Deployment Blueprint for Virtual Desktop Infrastructure on Microsoft Windows Server 2008 R2 SP1 and NetApp Unified Storage Solutions

Including production considerations and best practices

Solution Blueprint

Published: October 25, 2011

Abstract: NetApp unified storage can provide enhanced storage performance, storage efficiency, and storage manageability for VDI deployments on Microsoft® Windows® Server 2008 R2 SP1 Hyper-V™.

Executive summary

Virtual Desktop Infrastructure (VDI) delivers dedicated desktop environments to users from servers in the data center. This is in contrast to the traditional model in which each user has one or more physical desktop machines, each running its own operating system and applications. Some of the potential benefits of deploying VDI include:

- **Streamlined and centralized management of desktops**, including the ability to quickly provision new desktops based on one or more approved images.
- **Greater agility** through the ability to quickly provision desktops and deploy new applications.
- **Improved reliability, backup, and recovery of desktop data** by moving it to centrally managed, datacenter-grade infrastructure.
- Reduced application conflicts using a number of methods.
- **Reduced hardware and support costs** associated with providing each user with full-featured desktop or laptop computer, as well as longer hardware refresh cycles.
- **Enhanced productivity** by enabling users to log in to their desktop environment from anywhere with their last session preserved and from a variety of devices and platforms.

Windows Server 2008 R2 SP1 supports VDI deployments with the included Hyper-V hypervisor. To maximize VDI benefits, it is important to address the special requirements of VDI in terms of storage performance, storage efficiency, and manageability. NetApp unified storage provides a number of compelling capabilities that can solve these challenges in a Hyper-V environment:

- **Eliminate performance degradation due to boot storms:** When many virtual machines are booted simultaneously, VDI experiences a spike in storage input/output (I/O). Using NetApp Flash Cache technology, you can maintain VDI performance and reliability during I/O spikes without adding additional spindles.
- **Maximize efficient use of storage:** VDI moves data from relatively cheap desktop storage to relatively expensive datacenter storage. While this provides reliability, security, and manageability benefits, the resulting system must make efficient use of storage resources to minimize costs. NetApp unified storage provides onboard data deduplication to reduce the inherent redundancy of multiple virtual desktops running the same operating system and applications. Additionally, thin provisioning allows storage space to be allocated on a just-in-time basis, eliminating the space wasted by traditional upfront storage allocation.
- **Simplify storage management:** With SnapDrive and SnapManager for Hyper-V, NetApp provides tools that work closely with Hyper-V to reduce the cost and complexity of tasks associated with managing VDI storage systems, such as provisioning, data backup, and data recovery.

Introduction

Virtual desktop infrastructure (VDI) delivers a dedicated virtual desktop to each user in your organization. The operating system, applications, and data exist as virtual machines on servers in a datacenter rather than running directly on the client machine. Compared to deploying and managing individual computers—each with its own hardware, operating system, applications, storage, and so on—VDI can support a more agile and efficient IT infrastructure. Users can access virtual desktops via a traditional desktop computer, a laptop, a "thin client" (an inexpensive computer with low performance specifications that is specifically designed to run a computing environment delivered from a server), and even mobile devices such as tablets and smartphones.

Benefits of VDI include:

- **Centralized desktop management**, enabling IT professionals to maintain a limited number of images on which individual users' environments are based.
- **Enhanced agility** through faster provisioning of new desktops, centralized operating system upgrades, and easier deployment of applications.
- **More robust data security, backup, and recovery** because all data is centralized on datacenter-grade hardware.
- Reduced hardware deployment and support costs through the use of thin clients.
- **Reduced application conflicts** using a number of strategies such as application streaming and application virtualization.
- **Better experiences for remote workers**, as the desktop experience is not tied to a single machine but can be configured for access from multiple devices and locations.
- **Enhanced productivity**, as users can log in to their desktop environment from anywhere with their most recent session preserved.

To maximize the benefits of VDI deployment, it is critical to choose the correct storage system at the outset. This aspect of the project must be considered with utmost care, as changing storage systems after VDI is running is much more expensive and difficult than changing server hardware. VDI deployments of any scale place significant demand on storage systems. In the following section, we discuss key concepts that are critical to understanding the choice of technology for the solution blueprint.

Key Concepts

In this section, we cover in detail the key considerations in choosing the right storage system for an agile, scalable, and manageable Windows Hyper-V VDI environment.

Unified storage

The traditional model for enterprise storage requires a different storage system for each storage function. One storage architecture might be deployed for primary network-attached storage (NAS) and another for storage area networks (SANs), with additional platforms for secondary storage, archive, and compliance. In contrast, a unified storage approach enables the use of multiple storage protocols in one unified system. However, NetApp takes the concept much further with capabilities such as the following:

- **Single provisioning interface**: Centralized allocation, management, and dynamic sizing of data containers across block-level (SAN) or file-level (NAS) access.

- **Common management framework**: A single data model and toolset supports more efficient management across applications and workloads.
- **Policy-based automation**: Delegate some or all of the responsibility for provisioning and management tasks to specialist users and automate common tasks.
- **Shared data protection at the storage level**: A single data protection architecture enables everything from single file restore to full disaster recovery.
- **Incorporation of legacy storage**: Extend unified storage benefits to existing storage systems from a variety of vendors, including the ability to virtualize legacy storage.

These capabilities can make VDI more robust and scalable as well as easier to deploy and manage.

Storage Efficiency

NetApp unified storage provides two key technologies that enable you to reduce the storage footprint of virtual desktops: thin provisioning and data deduplication.

Deduplication

If you run 500 virtual desktops and each requires 20 GB of space, simple arithmetic tells us that you would need 10 TB of storage space. However, assuming that these virtual desktops are running the same operating system and using many of the same installed applications, much of that data is redundant. NetApp deduplication removes redundant copies of blocks within a volume hosting virtual desktops. This process is transparent to the application and user and can be enabled and disabled on the fly. NetApp is the only storage vendor that offers block-level data deduplication for live virtual machines.

Thin provisioning

Traditional provisioning preallocates storage; thin provisioning provides storage on demand. The value of thin-provisioned storage is that storage is treated as a shared resource pool and is consumed only as each individual VM requires it. NetApp thin provisioning allows LUNs that are presented as physical disks or pass-through disks to be provisioned to their total capacity, yet consume only as much storage as is required to store the virtual hard drive files.

Storage performance

Unlike physical desktops in which each machine has its own local disks, hosting virtual desktops on shared storage requires high-performance storage systems. During typical operations, storage I/O varies within a fairly predictable range. However, there are conditions under which many virtual desktops may need to access storage simultaneously, creating a burst of read/write activity, often known as a "boot storm." These situations include antivirus scanning, applying patches and upgrades that cause virtual desktops to reboot, and the start of the workday when many users boot up their virtual desktops at once. Given the disparity between disk size and disk performance today, the traditional method for addressing boot storms is to increase the number of spindles. This approach wastes storage capacity, rack space, and power. Solid state disks have the potential to solve this problem, but they are relatively expensive compared to ordinary disks.

NetApp Flash Cache

NetApp Flash Cache modules give you a new way to optimize performance for random-read intensive workloads such as VDI. Flash Cache is implemented with PCIe-based modules containing 256GB or 512GB of Flash memory per module. Multiple modules may be combined in a single system and presented as a single unit. In addition to the modules themselves, Flash Cache works with software features in Data ONTAP® that enable you to configure and manage Flash Cache for your environment to obtain maximum performance. This technology allows sub-millisecond access to data that would previously have been served from disk at averages of 10 milliseconds or more.

Storage reliability and recovery

With many desktops relying on shared storage, reliability becomes more important, as do data backup and data recovery. Data loss or a system outage can leave scores of users with no access to their most important productivity tools.

Storage manageability

VDI often uses both network attached storage (NAS) and storage area network (SAN) protocols. Choosing a system that allows you to manage both through one interface can create management efficiencies. Additionally, with the need to provision and manage storage for large numbers of virtual desktops, a well-designed set of management and monitoring tools is essential.

Dynamic memory

Hyper-V in Windows Server 2008 and Windows Server 2008 R2 employed a static memory model where the administrator of a Hyper-V host would manually assign physical memory for each virtual machine on the host. The assigned memory allocated to the virtual machine could not change unless the virtual machine was shut down and had a different amount of memory assigned to it. This could limit the VDI capacity of Hyper-V.

To address this issue, Microsoft has introduced a new feature called Dynamic Memory in Service Pack 1 for Windows Server 2008 R2. Dynamic Memory allows the amount of physical memory allocated to a virtual machine to be dynamically adjusted in response to the changing workload needs of the virtual machine. Instead of assigning a specific amount of memory to a virtual machine, the administrator configures a range of memory, memory priority, and other settings that Hyper-V uses to determine how much memory to allocate to the virtual machine in real time.

Solution

Here we discuss in detail the components that make up the VDI solution tested in our lab and the associated benefits of each.

Overview

The solution consists of a VDI implementation based on Windows Server 2008 R2 SP1 Hyper-V accessing a Brocade SAN via NetApp FAS3170 storage devices. VDI is inherently complex and involves many technology components. Here is a summary of key VDI components used in the solution blueprint:

- **Storage network:** The network where data is stored. In this case, we are using NetApp FAS3170 storage controllers to access storage disks using SAN protocol over a Fibre Channel over Ethernet (FCoE) network. The FAS3170 run Data ONTAP. They are configured with Flash Cache for enhanced read/write performance during heavy loads and use onboard data deduplication and thin provisioning.
- **Server network**: This is the network of server hardware that hosts the virtual desktops and associated applications, profiles, and so on. The servers in this case run Windows Server 2008 R2 SP1.
- Hypervisor: This technology separates the virtual machine from the underlying hardware by providing it with an emulated hardware environment. In this case, we are using Windows Server Hyper-V, which is included with Windows Server 2008 R2 SP1.
- VDI solution: Citrix XenDesktop 5 serves as a connection broker and virtual desktop provisioning service in our solution.
- Management and monitoring tools: Running VDI requires various tools to provision virtual desktops and associated storage resources; backup and restore data; manage the storage and server environments; and monitor the environment. This blueprint employs a number of tools for these purposes, including:
  - Microsoft System Center Virtual Machine Manager (SCVMM)
  - Microsoft System Center Operations Manager (SCOM)
  - NetApp SnapManager for Hyper-V
  - NetApp SnapDrive
- Desktop operating systems: We used Login VSI to simulate Windows client workloads. In production, VDI can serve any number of desktop operating systems. Operating systems for VDI virtual machines must typically be licensed as if they were running on individual desktops, although the cost may be lower depending on the vendor.
- Display protocol: This is the service that enables the client to display the virtual desktop. Login VSI uses Microsoft Remote Desktop Protocol (RDP). You can choose from among several protocols. RDP is a mature solution for virtualization of Windows desktops.
- **Network:** The network used to connect the NetApp storage with the Servers was established using Brocade 8000 and Brocade VDX™ data center switches. The Servers utilized Brocade dual-port 10 GbE Converged Network Adapters (CNA).

Use cases

An architecture similar to the one used in this blueprint would be appropriate for a wide range of VDI deployments using Windows Server 2008 R2 SP1 Hyper-V and Windows Desktops. Unified storage from NetApp and the Citrix XenDesktop connection broker are highly flexible technologies, enabling you to customize the basic architecture to use your choice of storage protocols to deliver many types of virtual desktops. In addition to being a capable and flexible hypervisor, Hyper-V has the advantage of being included with Windows Server 2008 R2 SP1, making it a highly cost-effective choice for VDI deployments. All of the major components of the solution—NetApp storage and associated tools, XenDesktop 5, Windows Hyper-V, and Windows management and monitoring software—are designed to work together, simplifying the VDI deployment process.

Solution components

In this section we discuss features and benefits of key components of the solution detailed in the blueprint, with special emphasis on those capabilities that enhance VDI scalability and manageability.

NetApp

NetApp unified storage

NetApp unified storage provides a simpler way to manage multiple storage protocols on the same system while also making it easier to accomplish all the tasks associated with storage management. The unified Data ONTAP operating system includes built-in that can provide substantial reductions in total cost of ownership for VDI implementations.

Features/Capabilities	Benefits
Single operating environment	The stable and mature NetApp Data ONTAP operating environment is the software foundation that underlies every NetApp storage system.
Integrated multiprotocol access	With a single software foundation, it is possible to meet the storage needs of the entire data center with one basic approach, which translates to less training, more staff productivity, and improved business agility.
Single, scalable architecture	NetApp offers a single storage architecture across storage systems that range from workgroup or departmental solutions to full enterprise systems.
Multivendor Array Support	NetApp unified storage enables you to consolidate storage arrays from other vendors and apply technologies such as thin provisioning and data deduplication to legacy systems.
Snapshot copy technology	Space-efficient, scalable Snapshot technology can yield dramatic reductions in backup and recovery time across files, LUNs, databases, and applications.
Disaster recovery	Snapshot supports asynchronous data transfers that replicate only changed blocks (rather than whole files) for efficient use of network bandwidth.

Features/Capabilities	Benefits
Unified data protection architecture	All NetApp data protection software can be configured, managed, and monitored across all storage systems and all storage protocols from a single management console.
Quality of Service (QoS)	Data ONTAP includes QoS capability as a standard feature so you can safely support applications with different priorities on a single, consolidated storage system.
Flexible Storage Tiers	The NetApp unified storage architecture supports provisioning of multiple tiers, including high-performance, low-cost, and archival storage tiers, in a single solution.
Flexible, dynamic provisioning	Data ONTAP enables you to non-disruptively grow and shrink data volumes as needs change.
Thin provisioning	Thin provisioning allows just-in-time provisioning of new storage capacity, resulting in more efficient use of storage.
Consistent skill set across systems	A storage administrator who is trained on Data ONTAP can use those skills across NetApp products.

NetApp Deduplication

NetApp deduplication, an integral feature of NetApp unified storage and Data ONTAP, reduces the storage requirements of VDI with minimum impact on performance.

Features/capabilities	Benefits
Maximize efficiency	Primary data, backup data, and archival data can all be deduplicated with nominal impact on data center operations.
Smart scheduling	You can schedule deduplication to occur during off-peak times so that applications can sustain critical performance but still realize significantly reduced storage capacity requirements.
Simple to use	Once deduplication is enabled and scheduled, no other action is required.
Highly configurable	Select which datasets to deduplicate with our tools to evaluate those datasets and help point out the areas that will provide the greatest return.
Full validation	Perform a full byte-for-byte validation before removing any duplicate data for worry-free deduplication

NetApp Flash Cache

NetApp Flash Cache provides a cost-effective, high-performance solution to the challenges created by bursty read/write activity inherent in VDI deployments.

Features/capabilities	Benefits
Solve common datacenter challenges	Scale your VDI while conserving power, cooling, and space.
Optimize performance without adding storage tiers	NetApp® Flash Cache gives you performance that is comparable to that of solid state disks (SSDs) without the complexity of another storage tier. You don't need to move data from tier to tier to optimize performance and cost.
Automatic caching	Active data automatically flows into Flash Cache because every volume and LUN behind the storage controller is subject to caching.
QoS support	Give caching priority to your most important volumes and LUNs when the load is heaviest by using our FlexShare® quality of service software in combination with Flash Cache cards.
Performance tuning	You can tune Flash Cache to match your specific workload with software settings that let you cache only metadata or cache new data when it is written to disk.

NetApp SnapDrive

NetApp SnapDrive integrates with your native file system and provides a layer of abstraction between application data and physical storage associated with that data. SnapDrive helps host administrators to provision and manage storage directly from the host and resize storage on the fly without any disruption of application service.

Features/Capabilities	Benefits
Increase storage utilization	Create virtual disks from pools of storage that can then be distributed among several storage appliances.
Manage volumes dynamically	Add, delete, map, unmap, and mirror virtual disks online and expand capacity with limited or no impact on application or system performance.
Dynamic pass-through LUN provisioning	Attach and detach of NetApp storage to and from the virtual SCSI controller on the VM without rebooting and without administration on the Hyper-V parent.
Pass-through LUN management	Create and connect LUNs via the Hyper-V virtual machine console. The SnapDrive console shows pass-through information from both the Hyper-V parent and Hyper-V virtual machine.

Features/Capabilities	Benefits
Cluster Shared Volume (CSV) Support	CSV is a new feature in Failover Clustering in Windows Server 2008 R2 that enables access to files on shared storage from all nodes at the same time. CSV in Hyper-V can be managed from SnapDrive.
Increase reliability	Simplify management of cluster resources, virtual disks, and Snapshot copies in your Windows environment, making failover more transparent.
Faster backup and restore	Point-in-time Snapshot images require minimal disk space, cause minimal or no disruption to service, and can be mounted as virtual disks.
Simplify management	SnapDrive makes management simple and intuitive in the Windows environment by allowing administration through the Microsoft Management Console or a command line.
SnapDrive Space Reclamation	Increase storage efficiency by conserving and reusing disk space already allocated to the Windows environment.

Microsoft

Windows Server 2008 R2 SP1 with Hyper-V

Windows Server 2008 R2 SP1 builds on the foundation of Windows Server 2008, enabling organizations to increase the reliability and flexibility of their server infrastructures with expanded technology and new features.

Features/Capabilities	Benefits

Features/Capabilities	**Benefits**
Dynamic Memory	Dynamic Memory supports larger numbers of virtual desktops per Hyper-V host for VDI scenarios and more efficient use of server resources, which can translate into easier management and lower costs. During our lab tests, when VMs were configured using a fixed memory size, host side memory was the limiting factor. When configuring each VM with 1.5 GB of memory the solution was only able to scale to 32 VDI instances per Hyper-V host. Using Dynamic Memory, each Windows 7 VM consumed approximately 750 MB of memory to boot. During the test runs memory consumption did not go over 1GB per VDI instance. As a result each VM was configured to use dynamic memory. A minimum setting of 512 MB was given to each VM and a maximum setting of 1.5GB was used. Using dynamic memory allowed us to scale beyond 32 VMs per host. When using Dynamic memory the limiting factor was host side CPU. In order to keep CPU utilization within acceptable limits the number of VDI instances was limited to 10 per core. By using dynamic memory we achieved a 25 % increase in VM density per host.
Built-in Windows Server 2008 R2 Hyper-V technology	Enables enterprises to easily leverage the benefits of virtualization without adopting a new technology.
Live virtual machine migration	Live migration is available to support business continuity during planned downtime and over distances.
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 types, and devices.
Symmetric Multiprocessor (SMP) support	Hyper-V supports SMP on virtual machines.
Host high availability	Windows Server 2008 R2 clustering provides high availability to virtual machines to minimize unplanned downtime.
Shared storage high availability	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.
Volume Shadow Copy support (VSS)	Provides a robust host-based API to enable virtual machine backup by leveraging the existing Windows VSS-based infrastructure.

Features/Capabilities	Benefits
Easy extensibility	Easy extensibility is available by using the standards-based Windows Management Instrumentation (WMI) interfaces and PowerShell.
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.

System Center Virtual Machine Manager

Microsoft System Center Virtual Machine Manager (SCVMM) enables management of heterogeneous environments, both physical and virtual, through a single interface.

Features/Capabilities	Benefits
Enterprise-class management suite	Manages both Hyper-V and VMware® ESX® virtualization environments.
Intelligent virtual machine placement	Provides support for the intelligent placement of virtual machines.
System Center Operations Manager 2007 R2 integration	Integrates with System Center Operations Manager 2007 to provide proactive management of both virtual and physical environments through a single console.
Performance and Resource Optimization (PRO)	Ties specific alerts from System Center Operations Manager 2007, such as when hardware fails or load thresholds are exceeded, to remediation actions in SCVMM.
Native P2V/V2V migration	Provides native capability for physical-to-virtual migrations and virtual-to-virtual migrations.
Failover integration	Provides integration with failover clustering to support high availability and the live migration of virtual machines.
Automation	Provides easy automation capabilities leveraging Windows PowerShell.
Microsoft System Center Virtual Machine Manager Self Service Portal 2.0	The self-service portal, built on top of Windows Server Hyper-V and System Center, is a free, extensible, turnkey solution that empowers datacenters to dynamically pool, allocate, and manage resources to enable private cloud computing on premises.

System Center Operations Manager

System Center Operations Manager 2007 R2 enables customers to reduce the cost of data center management across server operating systems and hypervisors through a single, familiar and easy to use interface.

Features/Capabilities	Benefits
End-to-End Service Management	Provides end-to-end service management that is easy to customize and extend for improved service levels across your IT environment. This enables operations and IT Management teams to identify, and resolve issues affecting the health of distributed IT services.
"Best of breed" for Windows	Includes expertise from the Microsoft server, client, and application teams, providing you with knowledge and capabilities to drive greater efficiency.
Increased Efficiency and Control	Automates routine, redundant tasks, and provides intelligent reporting and monitoring to help increase efficiency and enable greater control of your IT environment.

Citrix

XenDesktop 5

Citrix XenDesktop 5 offers a powerful and flexible desktop virtualization solution, allowing you to deliver virtual desktops to users anywhere, no matter what device they are using. XenDesktop 5 natively supports Windows Hyper-V and integrates with System Center Virtual Machine Manager through the use of snap-ins.

Features/Capabilities	Benefits
Superior user experience	Users are instantly provisioned with a pristine desktop that incorporates their personal settings and applications, regardless of the user device.
High definition performance and multimedia support	Citrix HDX network and display optimizations and performance boosting technologies deliver an optimal user experience over any network, including low-bandwidth and high-latency WAN connections.
Single image desktop management	Maintaining a single master desktop image in the data center provides users with an up-to-date, pristine desktop at each logon, drastically reduces patch and upgrade maintenance efforts, and cuts storage costs by up to 90 percent.
Built-in virtual applications	Using XenApp with XenDesktop allows you to separate applications from the desktop, resulting in fewer, simpler desktop images.

Features/Capabilities	Benefits
Control over data	Centralized control policies and support for multifactor authentication help ensure that only authorized users can connect to their desktops, while Secure Socket Layers (SSL) technology encrypts the virtual desktop connection for both local and remote users.
Profile management	Profile management provides an easy, reliable, and high performance method to manage user personalization settings in virtualized Windows environments.
Local peripheral support	XenDesktop users can insert a USB device locally and use it with their virtual desktops and applications as they would on a local machine
Active Directory multi-forest support	XenDesktop supports deployment across a range of Active Directory topologies, including multiple domains and multiple forests.
PowerShell scripting	XenDesktop provides an SDK based on PowerShell 2.0 snap-ins that allows you to efficiently perform management tasks including those you cannot perform from the Desktop Studio console such as assigning an IP address to a desktop rather than a user.

Brocade

Brocade Data Center Switches and Adapters

The blueprint leverages Brocade 8000 and Brocade VDX™ Data Center Switches, to provide the foundation for the Ethernet fabric. Brocade Data Center Switches are specifically designed to improve network utilization, maximize application availability, increase scalability, and dramatically simplify network architecture in virtualized datacenters. Installed within all servers are the Brocade 1020 Converged Network Adapters (CNAs). These adapters integrate 10 Gbps Ethernet Network Interface Card (NIC) functionality with Fibre Channel technology enabling transport over a 10 Gigabit Ethernet (GbE) connection through the new Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE) protocols. This provides best-in-class LAN connectivity and I/O consolidation to help reduce cost and complexity in next-generation data center.

Feature/Capability	Benefits
Brocade VCS technology	Simplifies interconnect configuration, provides automatic link failover with no interruption of traffic on unaffected links, and provides plug-and-play fabric scalability.
Ports on Demand	Flexibility to use 16-, 24-, 40-, 50-, and 60-port 10 GbE models built on Brocade sixth-generation fabric switching

	technology.
600 nanosecond port-to-port latency and hardware-based Inter-Switch Link (ISL) Trunking	Maximum performance through wire-speed ports.
Active-active server connections for Ethernet networks	Transparent, automatic failover for business-class reliability.
Scale-out Ethernet fabrics	Supports virtualized data centers.
End-to-end network convergence	Enables utilization of FCoE, iSCSI, and NAS storage along with IP LAN data traffic on one network.
Automatic Migration of Port Profiles	Simplifies management by maintaining network policies as virtual machines move throughout the data center.
Logical Chassis	Helps reduce operational costs and complexity by enabling the entire fabric to be managed as a single switch

Solution Blueprint

Architectural overview

This section details the configuration and tuning that was done to various components for a complete solution validation. Some of the design principles behind the architecture are highlighted below:

- The design includes a scalable and modular building block that can be validated internally.
- Run all the XenDesktop infrastructure components, SCVMM, and other supporting services like AD/DNS/DHCP in a virtual environment.
- The resilient architecture should be able to scale as well as maintain redundancy, maximum bandwidth and robustness. We achieved this using:
  - Dual SAN fabric
  - Multiple links to the network
  - Dual paths to storage
  - Dual storage controller (for maximum redundancy).
- Architecture should be based on established best practices for each partner component involved in the reference architecture.

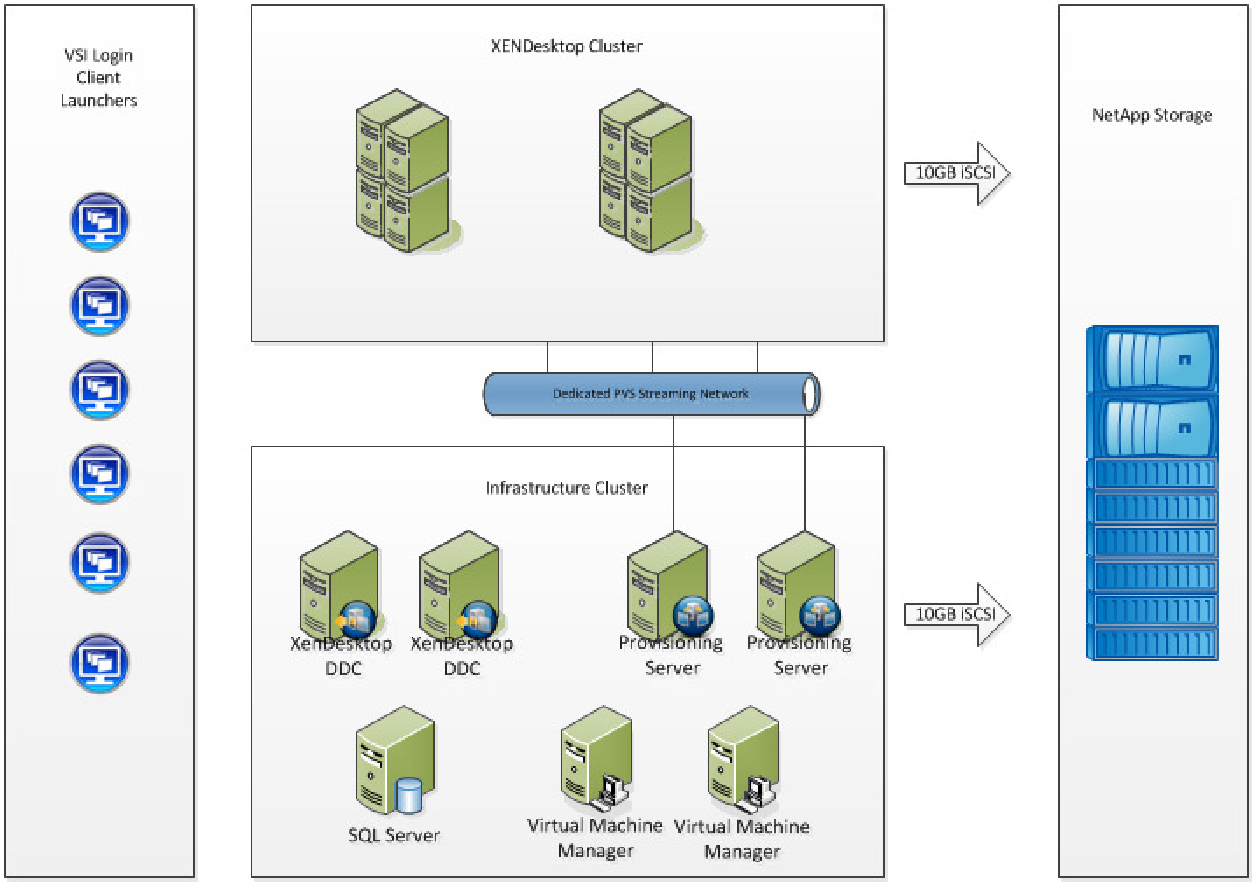

Setup detail

Server Components

All servers relating to the blueprint architecture were virtualized. The blueprint reference architecture used two separate Windows Server 2008 R2 SP1 Clusters. The first cluster included the following servers.

- 2 Desktop Delivery Servers (DDC)
- 2 Microsoft System Center Virtual Machine Managers Servers
- 2 Provisioning Servers (PVS)
- 2 Web interface servers
- 1 Licensing Server

According to NetApp virtualization best practices the infrastructure cluster was built with one CSV per cluster node. For additional information NetApp best practices for deploying Hyper-V with NetApp storage refer to NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V.

Network Components

The network layer was comprised of one Brocade VDX™ Ethernet Fabric switch, and one Brocade 8000 series switch. Each host server used one Brocade 1020 CNA. Each storage controller was connected via fibre channel to the FC ports in the 8000. This was used to SAN boot each of the hosts.

Each host used one port as a SAN boot port booting over FCoE, while the second port was provisioned for the host side networks. The VDX™ switch contained the following VLANS which were presented to each node in the windows cluster.

- Network 1: SCVMM and Hyper-V management traffic
- Network 2: Internal cluster heartbeat
- Network 3: Public - Host/Guest external traffic
- Network 4: Live Migration (cluster communication) and cluster metadata
- Network 5: Dedicated for SMB/Redirected I/O traffic
- Network 6: iSCSI

Storage Configuration and Best Practices

Estimating storage requirements for a Citrix XenDesktop solution on NetApp infrastructure includes the following steps:

- Gather essential solution requirements.
- Perform performance-based and capacity-based storage estimation.
- Plan logical architecture.

NetApp has detailed sizing tools specific to Citrix XenDesktop that can help architect Citrix XenDesktop deployments of any scale. The tools are designed to factor in all the NetApp storage efficiency and performance acceleration components discussed earlier. Contact your local NetApp sales engineer for assistance in storage sizing and architecture for your VDI project.

Gather Essential Solution Requirements

The first step of the storage sizing process is to gather the solution requirements. This is essential to sizing the storage system correctly in terms of the model and the number of required NetApp storage controllers, type and quantity of disk spindles, software features, and general configuration recommendations. The main storage sizing elements are:

- The total number of virtual machines for which the system has to be designed.
- The types and percentage of different types of desktops being deployed. For example, if Citrix XenDesktop is used, different desktop delivery models might require special storage considerations.
- Size per virtual machine (for example, 20GB C: drive, 2GB data disk).
- Virtual machine OS (for example, Windows XP, Windows 7, and so on).
- Worker workload profile (type of applications on the virtual machine, IOPS requirement, read-write ratio, if known).
- Number of years for which the storage growth has to be considered.
- Disaster recovery/business continuance requirements.
- Size of NAS (CIFS) home directories.

NetApp strongly recommends storing user data and roaming profile information CIFS shares. Using NAS home drives, companies can more efficiently manage and protect the user data and eliminate the need to back up the virtual desktops.

There is several storage sizing considerations specific to the deployment of Citrix XenDesktop including:

- Types of desktops that will be deployed for different user profiles.
- Data protection requirements for different data components (OS disk, user data disk, CIFS home directories) for each desktop type being implemented.
- For Citrix provisioning of server-pooled desktops, the write back cache size needs to be calculated based on how often the user reboots the desktop and what applications the user uses. When using a Hyper-V environment, Citrix recommends that a CSV be used for write back cache to conserve space.

This phase of storage evaluation is also an appropriate time to decide whether NetApp thin provisioning, deduplication, and NetApp snapshot should be used to improve storage efficiency and data protection for the user data disk.

Performance-Based and Capacity-Based Storage Estimation Processes

For any VDI implementation, the storage system should be able to meet the performance and capacity requirements of the project and be scalable to account for future growth.

The steps for calculating these storage requirements include:

- Determining the storage sizing building block.
- Performing detailed performance estimation.
- Performing detailed capacity estimation.

Provide the total capacity and performance requirements to your local NetApp technical resource to determine the appropriate storage system configuration.

Plan Logical Architecture

The logical architecture is the total number of templates and associated Flex Clone volumes that should be provisioned per storage aggregate. The recommendation is to provision fewer large aggregates over more, smaller aggregates. The advantage of larger aggregates is that the I/O has more disks to write across, therefore increasing the performance of all volumes contained within the aggregate. Based on the estimated volume size from the capacity calculations section earlier, determine the number of templates and associated Flex Clone volumes that can be hosted in the largest possible aggregate.

- Leave some room to grow the aggregates to handle situations when unexpected growth occurs.
- Disable scheduled aggregate Snapshot copies and set the aggregate snap reserve to zero.

- Make sure the number of disks in the aggregate satisfies the performance requirements for the proposed number of virtual machines for volumes to be hosted in the aggregate.

Blueprint Storage Layout

The storage layout for the blueprint solution used 2 NetApp FAS controllers. A NetApp FAS3170 was used for SAN to SAN boot each host server in the solution as well as house the virtualized environment for the infrastructure as well as the write cache for the provisioning servers as well as CSVs in the infrastructure cluster

Configuration

Storage Configuration

When designing a shared storage environment for XenDesktop, the solution must be highly available. NetApp uses an active-active controller design to provide data availability in business-critical environments, such as XenDesktop virtual environments. Active-active controllers provide simple, automatic, and transparent failover to deliver enterprise-class availability. Providing the highest level of availability of the shared storage is critical because all servers depend on it. For more information, refer to TR-3450: Active-Active Controller Configuration Overview and Best Practice Guidelines.

Aggr0 is only used for the root file system and is typically three drives. To meet the performance and capacity needs of this configuration, each controller has one data aggregate (Aggr1 for hosting production VMs) with the required number of spindles and enough spare disks that can be easily added later to the aggregates to deal with unknowns.

RAID-DP

NetApp RAID-DP® is an advanced RAID technology that provides the default RAID level on all storage systems. RAID-DP protects against the simultaneous loss of two drives in a single RAID group. It is very economical to deploy; the overhead with default RAID groups is only 12.5%. Data residing on RAID-DP is safer than data residing on RAID 5 and more cost effective than RAID 10.

Use RAID-DP, the NetApp high-performance implementation of RAID 6, for better data protection on all RAID groups that store virtual disks for the Hyper-V VMs. Data aggregates should have a RAID group size of no less than 12. A NetApp best practice is to create as large an aggregate as possible.

Aggregate Configuration

An aggregate is the NetApp virtualization layer, which abstracts physical disks on a storage device from logical datasets that are referred to as flexible volumes. Aggregates are the means by which the total IOPS available from all of the individual physical disks are pooled as a resource. The aggregate concept is well suited to meet users' differing security, backup, performance, and data-sharing needs, as well as the most unpredictable and mixed workloads. Aggregates can be 32-bit or 64-bit.

NetApp recommends that, whenever possible, a separate small aggregate with RAID-DP should be used for hosting the root volume. This aggregate stores the files required for running and providing GUI management tools for the NetApp storage system. The remaining storage should be placed into a small number of large aggregates; this provides optimal performance because of the ability of a large number of physical spindles to service I/O requests. On smaller arrays, it might not be practical to have more than a single aggregate because of the restricted number of disk drives on the system. In these cases, it is acceptable to have only a single aggregate.

Each storage controller was configured with two aggregates. AGGR0 was configured at the default size of 3 disks and the root volume for each controller was placed on this aggregate. The first storage controller has an additional 14 disk aggregate. This aggregate was used to house the CIFS share used for roaming user profiles as well as the infrastructure volume which contained the CSV's for each of the infrastructure cluster nodes.

The second storage controller had the AGGR0 (the root AGGR) as well as two additional aggregates. Each additional aggregate contained 14 disks. AGGR1 contained the volumes and LUNS used to SAN boot each of the Hyper-V servers. AGGR2 contained the volume containing the CSV LUNS used for PVS write-cache and VM gold images, as well as persistent desktops.

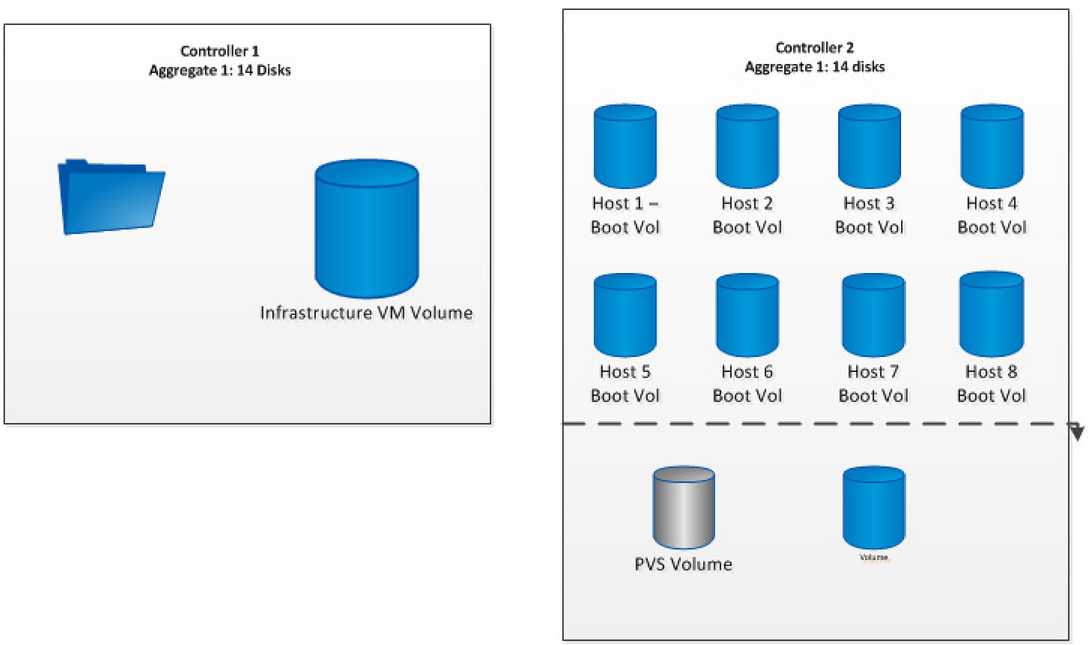

Volume Configuration

Flexible volumes contain LUNs that are accessed by Hyper-V servers over FC or iSCSI. They are virtual volumes that you can manage and move independently from physical storage, and they can be created and resized as your application needs change.

Each volume used in the solution blueprint was configured with the following features

- Thin Provisioned

- Auto-Grow
- Auto-Delete
- Space Reserve set to zero

LUN Configuration

LUNs are units of storage provisioned from a NetApp storage system directly to the Hyper-V servers. The Hyper-V server can access the LUNs using FC or iSCSI protocol as physical disks. The physical disks can serve the following purposes in a Microsoft virtualization environment:

- **Hyper-V parent partition**. In a boot-from-SAN deployment, Windows Server 2008 R2 can be installed on a LUN, provisioned, zoned (single path only during installation, return to MPIO after Hyper-V is installed), and connected to the physical Hyper-V server.
- **Windows failover clustering quorum or witness disk**. In order to use live migration and high availability for Hyper-V VMs, Windows failover clustering (WFC) must be configured. In order to complete the configuration of Windows failover clustering, a witness or quorum disk must be configured. All Hyper-V servers in the Windows failover cluster (WFC) must be able to access this disk; thus, provisioning a LUN from shared storage such as a NetApp storage array meets the requirements.
- **VHD storage**. A physical disk presented to the Hyper-V server can be used to store the virtual hard disk (VHD) files that a virtual machine uses as virtual disks for storage.

Virtual Hard Disks (VHDs)

Virtual Hard Disks (VHDs) allow you to assign storage to a VM while the actual storage is kept in a VHD file located on a physical disk, which is formatted with NTFS and attached to the Hyper-V parent partition. Often these physical disks are LUNs configured from shared storage, through Fibre Channel or iSCSI only. VHDs offer the benefits of increased manageability and portability associated with having the VM storage encapsulated in a single file.

Fixed-Size VHD

This type of VHD allocates the entire configured amount of storage at the time of creation; thus, it does not expand while the VM is operational under any circumstances. Therefore, the fixed-size VHD type offers the lowest performance overhead when compared to the other Hyper-V VHD types.

Dynamically Expanding VHD

This type of VHD doesn’t allocate the entire configured amount of storage at the time of creation; thus, it expands while the VM is operational and data is being written to the virtual disk, but it never expands beyond its configured size. This is done by including a single byte at the end of the file that causes .vhd files to be out of alignment with the disk subsystem. When a file is out of alignment, it causes the disk subsystem to perform additional I/O operations for each file change, degrading performance considerably. Because of alignment and performance issues associated with Dynamically Expanding VHDs, it is recommended that when using NetApp storage VHDs be created using fixed VHDs. The upfront data storage requirements of fixed-size VHDs can be mitigated by many NetApp technologies that help regain storage savings, such as thin provisioning and deduplication.

XEN Desktop and Configuration

The DDCs were virtualized on Hyper-V and some of the roles of the DDC were assigned to specific DDC machines, an approach commonly taken in Citrix XenDesktop and XenApp deployments. In this environment, 2 DDCs (2vCPU, 4GB RAM) easily sustained the farm size of 500 desktops and proved stable during all stages of testing. In order to integrate the DDC with the VMM server, PowerShell and the VMM Administrative console were required to be installed on each controller in the farm.

Citrix Provisioning Services

Citrix Provisioning Server (PVS) is part of the XenDesktop Enterprise and Platinum suites and was used in all tested scenarios, this allows hundreds or thousands of virtual machines hosted on hypervisor servers to PXE boot from and share a single gold Windows 7 Image.

Citrix Provisioning Server (PVS) for Use With Standard Desktops

PVS servers can be configured in a farm to provide high availability and resilience; connections are automatically failed over to a working server/s within the farm in the event of a failure without interruption to the desktop.

The windows desktop image is converted into a vDisk (.vhd) image; this is then locked in a shared (Read-only) mode and hosted on the PVS server's local disk or on a shared file location.

- Virtual desktops are then configured to PXE boot on Hypervisor server.
- PVS streams the vDisk image on start to the Hypervisor, and is loaded into RAM.
- PVS injects a Security Identifier (SID) and host name as each desktop boots to make them unique in AD. These object mappings are maintained and managed within the PVS server and are visible in the PVS Console under "Collections" view are initially created and mapped by the XenDesktop Setup tool.

Each virtual desktop is assigned a Write Cache (temporary file) where any delta changes (writes) to the default image are recorded and is used by the virtual windows operating system throughout its working life cycle. This is where ALL write I/O is conducted for the given virtual desktop instance. It is important, therefore, to consider where the Write Cache is placed when scaling virtual desktops using PVS server. There are several options as to where the Write Cache can be placed:

- PVS Server
- Hypervisor RAM
- Device Local Disk (an additional Virtual Disk for VDI instances)

In this configuration a 3GB virtual disk was assigned to the virtual machine templates used in the virtual desktop creation process. By creating the 3GB drives associated with the templates on FC volumes mounted on the hypervisors, we were able to give each VDI instances its own 3GB drive where the PVS Write Cache was placed. In addition the PVS Target device agent was installed in the Windows 7 image and automatically placed the Windows swap file on the same drive when this mode was enabled. The result is that both the PVS Write Cache and Windows Swap file were hosted on CSV mounted volumes hosted on NetApp storage.

Conclusion

Deploying scalable, reliable VDI requires careful consideration of multiple technology elements. Microsoft Windows Server 2008 R2 SP1 provides a proven virtualization solution with the inclusion of Hyper-V. To provide maximum scalability and reliability, it is critical to choose a storage system capable of performing under heavy loads. NetApp unified storage works with Hyper-V to provide the basis of a robust VDI infrastructure, with innovations that help utilize high-performance storage more efficiently while achieving the business goals of VDI. This blueprint has provided an overview of solution components, architecture, and best practices for deploying such a solution. For assistance in architecting and deploying your VDI solution on NetApp Storage contact your local NetApp and Citrix representative.

Additional information

Citrix XenDesktop 5:
http://www.citrix.com/English/ps2/products/product.asp?contentID=163057&ntref=prod_top

Configuring XenDesktop 5 and Hyper-V: http://support.citrix.com/article/CTX127578

NetApp Storage Best Practices for Microsoft Virtualization:
http://media.netapp.com/documents/tr-3702.pdf

NetApp SnapManager for Hyper-V: http://www.netapp.com/us/products/management-software/snapmanager-hyperv.html

NetApp SnapDrive for Windows: http://www.netapp.com/us/products/management-software/snapdrive-windows.html

Brocade solutions for Microsoft Applications:
http://www.brocade.com/partnerships/technology-alliance-partners/technology-alliances/Microsoft/index.page

Brocade VDX™ switches:
http://www.brocade.com/products/all/switches/index.page?network=CONVERGED

Brocade adapters:
http://www.brocade.com/products/all/adapters/index.page?network=CONVERGED