. It further confirms that the required networks are present at the secondary site in addition to the boot and data disks for the VMs. This ensures that all resources are present for a seamless recovery.
Staging and testing	Site Recovery enables VMs to be tested and staged at the secondary site to ensure that recovery time objectives and recovery point objectives can be met. The administrator can test possible DR scenarios and ensure confidence that the recovery will occur as planned. A snapshot or clone is created from the VMs and they are bought up in isolated networks in order to avoid disrupting the ongoing replication. When the testing is complete, the VMs can be discarded.

Failover	Failover can be enabled manually to give control to the administrator. Alternatively, an administrator can write a script that integrates with Microsoft Failover Clustering to set up an automatic failover.
Application integration	A software development kit is provided so that recovery services can be integrated directly with existing system management applications. This makes it easier than ever to protect key assets.
SCVMM integration	SCVMM can be used to create VMs. A PowerShell script can be used to provision storage for these machines yet keep SCVMM as the primary management interface. Furthermore, the recovered VMs at the secondary site can be started using SCVMM, making it easier for customers to leverage their investment in Microsoft System Center.

Table 8. StorageLink Site Recovery features

***Note**: HP Smart Bundles include Citrix StorageLink Express. Available at no charge, StorageLink Express supports two servers and one P4000 storage cluster. Customers that require additional support will need to upgrade to the full version of Citrix Essentials for the Hyper-V suite.*

Blueprint

The purpose of this blueprint is to provide an end-to-end, robust, and cost effective HA and DR protection solution that any moderately skilled administrator can implement for their organization without the expense of outside assistance. This section provides the basic information needed to configure the entire solution. To make things simple, an HP Hyper-V Smart Bundle can be purchased to provide the entire hardware and software requirements for building this exact blueprint.

The blueprint is a generic design that can be modified to fit any organization's needs. As it was developed to be tested in the lab to demonstrate HA and DR protection at hardware, datacenter, and campus levels, the lab setup required three labs for a realistic demonstration. Two labs were located in Santa Clara, California on the same campus and the third was located in Redmond, Washington. To make the lab work easy to follow and replicate, the blueprint uses the true locations, but an organization can make the appropriate substitutions.

Overview

The blueprint lab setup consisted of a SQL Server® workload running in a Hyper-V VM. This VM was made highly available through the use of WSFC, where each machine in the cluster was located in a separate physical room in the primary data center in Santa Clara, CA. This protected the workload against local storage, network, or power failures.

The Santa Clara site was protected against local natural or man-made disasters using StorageLink Site Recovery while the failover datacenter was located in Redmond, WA.

All storage in this blueprint was delivered by HP P4000 storage.

Setup detail

The following diagram shows the lab setup in Santa Clara and Redmond.

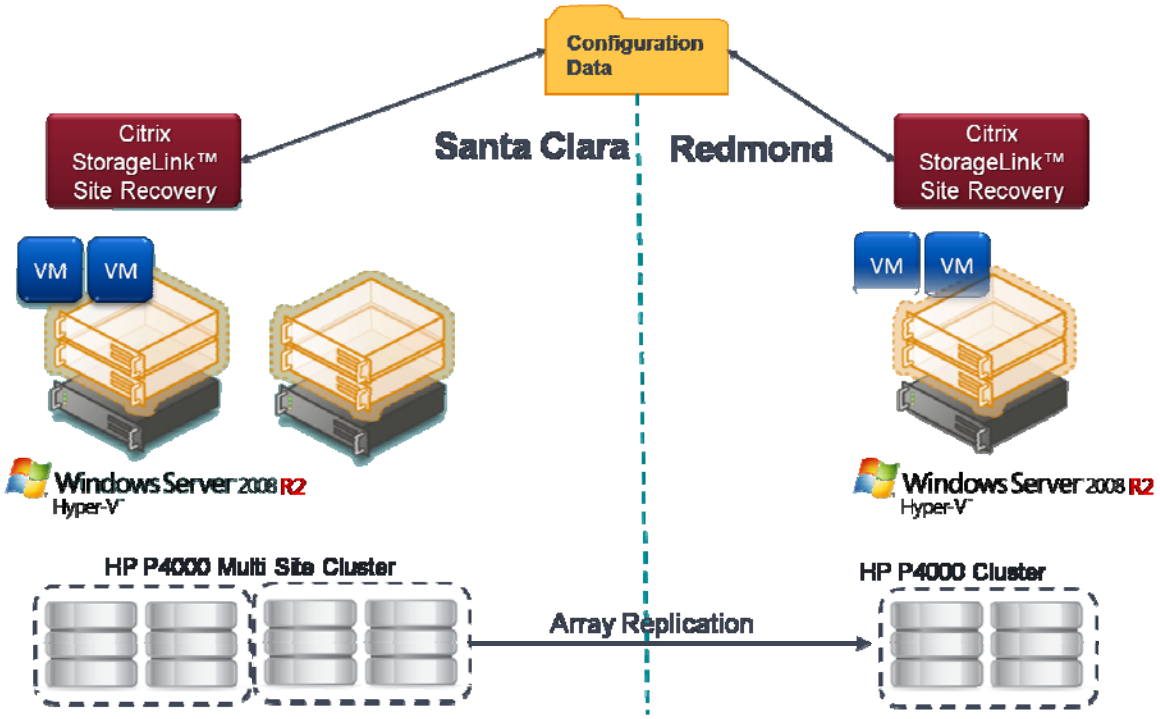

Figure 1: Blueprint overview as set up for lab testing

Figure 2 displays a more detailed diagram of the multi-room lab in Santa Clara. The diagram shows how both the failover server cluster and the storage cluster are spread across the two rooms. They will both work as single logical clusters, but in the event of an environmental problem in one of the rooms, the systems in the other room can take over.

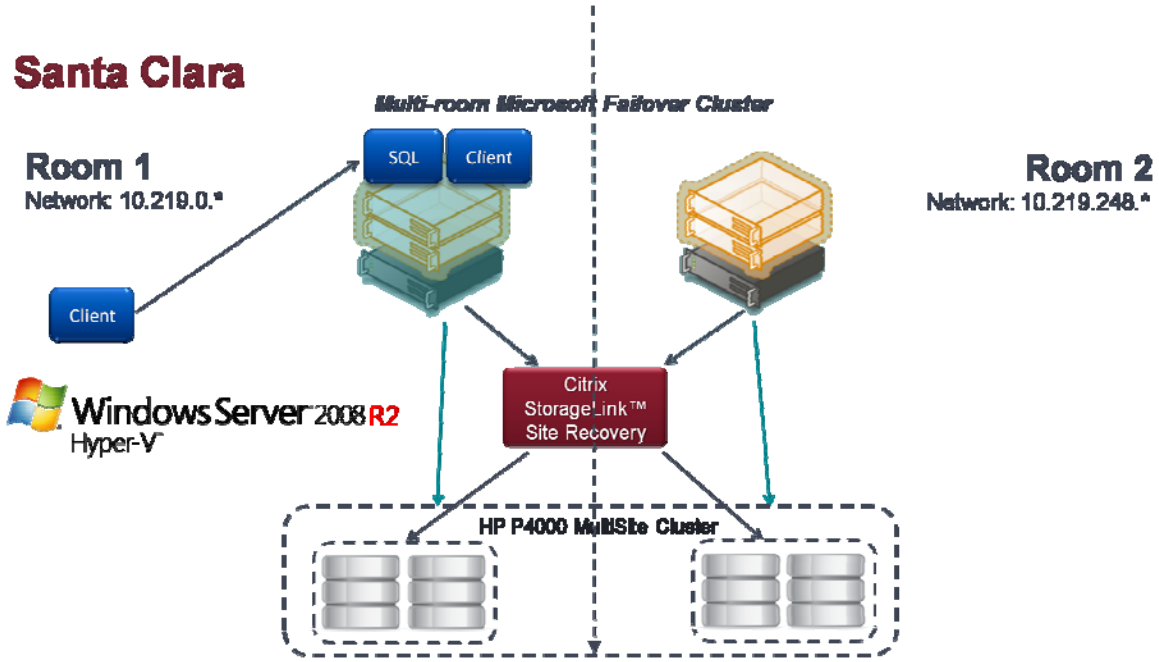

Figure 2: Santa Clara, CA multi-room lab detail

Figure 3 shows a detailed view of the Redmond datacenter. Redmond is configured as the DR site and has enough storage and server capacity to take over and run critical applications in the event of the datacenter in Santa Clara going offline.

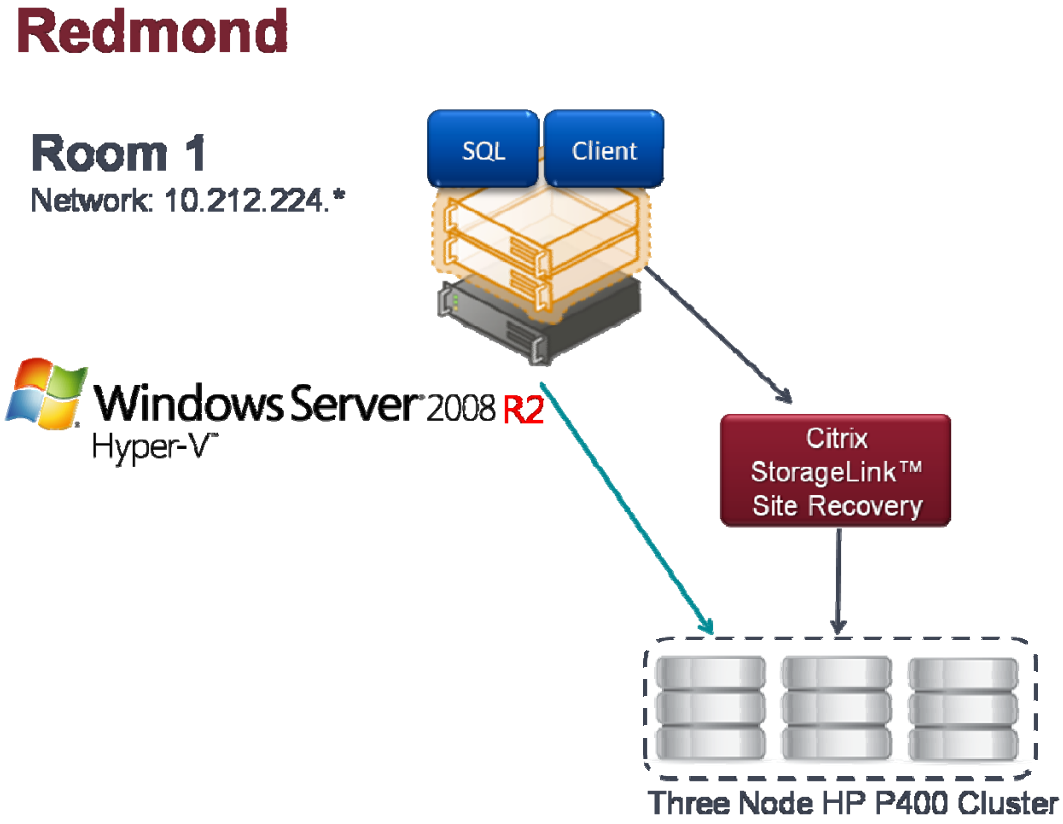

Figure 3: Redmond, WA multi-room lab detail

Workload

An application workload was selected to run based on Microsoft SQL Server 2008. The Microsoft DBHammer tool was used as a database client.

Database server details:

1. Microsoft SQL Server 2008 (64-bit), MSDN edition
   - Standard setup
2. Installed in a Hyper-V VM
   - Guest OS: Windows 2003 R2 in Windows Server 2008 R2 Hyper-V
   - 8 GB RAM
   - Passthrough 60 GB OS and Application Virtual Disk
   - 4 vCPUs
3. Database created
   - Name “Test”
   - Windows credentials
   - Filled with DBHammer LoadMaster application resulting in 10,000,000 rows in the database.
4. Windows Server 2008 R2 Hyper-V details
   - Installed on a HP DL380 G5
   - 8 GB physical memory
   - 2 x Intel® Xeon® E5310 1.6Ghz 4-core CPU

Database client details

DBHammer details:

- Provided in Microsoft SQL 2000 Resource Kit
- Database setup with LoadMaster
- Database test runs with ProcMaster
- Default settings

Sites

Data centers in two physical locations were selected to provide redundancy to the above mentioned workload. The primary location was in Santa Clara, CA and the failover location was in Redmond, WA. Both sites are located in the Citrix Systems, Inc. engineering labs.

The Santa Clara datacenter setup had two physical separate rooms with a physical machine in each room running in a two node WSFC. The Redmond datacenter was located in a single room setup with one physical computing machine running Microsoft Windows Server 2008 R2.

Storage

The storage in the Santa Clara datacenter was provided by a multi-site SAN consisting of two HP P4000 storage arrays in each room, forming a four node storage cluster. Each node contained eight physical 420 GB disks, forming a total storage capacity of 2.81 TB of useable space per storage node.

One separate HP P4000 storage array was added to the management group as a failover manager. This array did not provide any storage services, but merely acted as a tiebreaker in case of a two out of four node failure.

The software version installed on all nodes was 8.0.00.1724.0.

The Redmond datacenter had a three node, single site HP P4000 storage cluster. Each node contained six physical 136 GB disks, forming a total storage capacity of 662 GB per storage node.

Network

The storage was connected to the computing nodes though one dedicated gigabyte Ethernet network. Two subnets were used in the two rooms in Santa Clara. This arrangement provided some performance benefits in terms of network preferences within the SAN.

The datacenters in Santa Clara and Redmond were connected by a company VPN tunnel over public internet connections with a maximum throughput of about 30 Mbit/s.

Hyper-V Servers

The two servers in the Santa Clara datacenter were HP DL-380 G5 machines with 8 GB of physical memory. Microsoft Windows Server 2008 R2 with the Hyper-V role enabled were installed on both servers.

Each machine was connected through two Ethernet connections; one to the room 1 network and one to the room 2 network.

Citrix StorageLink

The Citrix StorageLink gateways in both nodes were set up inside a VM installed on Windows Server 2008 R2. There was a network file share configured so both nodes can read and write files to this file share that was used to transfer the Site Recovery configuration files between the two sites.

The latest StorageLink release, StorageLink v2.1 plus Hotfix 1, was used for this setup.

Site Recovery

Site Recovery was set up with the following parameters:

Primary site:

- Santa Clara datacenter
- Configuration Data path:
  Default (`C:\ProgramData\Citrix\StorageLink\Server\SRPExports`)

Secondary site:

- Redmond datacenter
- Configuration Data path:
  `C:\Configuration Data` (location of network share)

VM protection was set up using default arguments and parameters. Manually, a remote snapshot schedule was established to create a new remote snapshot once every 30 minutes while preserving one snapshot on both primary and secondary sites.

Integration

The setup detail for this blueprint is covered in this section. Once the infrastructure is set up, the initial VMs must be created using StorageLink to establish the connectivity to the storage arrays. The VMs are then managed as usual with SCVMM. Once the basic blueprint setup is complete, the site recovery plan must be set up.

Setup procedure

Many components need to be configured in a detailed setup. In this case, the following steps were taken:

1. Multi-site HP cluster setup and configuration.
2. Two node MSFC setup spanning two physical rooms in the Santa Clara datacenter with one physical computing node in each room.
3. Citrix StorageLink Gateway service installation together with a StorageLink Manager for managing the storage and computing nodes.
4. Windows Server 2008 R2 operating system installation in a VM.
5. SQL Server 2008 database application installation and configuration.
6. DBHammer tool setup to be the database client of the SQL Server 2008 application.
7. Redmond datacenter procurement, installation, and configuration:
    a. A computing node installed with Windows Server 2008 R2.
    b. Windows Server 2008 R2 Hyper-V VM installed for the Citrix StorageLink gateway.
    c. Attained access to the existing HP storage cluster.
8. System testing of progressing failover of all involved components.

Virtual Machine setup

The initial VMs must be created by the StorageLink Gateway Manager and Service utilizing StorageLink Storage Adapters for deep integration with the HP P4000 storage arrays. The adapters allocate the right amount of storage, snapshot the right golden image volumes, and assign the right LUNs to the right Microsoft Hyper-V host. In short, Citrix StorageLink provides a very user friendly way of provisioning and spinning up new VMs.

After the VMs have been created with Citrix StorageLink, the VMs are managed through SCVMM as normal.

Note: *To get the HA and DR benefits of StorageLink, the initial VMs must be created by StorageLink to integrate the storage connectors. The VMs can then be managed through SCVMM as normal.*

Site Recovery plan

The Citrix StorageLink Site Recovery option is used to preprogram the Site Recovery plan and set up replication between the HP P4000 storage clusters in geographically dispersed data centers. If a disaster occurs, the StorageLink big red button is pushed to automatically recover the production application at the DR site.

Lab Testing

The purpose of this blueprint is to provide robust HA and DR protection, including:

- **Disk RAID 5:** Set within each storage node to protect against disk failures.
- **P4000 Network RAID, level 2:** To protect against storage node failures.
- **P4000 multi-site storage clusters:** To protect against network or power failures in separate locations in a datacenter or at a local campus.
- **Microsoft Windows Server Failover Cluster:** To protect application workloads against server failures.
- **Citrix StorageLink Site Recovery:** To protect a datacenter against man-made or natural disasters.

To demonstrate the value of each component, the blueprint was set up in a lab to test and demonstrate each of these components.

Testing procedure

In sequential order, the following tests were performed in this lab:

1. Pulled one disk from a storage node from the primary site of the HP cluster.
   - Disk failure - Disk RAID redundancy
2. Pulled another disk from the same storage node.
   - Node failure – Network RAID redundancy
3. Failed both storage nodes in the primary site.
   - Site Failure – Multi-site HP cluster failover
4. Failed the computing node in the primary room.
   - Node Failure – Windows Server Failover Cluster
5. Failed storage and computing nodes in the secondary room
   - Datacenter Failure – Site Recovery failover
6. Failback to the Santa Clara datacenter
   - After all computing and storage resources were recovered, the failback procedure was used to move the current version of database data back to the Santa Clara datacenter.

Each step created a more serious failure, going from a single disk failure to a complete physical site failure. To display the outcome of these steps, results were taken and saved using the Microsoft Windows Performance Manager application. The performance manager was configured to track the number of MSSQL transactions per second.

Test 1: Pull one disk

The green arrow in the following screenshot marks the point in time where the first disk was pulled. This caused a slight blip while the array detected that the disk was inaccessible causing I/Os to rely on disk RAID 5 protection. Database access from the DBHammer client was not interrupted.

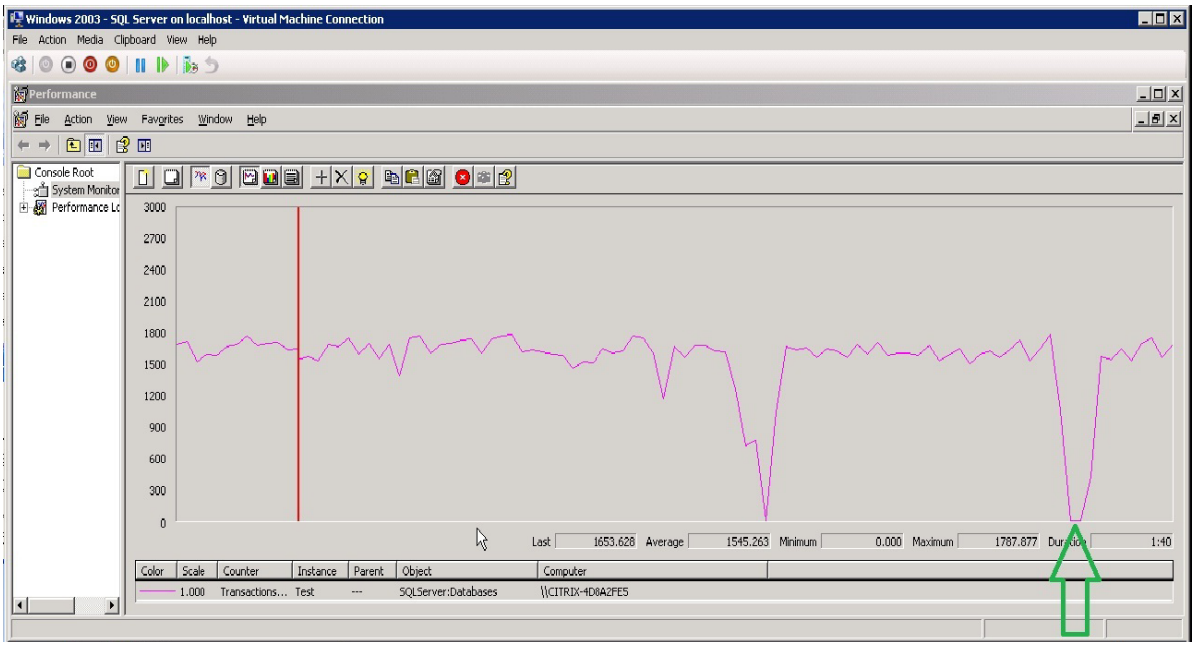

Figure 4: Test 1 performance results

Test 2: Pull second disk

The green arrow in the following screenshot marks the point in time where the second disk was pulled. This caused a bigger transaction blip while the cluster reconfigured itself to not use any storage from the first storage node. During this reconfiguration, database access from the DBHammer client was not interrupted.

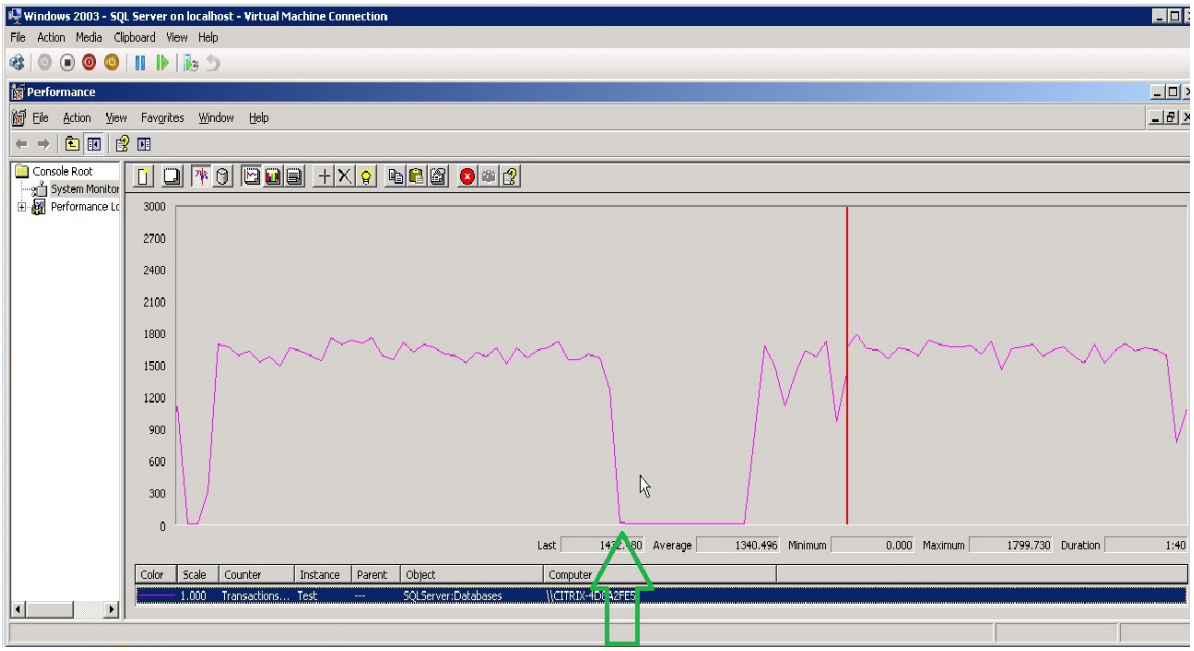

Figure 5: Test 2 performance results

Test 3: Disabling storage in room 1

The green arrow in the following screenshot marks the point in time where the power was pulled to all storage in Room 1. This caused a bigger transaction blip while the iSCSI session reconfigured itself to use the virtual IP of the second site for storage traffic. During this reconfiguration, database access from the DBHammer client was temporarily paused, but no connection failures occurred.

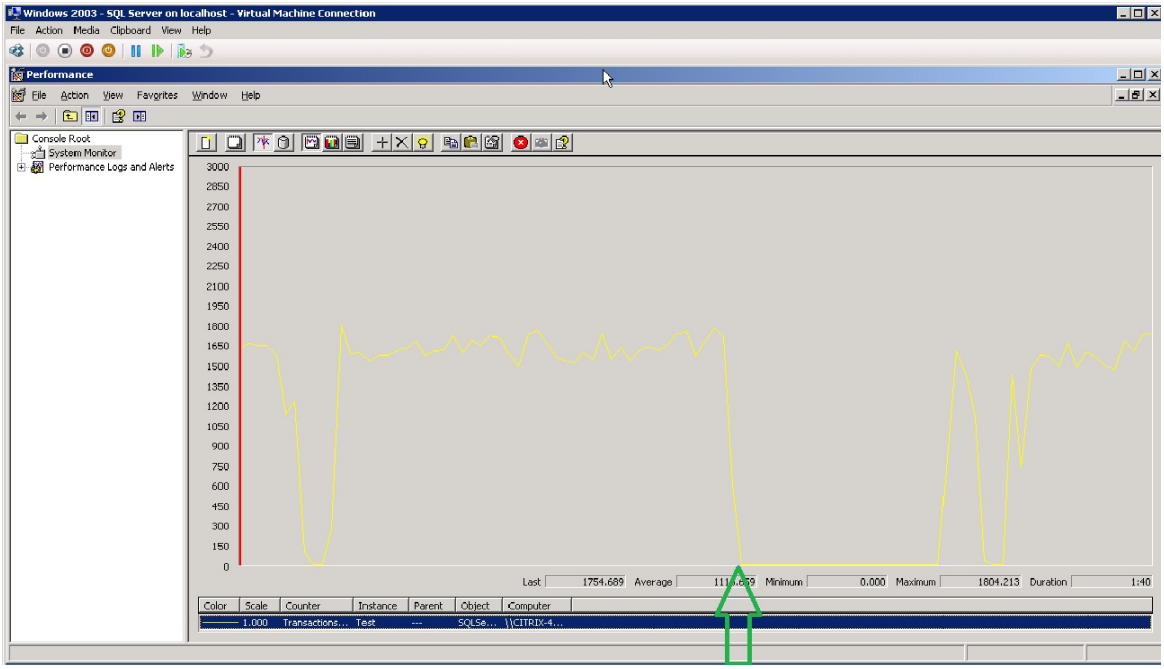

Figure 6: Test 3 performance results

Test 4: Disabling computing node in room 1

In this test, the power was pulled to the computing node in Room 1 causing the VM running on this machine to fail and subsequently the workloads running inside the VM to fail. This failure is detected by WSFC. WSFC reacted to this failure by restarting the VMs on the computing node in Room 2. Because StorageLink already had configured all the VM related storage to be available on the second node, this failover happened automatically without any user intervention.

Because the failover was a full failover, meaning the VMs were shut down and restarted on the computing node in Room 2, it was not possible for the DBHammer client to continue working without interruptions. Once the VM and Microsoft SQL Server started, the DBHammer client was restarted.

Test 5: Failing all resources in room 2

In this step, all computing and storage resources in the second room were powered off. This caused a complete site failure of the Santa Clara location.

Because a Citrix StorageLink site recovery plan was already in place, failing over the services provided by the VMs to Redmond was very quick and easy. This is because the recovery plan sets up the storage replication between the two P4000 clusters in California and Washington in addition to providing an easy way to manage the VM configuration data.

Before this man-made disaster happened, the recovery plans were validated and tested using the Citrix StorageLink Manager connected to the datacenter in Redmond. The validation checks whether a particular VM resource can indeed failover to the managed datacenter. It does this through a series of checks including computing resource availability, the presence of the right networks and most importantly, if the replicated storage can be discovered on the storage systems in the datacenter. All results were categorized and displayed as shown in Figure 7.

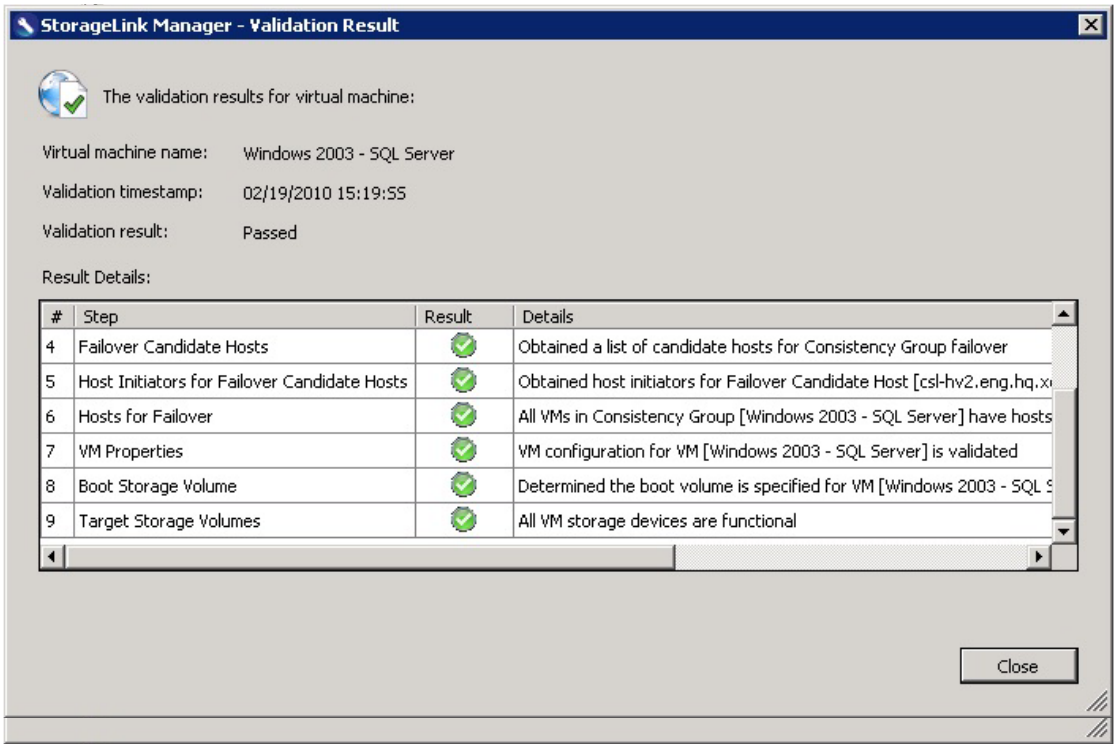

Figure 7: Test 5 validation results

After the power was turned off for the assets in the Santa Clara data center, the failover of the VM resources to the Redmond datacenter was performed by selecting the Failover Recovery Configuration operation as shown in Figure 8.

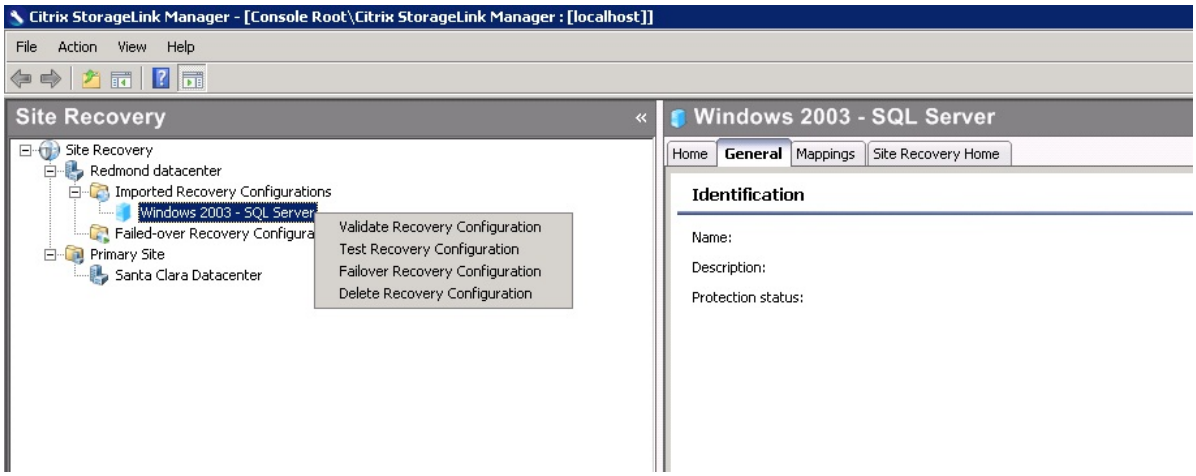

Figure 8: Failover Recovery Configuration operation

Once the failovers were complete, the VM was restarted and the DBHammer client was kicked off again to continue testing database access.

Test 6: Failback to the original datacenter in Santa Clara

In this step, the computing and storage resources in the Santa Clara datacenter were brought back up and the HP Centralized Management Console was used to commence a failback operation for all the volumes that were failed over to Redmond in Step 5. For the failback operation, the *move the primary volume* workflow in the Failover/Failback Volume wizard was used.

After the failback operations were complete and the remote snapshot schedule from storage in Santa Clara to the storage in Redmond was reestablished, the StorageLink service was started in the primary site. To demonstrate that the synchronization from Santa Clara to Redmond was complete, the protection status changed back to 'Active'. This status indicts that all application workloads in the Santa Clara datacenter are protected against natural and man-made disasters.

Test results

All tests were performed successfully and all available redundancy protections were sufficient to withstand the increasing levels of failure introduced for this blueprint lab setup.

Best Practices

At the conclusion of the lab testing, the lab personnel expressed that it was relatively straight forward to set up all the components of the lab blueprint. They further stressed that they only had a reasonable background in Windows administration and did not have deep knowledge of the HP P4000 storage arrays, WSFC, or Microsoft SQL Server. They concluded that if they had to set up the lab again, they could probably do it within one week of dedicated lab time.

Here are some of the best practices they used to successfully set up the HA and DR environment:

- **Install HP P4000 MPIO DSM on all Microsoft Hyper-V machines**: This is to enable multipathing to the HP P4000 arrays in the cluster to provide protection against path failures and increase performance by using multiple active paths for I/O. The most current DSM for Microsoft MPIO can be downloaded from the HP Web site.
- **Ensure that the virtual IPs for both rooms are entered in the targets for the Microsoft iSCSI initiator**: An HP P4000 multi-site SAN has multiple virtual IP addresses; one for each of the sites defined in the HP Management Console. Be sure to add each of these IP addresses as Target Portals on the Discovery tab of the iSCSI initiator management utility.
- **Set up a suitable HP remote snapshot schedule**: Set up an HP remote snapshot schedule suitable for the application workloads that are protected using Citrix StorageLink Site Recovery according to your best practices.
- **Add a failover manager**: One separate P4000 array was added to the management group to provide tiebreaker capabilities in case of a situation where two out of four nodes are down. This prevents the cluster from forming a quorum. This can also be accomplished by adding the Failover Manager software service. This component is provided free-of-charge with the P4000 SAN and mitigates the need to run an additional node.

The following HP documentation was used for the setup and testing of this lab:

- Building high-performance, highly available IP storage networks with HP SAN/iQ® Software
  http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA2-5615ENW.pdf
- Multi-Site SAN User Manual
  http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c01806779/c01806779.pdf
- Remote Copy Networking Requirements
  http://bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c01750200/c01750200.pdf
- Remote Copy User Guide
  http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c01727843/c01727843.pdf
- HP StorageWorks HP LeftHand P4000 Multi-Site HA/DR Solution Pack - User Manual
  http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c01727773/c01727773.pdf
- HP LeftHand P4000 Windows Solution Pack - HP LeftHand P4000 DSM for MPIO - User Guide
  http://bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c01727956/c01727956.pdf

Conclusion

This blueprint delivers an end-to-end affordable HA and DR solution thanks to virtualization. The combination of server virtualization and management tools from Microsoft, advanced storage virtualization technologies from HP, and one touch multi-site failover capabilities from Citrix enables the creation of highly efficient, flexible, and easy-to-manage virtualized environments that provide HA and DR capabilities even for the non-critical applications.

Although all tests were performed successfully as expected, the more interesting observation made by the test team was the simplicity of the setup. Their feelings were that an administrator of moderate experience should be able to set this blueprint up without the need for additional and often expensive outside assistance. Furthermore, as everything required was provided by the HP Hyper-V Smart Bundle, there was very little planning needed. Their final observation was that the automation this configuration provided created a very easy-to-use environment.

HP developed an innovative SAN that reduces storage costs, cuts storage administration time, and provides continuous data availability. By packaging this solution in HP Smart Bundles for Hyper-V that provide the entire virtualization infrastructure, HP has simplified and shortened the entire implementation process.

Additional Information

For those interested in implementing the blueprint provided by this paper, see the links below for additional information. To take the next step, contact a HP sales representative at (800) 888-9909.

For more information, see:

- Microsoft Virtualization Solutions:
  http://www.microsoft.com/virtualization/solutions
- HP P4000 SAN Overview
  http://www.hp.com/go/P4000
- Citrix StorageLink Site Recovery
  http://www.citrix.com/English/ps2/products/subfeature.asp?contentID=1859615