Best practices for deploying Citrix XenDesktop 4 on Microsoft Windows Server 2008 R2 Hyper-V on HP Converged Infrastructure

Technical white paper

Table of contents

Executive summary

The expectations of today's workforce reach beyond the confines of a traditional office environment. In the past, IT departments typically expected to support desktops located within the corporate network. Today, however, increasingly disparate IT customers are connecting to the corporate infrastructure: office-, mobile-, and tele-workers; contractors, and outsourcing partners all require access and desktop support as if they were under the same roof.

The mobility of today's workforce is challenging IT to provide traditional support functionality to widely dispersed customers using a broad range of client devices. This burden can be eased by deploying a centrally-hosted desktop that is served up via a web browser and is accessible from any geographic location to any client. Using a virtualized desktop infrastructure (VDI) such as the one described in this white paper, you can easily convert your most common desktops to a few virtual images that are deployed to thousands of users, allowing you to move beyond the traditional desktop support model. Now you can easily support and manage virtual desktops (vDesktops) while ensuring that sensitive information remains within the confines of the data center.

HP Converged Infrastructure provides a blueprint for the data center of the future, eliminating costly, rigid IT silos so you can spend much more of your IT budget on business innovation. This infrastructure converges servers, storage, and networks with facilities – all managed through a common management platform. The result is a new level of simplicity, integration, and automation that can accelerate the business outcomes that matter most.

This white paper describes best practices for the deployment of Citrix XenDesktop 4 (referred to as XenDesktop for the remainder of this document) on Microsoft® Windows® Server 2008 R2 Hyper-V (Hyper-V) on HP Converged Infrastructure. For more information on a reference architecture created by HP for this Converged Infrastructure solution, refer to the companion white paper, "HP Converged Infrastructure for client virtualization: Citrix XenDesktop on Microsoft Windows Server 2008 R2."

Target audience: This white paper is intended for solution architects, engineers, and resellers.

This white paper describes testing performed in October 2010.

Deploying XenDesktop on Hyper-V on HP servers, storage, and networking

The solution described in this white paper delivers vDesktops to any location on any client device via a highly-available web browser, allowing you to support the desktop and secure data without leaving the data center.

Overview

The powerful combination of XenDesktop, Citrix Provisioning Services, Citrix Web Interface, and Microsoft Hyper-V allows desktops to be virtualized as virtual disks and delivered to users on any client device. Deploying this VDI solution on HP Converged Infrastructure (as shown in Figure 1) provides an architecture that has been designed for virtualization from the ground up.

Figure 1. Architectural overview of XenDesktop on Hyper-V on HP Converged Infrastructure

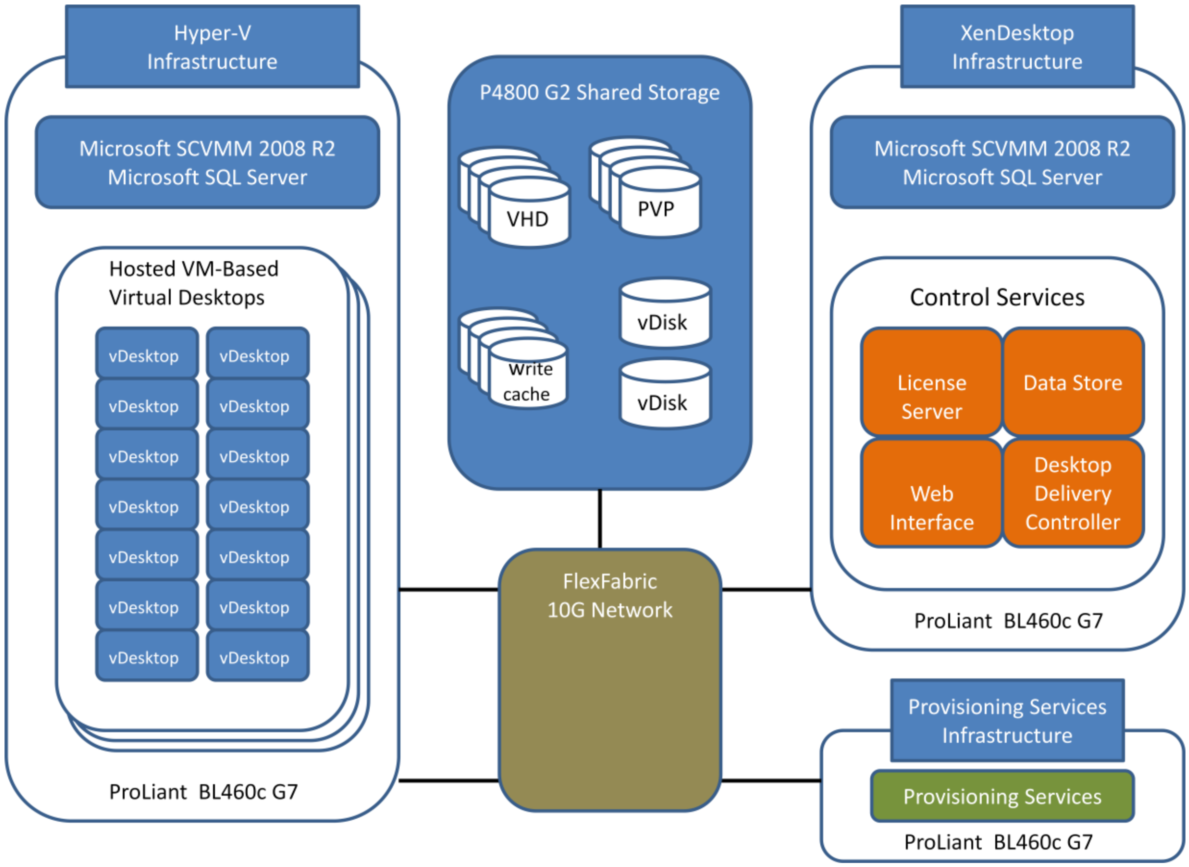

SCVMM = Microsoft System Center Virtual Machine Manager

Solution components

To understand how XenDesktop on Hyper-V on HP Converged Infrastructure works, it is helpful to become familiar with the various solution components, as outlined in this section.

Hardware components

The VDI solution, which only occupies 20U of rack space, includes the following HP hardware:

- HP BladeSystem c7000 enclosure
- HP ProLiant server blades
- HP P4800 G2 SAN
- HP FlexFabric

HP BladeSystem c7000 enclosure

The reference architecture created by HP for a Converged Infrastructure solution[1] has validated that 10 slots of a c7000 enclosure can be used to support 800 vDesktops. In addition, two slots are used to support management servers; and four slots to support storage blades.

[1] HP Converged Infrastructure for client virtualization: Citrix XenDesktop on Microsoft Windows Server 2008 R2

HP ProLiant server blades

The powerful combination of modern HP ProLiant servers and Windows Server 2008 R2 leads to improved consolidation ratios along with a significant reduction in power consumption. HP ProLiant G6 and G7 servers use half the power of earlier generations of servers, support more processor cores, and are certified to support Windows Server 2008 R2 advanced power management.

HP ProLiant BL460c G7 and BL490c G7 server blades make ideal platforms for a XenDesktop environment, delivering features that include:

- Latest Intel® Xeon® processors
- Intel Virtualization Technology (Intel VT)
- Support for CPU core parking, which allows Windows and Hyper-V to consolidate processing on to the least possible number of processor cores and suspend inactive processor cores.
- Boot from virtual hard drive (VHD) support, allowing Windows to boot from a VHD
- Maximum memory configuration of 192 GB
- Wide variety of I/O options
- Innovative HP Virtual Connect technology
- Multiple options for network connectivity
- HP Integrated Lights-Out (iLO) management
- Broad range of options for providing redundancy
- Embedded RAID controller

The Intel Xeon processors deployed in BL460c G7 and BL490c G7 server blades provide features that can optimize Hyper-V execution; for example, Intel VT offers hardware assistance for virtualization. In addition, second-level address translation (SLAT) capability reduces hypervisor processor time and saves about 1 MB of memory per virtual machine (VM) by using the processor to carry out VM memory management functions. Delegating functionality to the processor leads to enhanced scalability and performance.

Running exclusively on 64-bit processors, Windows Server 2008 R2 delivers enhanced support for physical processor and memory resources; for example, up to 256 processors can now be supported on a single server.

This operating system also increases support for virtual processor resources. Hyper-V R2 (hereafter referred to as Hyper-V in this document) now supports up to 64 logical processors (cores or hyper-threaded CPUs). Accepting processor overcommit, Hyper-V R2 allows for 512 virtual CPUs per physical host.

Furthermore, integration between Hyper-V and Microsoft System Center helps to improve availability and simplify management:

- Easily consolidate multiple physical servers onto virtual hosts
- Rapidly provision and optimize new and existing virtual machines
- Performance and Resource Optimization (PRO) enables the dynamic management of virtual resources through management packs that are PRO-enabled.

Note
For information on best practices for implementing Windows Server 2008 R2 Hyper-V on HP ProLiant servers, refer to the HP integration note, "Implementing Microsoft Windows Server 2008 R2 Hyper-V and Microsoft Hyper-V Server 2008 R2 on HP ProLiant servers".

HP P4800 G2 SAN

An integral component of the VDI solution is the HP P4800 G2 63TB SAS BladeSystem SAN, which provides scalable, shared storage for a converged server, storage and networking environment. This SAN can deliver a maximum of 63 TB of storage capacity via four P4000sb storage blades connected to 140 disk drives.

Installed in an HP BladeSystem enclosure, storage blades run proven SAN/iQ software to provide SAN controller functionality, supporting high levels of SAN scalability and data availability.

Storage is provided by two HP 600 Modular Disk System (MDS600) enclosures, delivering over 22 TB of useable space when configured for the reference architecture. This robust storage can accommodate boot storms, login events, and sustained write activity.

Note
Since storage needs in your environment may differ from the reference architecture, HP provides a downloadable sizing tool to help you design your storage infrastructure. For further information, refer to the HP website.

Fully transparent to Hyper-V, the iSCSI-based P4800 G2 SAN provides high availability without the need to buy additional synchronous replication software.

Moreover, Network RAID, which mirrors data blocks across storage nodes in the SAN, can be combined with traditional hardware RAID to extend fault tolerance to new levels. With Network RAID, you can create up to four copies of the data. This level of protection means that you can configure a P4800 G2 SAN to survive the failure of one or more drives, one or more network interfaces, and one or more storage nodes. HP refers to this capability as hyper-redundant clustered storage.

A P4800 G2 SAN includes the following:

- **SAN/iQ software** (see also SAN/iQ):
  - SAN/iQ Storage Clustering
  - SAN/iQ Network RAID
  - SAN/iQ Thin Provisioning
  - SAN/iQ Snapshots
  - SAN/iQ Remote Copy (asynchronous replication)
  - SAN/iQ Multi-Site/Disaster Recovery Solution (synchronous replication)
  - SAN/iQ Solution Pack for Microsoft Windows

- **SAN hardware**:
  - Four redundant active-active scalable P4000sb storage blades for BladeSystem c-Class enclosure
  - Two MDS600 enclosures with 140 450 GB 15,000 rpm SAS disk drives
  - Two 3 Gb 8-port SAS 2.0 BL-c switch modules for BladeSystem c-Class enclosure
  - Eight 2m mini SAS cables
  - 24 GB RAM
  - 2,048 MB battery-backed cache
  - Support for hardware RAID 5, 6, and 10
  - Support for network RAID 0, 5, 6, 10, 10+1, 10+2 (synchronous replication)
  - Eight 10 Gb network interface cards (NICs)
  - iLO 2 management

The SAN can be managed via the P4000 Centralized Management Console (CMC). (See CMC.)

For more information, refer to Appendix A: Installing the P4800 G2 SAN.

HP Flex-10

The Virtual Connect Flex-10 Ethernet Module simplifies server connections by cleanly separating the server enclosure from LAN, and simplifies networks by reducing cables without adding switches to manage. This blade interconnect allows you to change servers in just minutes, and fine-tune network speeds based on application needs. With Virtual Connect, you can reduce costs and simplify connections to LANs and SANs, consolidate and precisely control their network connections and enable administrators to add, replace and recover server resources on-the-fly.

Flex-10 technology is an HP hardware-based solution that enables you to partition a 10 GbE connection and manage data speed in each partition. Flex-10 allows Virtual Connect to re-configure a single 10 GbE network port in an HP ProLiant server blade, creating four physical network interface controllers (NICs) – also known as FlexNICs – with a total bandwidth of 10Gbps, as shown in Appendix B: Designing FlexNICs.

The bandwidth of each connection can be fine-tuned in 100 Mb increments up to the 10 Gb limit to meet the needs of a changing workload.

Note
HP ProLiant servers can support up to 40 network connections with significantly less investment in expensive network equipment on the server, in the enclosure, and in the corporate network compared to a conventional implementation.

For more information, refer to Appendix B: Designing FlexNICs.

Software components

The VDI solution includes the following software components:

- Hyper-V
- SAN/iQ
- CMC
- Microsoft System Center Virtual Machine Manager R2 (SCVMM)
- HP Insight Control for Microsoft Systems Center

- Citrix Desktop Delivery Controller (DDC)
- Citrix Provisioning Services (PVS)

The VDI solution uses SCVMM, Hyper-V, DDC, and PVS to capture a reference image and deploy this image to multiple vDesktops[2].

Hyper-V

First released with Windows Server 2008, Hyper-V is a hypervisor-based platform for x64 servers that supports the virtualization of guest operating systems.

Windows Server 2008 R2 (Standard, Enterprise, and Datacenter Editions) introduced new Hyper-V features. For example, live migration allows you to migrate running VMs from one physical host to another and add storage to or remove storage from a VM while it is running.

Note
Hyper-V only supports direct-attached storage (DAS) or block-level SAN storage such as the iSCSI storage provided by the P4800 G2 SAN.

Windows Server 2008 R2 Hyper-V provides enhancements for the dynamic virtual data center in the following areas:

- Availability
- Management
- VM performance
- Hardware support
- Improved virtual networking performance
- Simplified deployment of physical servers and VMs using .vhd files

Microsoft and Citrix provide resources that can help you deploy a VDI solution. For example, Citrix has developed a guide, " XenDesktop Design Guide For Microsoft Windows 2008 R2 Hyper-V," that offers best practices for deployment[3]; in addition, Microsoft offers a no-cost Virtual Lab that walks you through a vDesktop deployment.

Versions

The following versions of Hyper-V are available:

- **Standard installation** – Available with any release of Windows Server 2008 R2
- **Server core** – Special installation available with any release of Windows Server 2008 R2
- **Microsoft Hyper-V Server 2008 R2** – No-cost version of Windows Server 2008 R2; limited to Hyper-V and clustering components

Note
HP used Windows Server 2008 R2 as the basis for Hyper-V deployments in this VDI solution.

[2] For further information, refer to the Microsoft website.
[3] For further information, refer to the Citrix website.

SAN/iQ

At the core of the P4800 G2 SAN is SAN/iQ storage software, which provides SAN management features such as clustering, application-integrated snapshots, thin provisioning and remote copying (asynchronous replication). SAN/iQ supports standard, open iSCSI protocols and is compatible with almost any operating system that provides an iSCSI initiator.

SAN/iQ also includes unique Network RAID capability, which protects against disk, controller, storage node, power, network, or site failure while keeping data volumes accessible and applications online.

New with SAN/iQ 9.0

Backward-compatible with SANs shipped in 2007 and later, SAN/iQ 9.0 introduces capabilities such as the following:

- **Storage management cost control**

  Global configuration policies and automated online upgrades, along with HP Insight Remote Support Software, and HP Systems Insight Manager (HP SIM) make managing a very large SAN with 30 nodes as easy as managing a small SAN with just two nodes.
- **Storage availability and utilization**

  SAN/iQ can simultaneously maintain high data availability and improve capacity utilization. Both Network RAID 5 and 6 offer improved capacity utilization (up to 33%) while protecting against single and multiple disk failures.
- **Extending the benefits of HP BladeSystem to storage**

  SAN/iQ helps extend the advantages of an HP BladeSystem infrastructure to storage, as in the following examples:
  - HP Insight Software can now be used for all server, network, power/cooling, and storage hardware management.
  - Day-to-day management of the HP BladeSystem infrastructure and SAN can be performed using just two tools: BladeSystem Onboard Administrator and CMC.
  - The power-throttling and capping that help make an HP BladeSystem infrastructure so efficient have been extended to the SAN.

CMC

The CMC provides a user interface for SAN/iQ software, allowing you to provision and manage the P4800 G2 SAN from a single interface.

CMC capabilities include:

- Configure P4800 G2 SAN clusters
- Provision the clusters into LUNs
- Assign LUNs to the servers requiring access

SCVMM

SCVMM allows you to manage Hyper-V hosts via administrator consoles installed on the DDC and PVS servers.

SCVMM helps you centralize the management of physical and virtual IT infrastructures, increase server utilization, and dynamically optimize resources across Hyper-V and other virtualization platforms. It includes end-to-end capabilities such as the planning, deployment, management, and optimization of a virtual infrastructure.

HP Insight Control for Microsoft Systems Center

HP Insight Control for Microsoft System Center provides seamless integration of the unique ProLiant and BladeSystem manageability features into the Microsoft System Center consoles. Licensed as part of HP Insight Control, these extensions deliver comprehensive system health and alerting, configuration management, remote control, and proactive virtual machine management in a flexible, easy to install package.

DDC

DDC provides desktop brokering functionality for this VDI solution. In essence, DDC maps users to an assigned desktop group and then assembles an appropriate vDesktop based on the access assigned to these users in Active Directory. The desktop group publishes the vDesktop to users.

Desktop brokering services include the following:

- Validate the user and provide an end-to-end Citrix Independent Computing Architecture (ICA) connection
- Interface with Active Directory to identify user access rights and desktop group assignments
- Dynamically assemble vDesktops, assigning the appropriate virtual disk, memory, and CPU resources, etc.
- Power vDesktops up and down
- Monitor vDesktop availability
- Manage the reassignment of a vDesktop when the user disconnects or logs off

Note
After being brokered by DDC, there is a direct connection between the user's device and the hosted vDesktop. Established connections are not affected if DDC is re-booted.

See also Using DDC to mitigate boot storms.

PVS

PVS allows many users to share a single desktop image by streaming the operating system to the vDesktops in the form of a vDisk.

Note
You can build and assign as many vDisks as you need to meet the requirements of your user community. For example, one user group may need a Windows 7 vDisk while another needs a Windows XP vDisk.
With the powerful combination of DDC and PVS, the desired vDesktop can be dynamically assembled in minutes.

The following scenarios demonstrate the versatility of PVS.

Without PVS

Without the help of PVS when creating a vDesktop, you would have to individually manage and maintain each virtual machine in your virtual desktop environment.

Using PVS enables you to manage each aspect of the vDesktop using a single common image which can be rolled back in seconds in the event of an issue with a deployed image.

With PVS

With PVS, you do not need to mount a disk; instead you use XenDesktop Setup Wizard to specify which vDisk should be used. At boot time, PVS streams the vDisk to the vDesktop over the network; thus, a single vDisk can be accessed by thousands of VMs.

In this scenario, the user's configuration could rapidly be re-provisioned with another operating system; patches or programs can easily be backed off or rolled back.

Note
Since drivers are specific to the particular hypervisor, a vDisk is specific to the hypervisor on which it was created. Thus, a vDisk created on Hyper-V cannot be used on Citrix XenServer.

Having outlined the components of this VDI solution, this white paper now describes best practices for designing and deploying the solution.

See also Optimizing the PVS configuration.

Best practices for hardware deployment

This section suggests best practices for deploying the hardware components of the VDI solution that were described earlier.

HP ProLiant server blades

This section details best practices for configuring HP ProLiant server blades to optimize scalability.

Updating BIOS settings

Server scalability can be impacted by several BIOS settings, which may be accessed via the blade's ROM-based Setup Utility (RBSU). This section describes the settings and recommendations.

Using HP Power Regulator for ProLiant

Power Regulator is an innovative operating system-independent power management feature that can be used to manage server power consumption and system performance[4].

You can configure Power Regulator for the following modes of operation:

- HP Static High Performance
- HP Static Low Power
- HP Dynamic Power Savings
- OS Control

For optimum scalability in a XenDesktop environment, select **HP Static High Performance Mode**[5].

Enabling virtualization support

Use the RBSU to enable hardware-assisted virtualization and no-execute (NX).

Important
You must enable hardware-assisted virtualization before installing Hyper-V.

Enable Hyper-Threading

Intel Hyper-Threading Technology (Intel HT Technology) can deliver significant performance gains. Although HP has not yet characterized these gains in a VDI environment, Figure 2 summarizes the results of testing performed in a Citrix XenApp environment, which indicates that an HP ProLiant DL380 G7 server blade was able to support 36% more users after Intel HT Technology had been enabled[6].

[4] For further information, refer to the HP website.
[5] This setting can also be selected via HP BladeSystem Onboard Administrator.
[6] For further information, refer to the HP ActiveAnswers website: http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=4AA2-1450ENW&cc=us&lc=en.

Figure 2. Showing the performance impact of enabling Intel HT Technology in a XenApp server[7]

Optimizing DDR3 memory performance

HP ProLiant BL460c G7 and BL490c G7 server blades – the servers recommended for the VDI solution – support third-generation DDR memory, DDR3[8].

These blades use non-uniform memory architecture (NUMA), whereby each processor has its own memory controller and can directly control three memory channels; no intermediate transport mechanism is required.

Note
The memory provided for a particular processor is still part of the overall system memory. Other processors can access this memory using Intel QuickPath Interconnect links configured between processors and between processors and the chipset.

Selecting the appropriate DIMMs

BL460c G7 and BL490c G7 server blades support Registered DIMMs (RDIMMs) and Unbuffered DIMMs (UDIMMs); however, RDIMMs and UDIMMs cannot be used together in the same system.

Table 1 compares the maximum memory capacity of these blades.

Table 1. Comparing capacities

Capacity	BL460c G7	BL490c G7
DIMM slots	12	18
Memory channels (per processor)	3	3
DIMMs per channel	2	RDIMM: 3 UDIMM: 2
Maximum capacity	RDIMM: 384 GB (PC3-10600 at 800 MHz) UDIMM: 48 GB (PC3-10600 at 1,333 MHz)	RDIMM: 192 GB (PC3-10600 at 800 MHz) UDIMM: 48 GB (PC3-10600 at 1,066 MHz)

[7] x/y/z denotes a configuration with x VMs, each featuring y virtual CPUs and z GB of RAM.
[8] For further information, refer to the HP website.

BL460c G7 and BL490c G7 server blades ship standard with RDIMMs, the DIMM type needed to achieve maximum memory capacity. Thus, you should consider using RDIMMs when you implement a VDI solution.

HP simplifies memory selection by providing the online DDR3 Memory Configuration Tool, which can help you configure server memory and provide an orderable parts list.

Populating memory to optimize bandwidth

The key to optimizing memory performance is to populate as many channels as possible, thus maximizing the bandwidth available to the system. For example, you should install a minimum of six DIMMs (one per channel) in a 2P BL460c G7 blade, populating DIMM slots 2, 4, and 6 for each processor.

Figure 3 presents a simplified view of fully–populated memory in a BL460c G7 blade, with a total of six memory channels and 12 DIMMs.

Figure 3. Fully-populated memory configuration in a 2P BL460c G7 blade

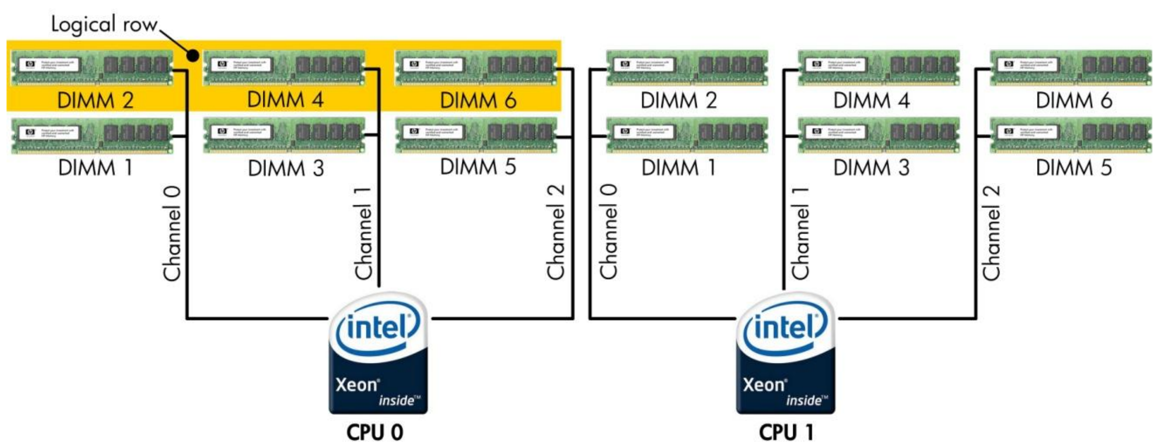

If all these DIMMs are identical, the processors' memory controllers can use channel interleaving to map consecutive eight-byte portions of the system's logical memory map to different physical memory channels. This increases the probability that memory accesses will be spread more evenly across the memory channels, allowing the potential composite bandwidth to be achieved.

Populating memory in logical rows and in groups of either three or six DIMMs helps maximize system throughput.

Balancing memory

For most applications, the optimal configuration for DDR3 memory is to balance memory both across memory channels and processors. This scenario optimizes channel and rank interleaving, ensuring maximum memory throughput.

Balancing the installed memory across processors ensures consistent performance for all threads running on the server. If more memory were installed on one processor than the other, threads running on the processor with more memory would perform better, creating a performance imbalance that could degrade overall system performance, particularly in a virtualized environment.

Best practices

When configuring memory in Intel-based HP ProLiant G6 and G7 servers, you should observe the following guidelines:

- If possible, balance DIMMs across a processor's memory channels (for example, use three identical DIMMs in the same logical row)
- When installing memory in a 2P system, balance memory between CPUs
- Mixing DIMM models within a memory channel should not adversely affect throughput; however, you should not mix DIMM models in a logical row since this would disrupt channel interleaving and, potentially, overall system performance

P4800 G2 SAN storage

This section suggests best practices for deploying the VDI solution's storage environment.

Note
While any P4000 G2 SAN solution can be used in a VDI environment, HP recommends that you do not cluster P4800 G2 SANs with other models. In such a scenario, the capacity of the smallest SAN would automatically become the maximum capacity available from any other SAN in the cluster.

Maintaining the relationship between storage blades and storage

P4000sb storage blade placement within the HP BladeSystem c7000 enclosure is critical. Figure 4 shows the relationship between storage blades (A, B, C, and D) and their MDS600 storage drawers.

Figure 4. Showing the relationship between storage blades (in the c7000 enclosure) and storage drawers

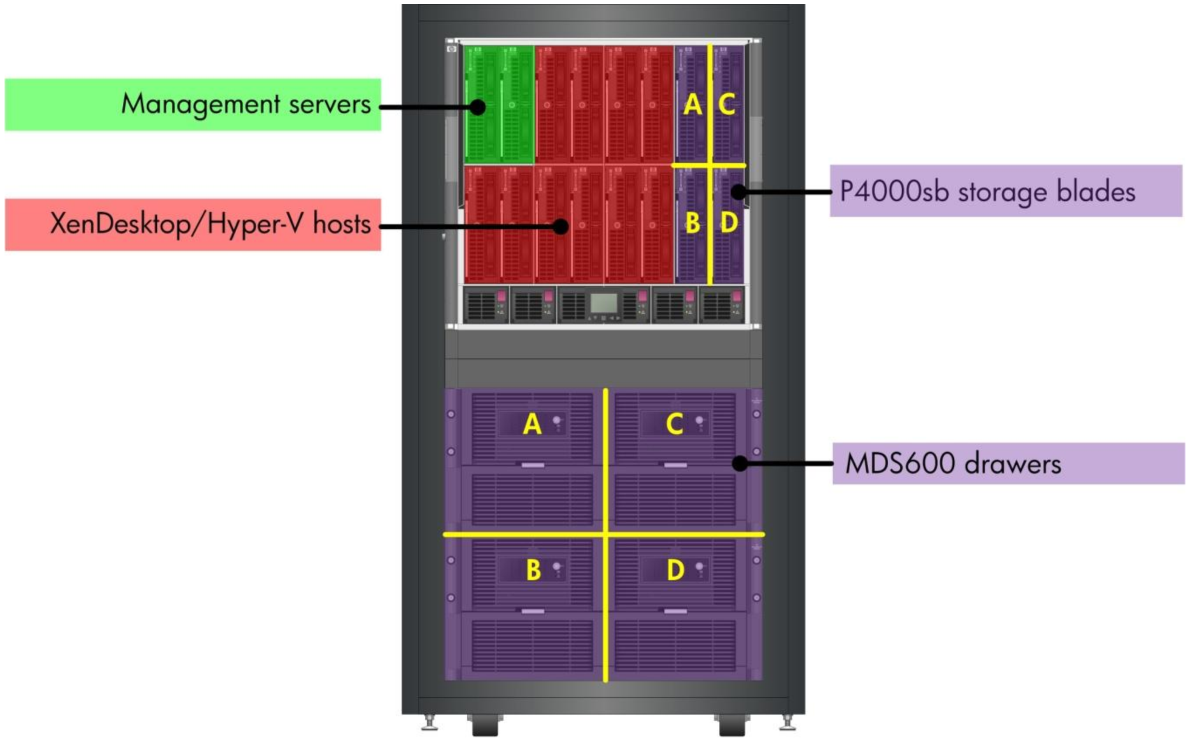

Each storage blade features an HP Smart Array P700m controller that is redundantly connected to its associated disk drawer via SAS switches. Each drawer contains 35 disks that are subdivided into seven groups of five; protection is provided at the controller level via hardware RAID 5.

SAN/iQ software creates a user-selectable Network RAID configuration that encompasses all four drawers, with each written block being replicated two, three, or four times to achieve the desired levels of data integrity and availability.

Setting up the P4800 G2 SAN

Table 2 details the capacity and effective yield of the four-node P4800 G2 SAN.

Table 2. P4800 G2 SAN capacity and yield

Capacity	Yield
Raw capacity	
Total	63.00 TB
Total per node	15.75 TB
Overhead per node	
Disk formatting	0.78 TB
7 x RAID 5 set (28 active / 7 parity disks)	3.00 TB
SAN/iQ overhead (5%)	0.72 TB
Total	4.50 TB
Usable capacity	
Total per node	11.25 TB
Total per SAN	45.00 TB
Total per SAN after Network RAID 10	22.50 TB

Prior to installing Hyper-V and the Citrix solution components, you should perform the following steps to configure the P4800 G2 SAN and install iSCSI software on Hyper-V servers to provide SAN connectivity.

Prerequisites

- Obtain the latest SAN/iQ software (which includes the latest SAN/iQ firmware) and download the P4000 SAN Solution User Guide.
- Install the P4800 G2 SAN on your network.
- Record the IP address or host name you configured when you installed the SAN.
- Install P4000 Central Management Console (CMC) software on a system that can connect to the SAN over the network.
- Install the latest firmware.

Configure the LUNs on the P4800

Use the CMC to configure the following LUNs on the P4800 for the XenDesktop setup:

- PVS LUN, which hosts the vDisk's virtual hard drive (**.vhd**) and PVS virtual properties (**.pvp**) files
- LUN for each Hyper-V server to host the vDisk configuration, PVS write cache, and the saved-state file

Grant servers appropriate access to these LUNs; for example, grant the PVS server access to the PVS LUN.

For more information, refer to the P4000 SAN Solution Guide.

Installing iSCSI software

Install the following on Hyper-V servers:

- An iSCSI initiator, such as the latest version of Microsoft iSCSI Initiator
- HP P4000 Device Specific Manager (DSM) for Multipath I/O (MPIO)

Important
P4000 DSM for MPIO is vendor-specific software that uses the Microsoft MPIO framework to enhance multi-pathing in a P4000 SAN. This software supports I/O path redundancy, failover, and load balancing; for example, it can be used to create fault-tolerant paths to storage nodes, while increasing available bandwidth.
P4000 DSM for MPIO software can be found in the HP P4000 Windows Solution Pack.

Best practices for deploying Windows Server 2008 R2 and Hyper-V

This section suggests best practices for deploying Windows Server 2008 R2 and Hyper-V in the VDI environment.

Begin by ensuring that the particular server's BIOS has been updated to the latest version; use the server's RBSU to enable support for No-Execute and hardware-assisted virtualization. Now you can carry out the installation in the following order:

1. Install Windows Server 2008 R2.

Note
Windows Server 2008 R2 should be installed on all servers in this VDI environment, except the Citrix XenDesktop 4 DDC server, which is only supported on Windows Server 2003.

2. Install HP ProLiant Support Pack (PSP) 8.30 or later – without Network Configuration Utility (NCU).

Note
The NCU should be installed after the Hyper-V role on the Windows 2008 R2 server is enabled.

Note
Do not install PSP on the vDesktops.

3. Enable the Hyper-V role.
4. Install the NCU.

Installing management software

When you are installing HP Systems Insight Manager (HP SIM) or HP Insight Control for System Center, make sure that Simple Network Management Protocol (SNMP) has been configured and started.

Install HP Insight Control for Microsoft System Center.

Enabling NIC teaming

After enabling Hyper-V, install and enable HP ProLiant Network Teaming Software.

Best practices for deploying the Citrix components

This section suggests best practices for deploying the Citrix components required by the VDI environment.

Installing XenDesktop

The following steps have been adapted from Citrix recommendations for a new installation of XenDesktop[9] components. HP suggests performing these steps in the order specified:

1. Install a license server and obtain appropriate licensing.
2. Create the data store for the farm.
3. Install the SCVMM server.
4. Install the SCVMM console on servers that will host DDC and PVS functionality.
5. Install DDC on a single server and create a farm.
6. Configure Active Directory.
7. Add DDC controllers to your farm, if necessary.
8. To enable the remote management of this solution, install management consoles on the appropriate servers.
9. Start Citrix Access Management Console and run the discovery function.
10. Install PVS, noting the following caveats:
    - Although the PVS server can be virtualized, the networking load created by PVS may necessitate the use of a physical server.
    - Due to its ability to provide more system cache allocations and, thus, provide faster vDisk access, a 64-bit edition of Windows Server 2008 will provide the best performance for PVS.
11. Install XenDesktop Setup Wizard and Virtual Desktop Agent on the PVS server.

Note
You cannot mix XenDesktop and XenApp farms.

Using DDC to mitigate boot storms

The largest load in a XenDesktop environment – known as a **boot storm** – is incurred during vDesktop boot. This process not only presents the largest load in terms of disk, processor, and network utilization but is also exacerbated by the large numbers of vDesktops booting at the same time.

[9] For further information, refer to the Citrix website.

You can mitigate the effects of boot storms by using DDC to automatically power up the vDesktops prior to the scheduled user start time (also known as pre-staging the vDesktops).

For a smaller environment featuring a single DDC, the **Idle-Pool Settings** from the Citrix Delivery Services Console can be used to stagger logons for each Desktop Group. These settings control the number of idle vDesktops that are powered on and awaiting user access, allowing you to schedule DDC to automatically power up desktops to fit particular user logon patterns. By using DDC to automate this process, you can ensure that adequate resources are available to accommodate user requirements, while avoiding the resource congestion caused by desktop boot storms.

For a larger environment featuring multiple DDCs, it may be necessary to throttle the number of VMs that can be powered on at any one time[10]. To modify the maximum number of concurrent requests allowed, edit the configuration of each DDC as shown in the following example, where the maximum number of concurrent commands is being set to 20.

Setting the maximum number of concurrent commands

Note
The maximum number of concurrent commands varies depending on the particular environment. The maximum number for VMware without Distributed Resource Scheduling (DRS), XenServer, or Hyper-V is 40. The maximum number for VMware with DRS is 20. Citrix recommends testing your environment to obtain the best results.

In each DDC server, set the maximum number of concurrent commands.

Open the file, **C:\Program Files\Citrix\VmManagement\CdsPoolMgr.exe.config**. In the **appSettings** section of this file, add the line shown below in bold text:

```
<appSettings>
<add key="LogToCdf" value ="1"/>
<add key="LogFileName" value ="C:\cdslogs\VMManager.log"/>
<add key="LogDebug" value="1"/>
<add key="MaximumTransitionRate" value="20"/>
</appSettings>
</configuration>
```

[10] For further information, refer to the Citrix website.

Best practices for designing the storage infrastructure

Correctly designing storage in a VDI environment is imperative to ensure a successful implementation. It is important to size not only for capacity but also for performance. This section provides guidelines for optimizing I/O performance in a VDI environment.

XenDesktop on Hyper-V storage requirements

In a XenDesktop on Hyper-V environment, PVS assembles the vDesktop from multiple areas of the storage infrastructure; thus, it is important to understand the requirements of each area when you are planning the infrastructure. You should also keep future growth in mind so that your solution can scale as your storage needs grow.

Storage requirements in a XenDesktop on Hyper-V VDI environment tend to focus on capacity and performance, as outlined below.

Storage capacity

It is important to understand the capacity requirements of your environment. While PVS greatly reduces the need for virtual desktop storage, users still need storage capacity for profiles, home drives, and any additional disk space that may be required.

Characterizing performance

For a suitable storage design, it is critical to characterize the workload of each vDesktop disk in Input/Output Operations per Second (IOPS).

Note
Each disk read or write constitutes a disk I/O operation.

The following factors can influence IOPS:

- Spindle speed – 10,000 rpm or 15,000 rpm
- Disk technology – SAS or SATA
- Hardware RAID configuration (RAID 0, 1, 5)
- Storage cache – memory is faster to access than disk
- Workload profile – percentages of read and write operations
- Client applications

Calculating the expected IOPS

In order to properly design the storage infrastructure, you must first be able to calculate the expected IOPS requirements.

IOPS calculations performed by Citrix Consulting Services (described below) utilize the parameters presented in Table 3[11].

Table 3. Parameters used in IOPS calculations

Parameter	Description	Values	
Raw IOPS (for SAS disk)	The speed at which the disk spins has a direct impact on how quickly the disk can read a particular sector.	**Disk speed**	
		15,000 rpm:	150 IOPS
		10,000 rpm:	110 IOPS
		5,400 rpm:	50 IOPS
Read/write ratio	I/Os are broken down into reads and writes. Some applications are read-intensive; others require more writes. The read/write ratio impacts the overall IOPS.	Most desktop implementations result in a read/write ratio of 1:4 (20%/80%).	
RAID penalty	The RAID configuration impacts how many write IOPS are available. The resulting write penalty reduces the overall IOPS for each disk spindle.	**RAID penalty**	
		RAID 0:	0 IOPS
		RAID 1:	2 IOPS
		RAID 10:	2 IOPS
		RAID 5 (four disks):	4 IOPS
		RAID 5 (five disks):	5 IOPS
IOPS requirements over the desktop lifecycle	Each desktop lifecycle includes six phases, each of which places a different load on the storage subsystem.	**Desktop lifecycle**	
		Boot:	26 IOPS
		Logon:	14 IOPS
		Work:	8 IOPS
		Idle:	4 IOPS
		Logoff:	12 IOPS
		Offline:	0 IOPS

Disk speed, read/write ratio, and RAID penalty values are included in the following equations, which can be used to calculate the **functional IOPS** value (which includes the RAID penalty) for a particular disk subsystem:

$$Total\ raw\ IOPS = Disk\ speed\ IOPS\ x\ Number\ of\ disks$$

$$Functional\ IOPS = \left(\frac{Total\ raw\ IOPS\ x\ Writes}{RAID\ penalty}\right) + (\ Total\ raw\ IOPS\ x\ Reads\ )$$

In the functional IOPS calculation, Writes and Reads are expressed as fractions of 1. Thus, for example, with a read/write ratio of 1:4, Writes would be expressed as 0.8; Reads as 0.2.

Using these formulas, you can calculate IOPS requirements on a server-by-server basis and for the entire infrastructure.

Example

In the following example, eight 15,000 rpm SAS drives are configured with RAID 10.

$$Total\ raw\ IOPS = 150\ x\ 8 = 1200$$

$$Functional\ IOPS = \left(\frac{1200\ x\ 0.8}{2}\right) + (\ 1200\ x\ 0.2\ ) = 720$$

[11] For further information, refer to the Citrix website.

Having obtained the functional IOPS for the disk subsystem, you can use the following formula to determine the number of desktops that can be supported during a particular phase of the desktop lifecycle.

$$Supported\ desktops = \left(\frac{Functional\ IOPS}{Desktop\ lifecycle\ IOPS} \right)$$

Thus, the maximum number of desktops that can be logged on simultaneously is 720 ÷ 14, which is approximately 51.

Additional considerations

Table 4 outlines additional recommendations for configuring storage.

Table 4. Recommendations for configuring storage

Component	Recommendation	Comments
Write cache IOPS	14 IOPS	Reserving 14 IOPS for the write cache provides the capacity needed to support intense logon storms, which cannot be controlled by other means. While boot storms have a larger performance impact, they can be managed via the XenDesktop configuration.
vDisk RAID	RAID 5	The vDisk only experiences reads. By optimizing the RAID configuration for reads (via RAID 5), vDisk blocks are delivered faster.
Disk type	Fixed	Dynamic disks present the following performance challenges, which are overcome by using fixed disks: • Fragmentation as the disks continue to expand • Simultaneous expansion of hundreds of disks during vDesktop startup • Misalignment of the storage, which results in each I/O operation in the disk requiring two I/O operations in the storage infrastructure as blocks within the disk cross block boundaries in the storage infrastructure For more information, see Fixed or dynamic vDisks?

IOPS in a VDI environment

Based on industry estimates for a VDI environment, HP selected an average of 20 IOPS per vDesktop when designing the storage configuration for the VDI solution. Thus, the maximum projected workload is 800 vDesktops x 20 IOPS per vDesktop = 16,000 IOPS.

Disk caching

It is always faster to access memory somewhere along the data path rather than accessing physical disk media; thus, caching can dramatically affect IOPS.

Note
VDI performance characterizations obtained using test harnesses may be skewed because automated harnesses tend to re-use information that may still be available in cache memory.

To enhance performance, HP provides 2GB battery-backed write cache (BBWC) within the P4800 G2 SAN.

In addition, SAN/iQ 8.5 introduced the **read cache enhancement** feature, which can speed up the processing of a read request with a sequential workload. Since a volume is striped across all storage nodes in a cluster, the system sends the adjacent node a data path request for read-ahead (known as a **hint**).

LUN provisioning for VDI

Since the P4800 G2 SAN supports multiple functions in a VDI infrastructure, a number of LUNs must be created and maintained. These LUNs perform different functions and should be optimized based on their particular usage.

For example, a PVS vDisk is read-only in a PVS standard-mode environment. Conversely, you should optimize the LUN hosting the PVS write-cache for disk writes, since all the vDesktop OS information is written to this LUN.

HP suggests the following best practices:

- **PVS LUN** (hosting vDisks)

  vDisks include the .vhd files used to store vDisk contents and the .pvp files used to store vDisk properties. These files must be located in the same folder; consider employing a RAID configuration – such as RAID 5[12] – that supports disk reads.
- **Hyper-V LUNs** (hosting vDesktops and associated files)

  You can configure a single LUN for each Hyper-V server to host the following:
  - .xml files containing the vDesktop configuration
  - .bin files containing saved-state information
  - PVS write cache drives

Best practices for networking

This section outlines best practices for configuring Virtual Connect Flex-10 technology in the VDI solution.

Prerequisite

This section assumes that you are familiar with Ethernet networking technology in general, as well as the features and operation of Virtual Connect Flex-10 technology. For more information, refer to the following sources:

- Unraveling the full potential of virtualization – Virtual Connect Technology Guide
- HP Virtual Connect Flex-10 technology: Convergence with FlexFabric components
- HP Virtual Connect technology for the HP BladeSystem c-Class
- HP Virtual Connect for the Cisco Network Administrator
- HP Virtual Connect for c-Class BladeSystem Version 3.10/3.15 User Guide

Guidelines

This section suggests a Virtual Connect Flex-10 network configuration that can support the VDI solution described in this white paper.

At a minimum, the following networks should be implemented:

- Management network – supports communications between Hyper-V hosts, HP Integrated Lights-Out (iLO) processors, and OA

[12] For further information, refer to the Citrix website.

- Production network – supports communications between users and outside network traffic, such as the corporate network and Internet access
- iSCSI network – supports the SAN and vDesktop deployments

Optionally, you can implement the following networks:

- Clustering network – supports Hyper-V high-availability
- Live migration network – moves running VMs between Hyper-V servers
- Backup network – supports server backup

Creating server profiles

This section describes how to create Virtual Connect Flex-10 server profiles to be assigned to either Hyper-V servers or P4000sb storage blades in an HP BladeSystem enclosure.

Comply with the following prerequisites:

- Ensure that the latest Virtual Connect firmware has been applied to the Virtual Connect Flex-10 modules[13]
- Prior to creating profiles, ensure that all servers in the enclosure have been powered off

Note
Profiles are not applied unless servers are powered off.

Perform the following procedures.

Defining and assigning profiles for P4000sb storage blades

1. In HP Virtual Connect Manager menu, click on **Server→Define Server Profile**.
2. Specify a name for the new profile.
3. Define the following profile on two connections (creating two new network connections), with the bandwidths shown:

Note
The profile includes a failover network, thus the second connection is required.

– iSCSI network	7 Gbps
– iSCSI network	7 Gbps

4. Assign the new profile to Bay 7.
5. Copy the profile, assigning a new name to each copy; assign the profiles to Bays, 8, 15, and 16, which are also housing P4000sb storage blades.

[13] For further information, refer to the HP website.

Defining and assigning profiles for Hyper-V servers

1. Create a new profile, selecting a name that reflects that this server is part of the management/hypervisor host schema.
2. Create four new network connections, for a total of six (the iSCSI networks having been created earlier). Assign bandwidths as follows:

– iSCSI network	7 Gbps
– iSCSI network	7 Gbps
– Management network	2 Gbps
– Management network	2 Gbps
– Production network	1 Gbps
– Production network	1 Gbps

3. Copy the profile, assigning a new name to each copy; assign the profiles to Bays 1 – 6 and 9 – 14, the bays housing Hyper-V servers.

Figure 6 shows how Virtual Connect Flex-10 technology was used to partition physical dual-port network interface cards (NICs) in each Hyper-V host into six (of a possible eight) FlexNICs to support the VDI solution.

Figure 5. Partitioning a dual-port Hyper-V server NIC to create six FlexNICs

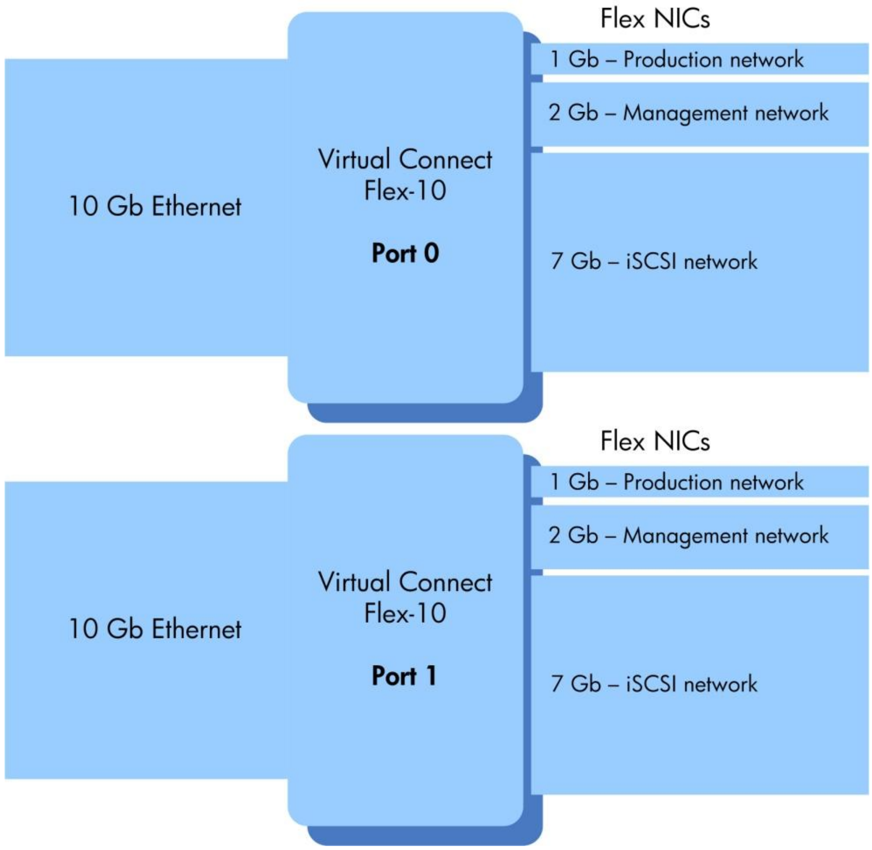

Note
See Appendix B: Designing FlexNICs for more information on how Virtual Connect Flex-10 partitions physical NICs into FlexNICs.

Optimizing the PVS configuration

Configuring PVS can present a challenge when you are implementing a VDI solution. To help you optimize your solution, this section describes how PVS works and suggests some best practices for its configuration.

Overview

Based on software-streaming technology, PVS allows you prepare a client for imaging by installing an operating system and any required software on that device. You can then create a vDisk based on the client's hard drive and save this vDisk on the network, to the PVS server or back-end storage, for example.

Once the vDisk is available on the network, the client does not require a local hard drive since it can boot directly from the network.

PVS streams the contents of the vDisk to the client on demand, in real time, allowing the client to behave as if it were running from its local drive. Unlike thin-client technology, processing takes place on the client[14].

Sequence of operation

The following steps describe how PVS allows a client to download a vDisk image, as shown in Figure 6:

1. The client contacts a Dynamic Host Configuration Protocol (DHCP) server, which ideally has a separate PVS scope.
   a) The client obtains an IP address.
   b) The DHCP scope[15] tells the client to download the network bootstrap program (NBP) – **ardbp32.bin** – from the Trivial File Transfer Protocol (TFTP) server, providing the appropriate IP address and the name of the file.
2. The client contacts the TFTP server.
   a) The PVS client on the client downloads the NBP, which is approximately 10 KB in size and contains the names and IP addresses of four PVS servers.
   b) The NBP installs a BIOS Int 13 disk I/O redirector and PVS protocol on to the client.
3. The client contacts the first available PVS server named in the NBP.

[14] For further information, refer to the Citrix website.
[15] Options 66 and 67

4. The PVS server recognizes the client by its Media Access Control (MAC) address and checks its database to see which vDisk has been assigned to the particular client.

   The vDisk contains an image of a desktop operating system, typically Windows XP or Windows 7.

   The vDisk is stored on a Windows file share, a SAN, network-attached storage, or the PVS server.

   PVS Streaming Service must run under an account that has access to the vDisk.

   vDisks use Microsoft VHD format and may be either fixed or dynamic (see Fixed or dynamic vDisks?).

Note
vDisk properties are stored in a .pvp file; vDisk I/O contents are stored in a .vhd file. The .pvp and .vhd files must always reside together in the same folder. If a vDisk is moved to a new location, the corresponding .pvp and .vhd files must also be moved to the same folder. The .pvp and .vhd files should be backed up as part of the normal PVS backup process.

5. The client mounts the vDisk over the network and begins to use the operating system as though the vDisk were local.

   The User Datagram Protocol (UDP)-based, customized, disk I/O redirector on the client sends hard drive requests to RAM, on to a NIC, and, lastly to the vDisk. Thus, the client is able to read the PVS vDisk over the network.

   If there are network issues, this communication process waits for network connectivity to be resumed, then continues where it left off.

6. The PVS disk driver on the client reads from the vDisk and writes to the PVS write cache file.

   PVS intercepts calls for requirements such as domain and computer names, interrogates the PVS database, and inserts the appropriate vDesktop name.

Figure 6. Client downloading a vDisk image

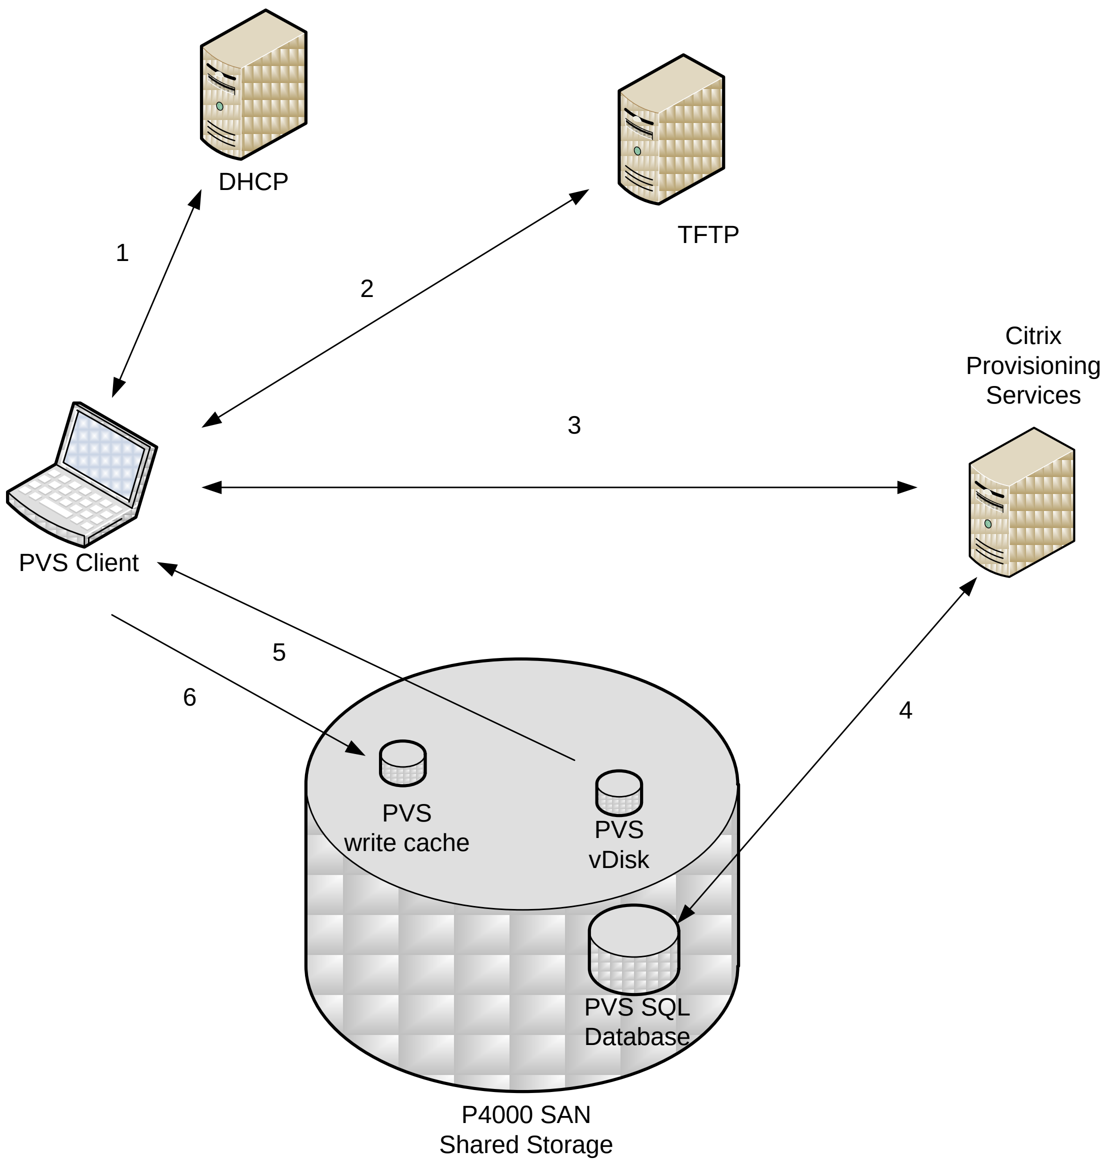

Best practices for configuring PVS on a network

Citrix knowledge base article CTX117374 provides best practices for configuring PVS on a network. You should follow these best practices if you have to address issues such as slow performance, image-build failures, lost connections to the streaming server, or excessive retries from the client.

PVS memory/cache best practices

To obtain optimum performance from a PVS server, ensure that the server is configured with sufficient memory to cache the vDisk(s). If enough memory is available, the PVS server will cache vDisk contents in Windows System Cache memory and service client vDisk requests from memory as opposed to reading from disk. The more vDisks that a PVS server supports, the greater amount of memory it will require for caching.

Citrix provides a white paper entitled Advanced Memory and Storage Considerations for Provisioning Services that includes formulas to accurately determine PVS memory requirements.

PVS storage requirements

In general, PVS storage requirements depend on the following:

- Number of vDisk images to be created and maintained
- Type of vDisks to be used
- Operating system installed in the image
- Components or applications installed in each image
- Plans for potential future application installations
- Location of the client write cache file[16]

vDisk access modes

The following vDisk modes can be used:

- **Standard mode**

  Standard mode allows the user to make changes to the information read from the vDisk image; however, the changes do not persist upon reboot. Thus, you receive a fresh image every time you log on.

Note
Write cache only clears when the client is gracefully shut down or rebooted.

  This is a one-to-many model whereby the read-only vDisk supplies a desktop image to many users who utilize the image as though it were a local hard disk. Changes are written to write cache.

Note
Many user-based preferences and customizations can be saved by means of roaming profiles. Consider installing Citrix Profile Management on the vDisk before setting Standard mode and using Active Directory to manage user profiles[17].

  Standard mode allows users to receive a standardized desktop and make temporary changes without affecting other users.

- **Private mode**

  Private mode places the vDisk into write mode, allowing the user to make changes to the vDisk. Rather than being written to write cache, changes take immediate effect on the vDisk.

  The user has exclusive control of the vDisk, making this a one-to-one model.

  Private mode is used by the administrator to create and make changes to the vDisk.

- **Differential mode**

  Essentially a hybrid of standard and private modes, differential mode features a read-only vDisk; however, changes written to write cache persist through reboot.

[16] For further information, refer to the Citrix website.
[17] For further information, refer to the Citrix website.

A disadvantage of this mode is that, if changes are made to the vDisk and write cache is subsequently deleted, the changes are lost.

Differential mode can be used in an environment where write-cache files are only deleted when they are no longer needed, such as a test and development or classroom system.

Impact on IOPS

In general, the selection of a suitable PVS disk access mode makes no difference to the quantity of disk IOPS. However, the selected mode impacts where disk write activity takes place and may affect vDesktop performance.

In private and differential modes, disk writes occur on the vDisk .vhd file; in standard mode, disk writes occur on the write-cache file.

Since the .vhd file and write-cache file may be located in different LUNs on the SAN and the LUNs may be optimized for either read or write performance, consider the implications when planning and designing the infrastructure.

PVS write cache

This section highlights best practices for the following:

- Locating write cache
- Assigning write-cache drives in bulk
- Configuring write cache

Locating write cache

The location of the PVS write cache can impact server and client performance, server scalability, and cost.

Table 4 outlines options for locating write cache and the associated consequences[18].

Table 4. Locating PVS cache

Location	**Comments**	**Benefits**	**Disadvantages**
PVS server's local disk	Exists as a PVS temporary file	• Easy to set up	• Slow, due to the time taken to traverse the network between the PVS server and the client
Client's memory	Exists as a temporary file in RAM	• Best performance	• RAM cannot be used to support workload • Difficult to size • Client crashes if write cache is full • RAM can be more expensive
Client's hard drive	Exists as a file on the hard drive	• Frees up PVS resources • Low cost	• Slower response time than RAM • Less reliable in a high-availability environment

[18] For further information, refer to the Citrix website.

To maximize performance, take the following steps to ensure that a disk subsystem hosting write cache is optimized for disk writes:

- Enable battery-backed write cache, if available
- Optimize disk hardware RAID (RAID 1 or RAID 10)

Assigning write-cache drives in bulk

In a XenDesktop on Hyper-V solution, write-cache drives must be assigned after vDesktops have been created, which can be accomplished from the SCVMM server.

Citrix provides a Windows PowerShell script that can be run from the SCVMM server to attach a VHD from the library within the SCVMM server to all machines matching a particular naming pattern. Attaching a write-cache drive to each VM takes about 30 – 90 seconds or perhaps longer, depending on storage and network speeds[19].

Configuring write cache

While typically recommending a 2 – 3 GB write cache,[20] Citrix also suggests testing your environment to determine the optimum sizing.

Note
Write cache contains the Windows pagefile – typically at least 1 GB – as well as the deltas created by the user from the vDisk.

Consider configuring write cache to clear during a graceful shutdown so that it does not continue to grow.

Setting up the desktop operating systems

This section focuses on setting up the desktop operating systems that are most commonly virtualized in a XenDesktop on Hyper-V VDI environment: Windows XP and Windows 7.

Note
Support for Windows Vista is also available.

You should consider the significant differences in Windows XP and Windows 7 implementations when selecting a desktop operating system.

For example, some of your applications may require a Windows XP environment or only support a 32-bit platform. Since Microsoft is currently indicating the expiration of Windows XP in 2014, you may consider using Citrix XenApp to virtualize and deliver applications as opposed to a Windows XP deployment.

The remainder of this section outlines best practices for setting up the desktop operating system.

[19] For further information, refer to the Citrix website.
[20] For further information, refer to the Citrix website.

Configuring the desktop NIC

The default network adapter for a Hyper-V VM is the synthetic adapter, which is optimized for a virtualized environment. However, this adapter is purely virtual and cannot be associated with any BIOS-level hardware, whether virtualized or not. Since the PVS pre-boot execution environment (PXE) boot process requires a BIOS-level interface, the synthetic adapter cannot be used. Instead, Hyper-V includes a legacy network adapter that includes BIOS-level integration and supports PXE booting. This legacy network adapter must be added to the VM and configured as the default boot device.

To optimize performance, the vDesktop image should include both adapters, allowing the legacy adapter to be used for PXE booting and the higher-performing synthetic adapter to be used for passing network traffic after the operating system has booted. Both adapters can be connected to the same logical network since the synthetic adapter has precedence over the legacy adapter in the route table.

When configuring both network adapters within a vDesktop, run **bindcfg** from the PVS installation folder to verify that the legacy adapter is bound to the PVS device driver before creating a vDisk image. If this device driver is bound to the synthetic adapter, the boot process cannot be completed[21].

Configuring the VHD partition disk offset with Windows XP

To optimize performance and minimize disk IOPS when deploying Windows XP in a virtual environment, use the **diskpart** utility to manually configure the VHD partition offset[22].

By default, the Windows XP setup process misaligns its partition by creating a signature on the first part of the disk and then starting the actual partition at the last few sectors of the first block. To mitigate this misalignment, manually set the partition offset to 1024 KB during the installation of the desktop golden image prior to virtualization, as follows:

1. Use **diskpart** to manually create a partition.

Note
When using diskpart in a Windows 7 environment, ensure you are using the Windows 7 diskpart utility.

2. Configure the partition to start at sector 128.

Note
Since a sector contains 512 B, the first sector of the partition starts at precisely 64 KB.

After re-alignment, each I/O from the partition results in a single I/O from storage.

[21] For further information, refer to the Citrix website.
[22] For further information, refer to the Citrix website.

Optimizing the vDesktop

To maximize scalability, it is important to optimize vDesktops prior to deployment. Citrix offers the following recommendations:

- Windows 7
  - Optimizing Windows 7 for FlexCast Delivery
  - Windows 7 Registry Optimizations
  - Optimize Windows 7 visual effects for Virtual Desktops
- Windows XP
  - Best practices for optimizing Windows XP

Optimizing scalability

HP, Citrix, and Microsoft collaborate at an engineering level to ensure that you can benefit from software, hardware, and service solutions that are jointly tested, certified, and tuned to deliver optimal server performance. Engineering teams have developed a range of recommended configurations and services for Citrix applications, focusing on particular business needs and technical scenarios.

This section provides an example of scalability in a XenDesktop on Hyper-V environment and offers guidelines for characterizing performance and optimizing scalability in this environment.

Reference architecture

HP has published a white paper[23] that describes a reference architecture[24] for desktop virtualization in the enterprise environment. This solution uses the following components:

- XenDesktop
- Windows Server 2008 R2 Hyper-V
- HP ProLiant BL460c G6 server blades
- HP P4800 G2 SAS BladeSystem SAN storage

The reference architecture was designed to support approximately 800 Microsoft Office 2007 users on Windows XP.

Results

Test results indicated that the HP reference architecture described earlier can support approximately 80 vDesktops per server, or 10 vDesktops per CPU core. CPU utilization was determined to be the limiting factor for this solution; the P4800 G2 SAN was able to deliver the capacity needed by this architecture when configured with optimal RAID levels.

Results were obtained by analyzing server metrics in order to determine the optimum number of users that each XenDesktop/Hyper-V host could support while delivering the resources needed to accommodate spikes in utilization. Windows Performance Monitor (Perfmon) metrics were utilized to determine when 80% CPU utilization – the level typically used by HP to specify when the optimal number of users is active – was reached; Login VSI Analyzer confirmed that the user experience would have been satisfactory at this time. Table 6 details the results.

[23] For further information, refer to the HP website.
[24] Based on HP Converged Infrastructure

Table 6. Capacity and performance of the reference architecture

	Single virtual desktop	Single host	Total
Disk capacity			
	31.5 GB	2.52 TB	25.2 TB*
Disk performance			
Virtual desktop creation	20 IOPS	1,600 IOPS	16,000 IOPS
Virtual desktop boot	15 IOPS	1,200 IOPS	12,000 IOPS
Sustained workload	2.5 IOPS	200 IOPS	2,000 IOPS
Network			
ICA connection	20 Kbps	1,600 Kbps	16 Mbps
Memory			
Per virtual desktop	512 MB	36 GB	512 GB
Operating system		2 GB[25]	
Number of vDesktops			
		80	800

* Configured for Network RAID 10/hardware RAID 5

The remainder of this section highlights best practices for characterizing and enhancing Hyper-V scalability in a XenDesktop environment.

Hyper-V performance counters

To help you characterize performance and optimize scalability, Microsoft has added specific Perfmon counters to monitor Hyper-V functions. However, using Perfmon to characterize Hyper-V performance is different from using Perfmon in a typical Windows environment. For example, in a typical Windows environment, the **\%Processor Time** counter captures all of the server's CPU activity. In a Hyper-V environment, you should monitor the **\Hypervisor: \Hyper-V Hypervisor Logical Processor(_Total)\% Total Run Time** counter.

Since Hyper-V uses the term "partitions" to denote VMs, ensure that you include partition monitoring when you are gathering Hyper-V performance data.

Using the new counters

You should monitor the new Hyper-V counters in addition to traditional counters in order to obtain an overall view of server performance.

The counters installed with Hyper-V include:

- Hyper-V Hypervisor
- Hyper-V Hypervisor Logical Processor
- Hyper-V Hypervisor Partition
- Hyper-V Hypervisor Root Partition
- Hyper-V Hypervisor Root Virtual Partition
- Hyper-V Hypervisor Virtual

[25] Microsoft recommends 2GB RAM or greater with additional memory for each running guest operating system.

- Hyper-V Legacy Network Adapter Processor
- Hyper-V Task Manager Details
- Hyper-V Virtual IDE Controller
- Hyper-V Virtual Machine Bus
- Hyper-V Virtual Machine Summary
- Hyper-V Virtual Network Adapter
- Hyper-V Virtual Storage Device
- Hyper-V Virtual Switch
- Hyper-V Virtual Switch Port

However, it is important to remember that these new counters do not provide information on memory utilization within a VM. You must run Perfmon within the VM to characterize memory utilization.

Health Summary status

The following Health Summary counters are provided:

- Health OK
- Health Critical

However, in the test environment, when a Hyper-V host became saturated and began to reboot VMs, this event was not reflected in the Health Status summary.

Tuning Hyper-V

Microsoft has developed a guide – "Performance Tuning Guidelines for Windows Server 2008 R2"[26] – for tuning Windows Server 2008 R2. This guide describes important settings that can be adjusted to improve the performance of Windows Server 2008 R2; in addition, potential effects are addressed to help you make an informed decision about the relevance of a particular setting to your system, workload, and performance goals.

Fixed or dynamic vDisks?

This section describes how the use of fixed or dynamic disks for a vDisk implementation can affect vDesktop performance.

Note
The vDesktop looks at a vDisk in much the same way a traditional PC looks at a hard drive.

Microsoft supports the following types of vDisks in a Hyper-V environment:

- **Dynamic vDisk**

  Dynamic vDisks are not set to a specific size. They are thinly provisioned, only using the space they currently require and expanding as needed.
- **Fixed vDisk**

  Fixed vDisks are set to a specific size and acquire that amount of space as soon as they are started.

Citrix recommends using fixed vDisks for the following reasons[27]:

[26] For further information, refer to the Microsoft website.
[27] For further information, refer to the Citrix website.

- Dynamic vDisks include 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.

  A fixed vDisk does not include the extra byte at the end of a .vhd file.
- Dynamic vDisks have an expansion algorithm[28] that generates significant overhead on the storage device when the drive expands.

Proof-of-concept

The best way to ensure success in a VDI environment is to implement a XenDesktop proof-of-concept (POC).

Available through My Citrix, Citrix provides a toolkit that can help you prepare a XenDesktop on Hyper-V POC to run in an isolated environment in your own lab or at a customer's site. Citrix also offers a checklist of items that should be completed prior to starting this POC[29].

[28] In 2 MB blocks

[29] For further information, refer to the Citrix website.

Appendix A: Installing the P4800 G2 SAN

This appendix outlines the installation of an HP P4800 G2 63 TB SAS BladeSystem SAN solution.

Components

The P4800 G2 SAN solution includes the following components:

Note
This SAN must be ordered with factory integration in an HP BladeSystem c7000 enclosure, two Virtual Connect Flex-10 modules, a rack, and a power distribution unit (PDU).

- Two HP BladeSystem 3 Gb SAS switches
- Two MDS600 systems
- Four P4000sb storage blades

When shipped, the SAN has all its cabling installed; SAS switch zoning is in place.

Important
Do not alter SAS switch zoning or the local array configuration of P4000sb nodes in enclosure bays 7, 8, 15 and 16 for any reason.

Network configuration

This section provides guidelines for configuring management and production networks, and iSCSI and SAS connectivity.

Management and production networks

The enclosure should be configured to meet the standards set by your IT organization.

Ideally, once configured, all interconnects and iLO ports within the c7000 enclosure should be accessible on your management network. At a minimum, you should connect your management network to at least one Onboard Administrator (OA) in the enclosure.

You are highly recommended to cable the enclosure as shown in Figure A-1.

Where Virtual Connect Flex-10 modules connect to the network core switches, define virtual LANs for the management and production networks.

Figure A-1. Core network connections (showing interconnect bay numbering)

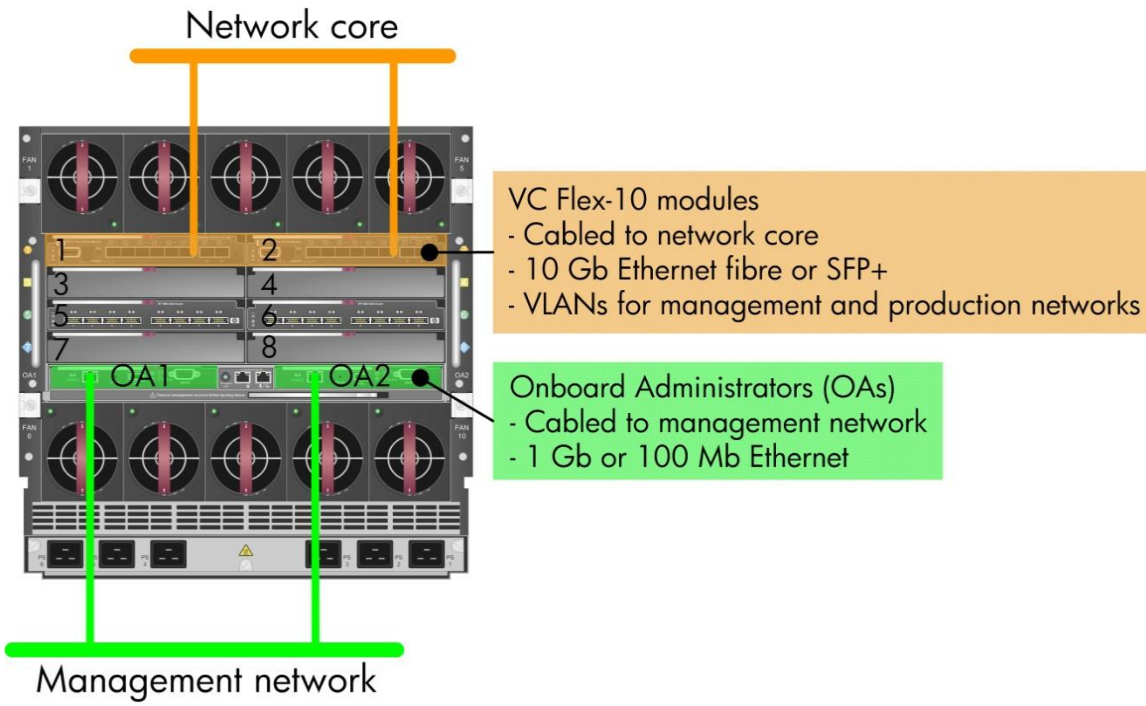

iSCSI traffic

This appendix assumes iSCSI traffic is internal to the c7000 enclosure.

If, however, you wish iSCSI traffic to exit the enclosure, you should cable an additional, dedicated 10 GbE link-pair to the network core, with one port of each Virtual Connect Flex-10 module assigned to this network. Do not share this cable and port with any other traffic.

For more information on configuring an external iSCSI network, refer to Chapter 4 of the HP document, "HP P4000 SAN Solution User Guide."[30]

[30] http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c02063204/c02063204.pdf

SAS configuration

Figure A-2 shows SAS connections pre-configured at the factory by HP.

Figure A-2. SAS cabling, as pre-configured by HP

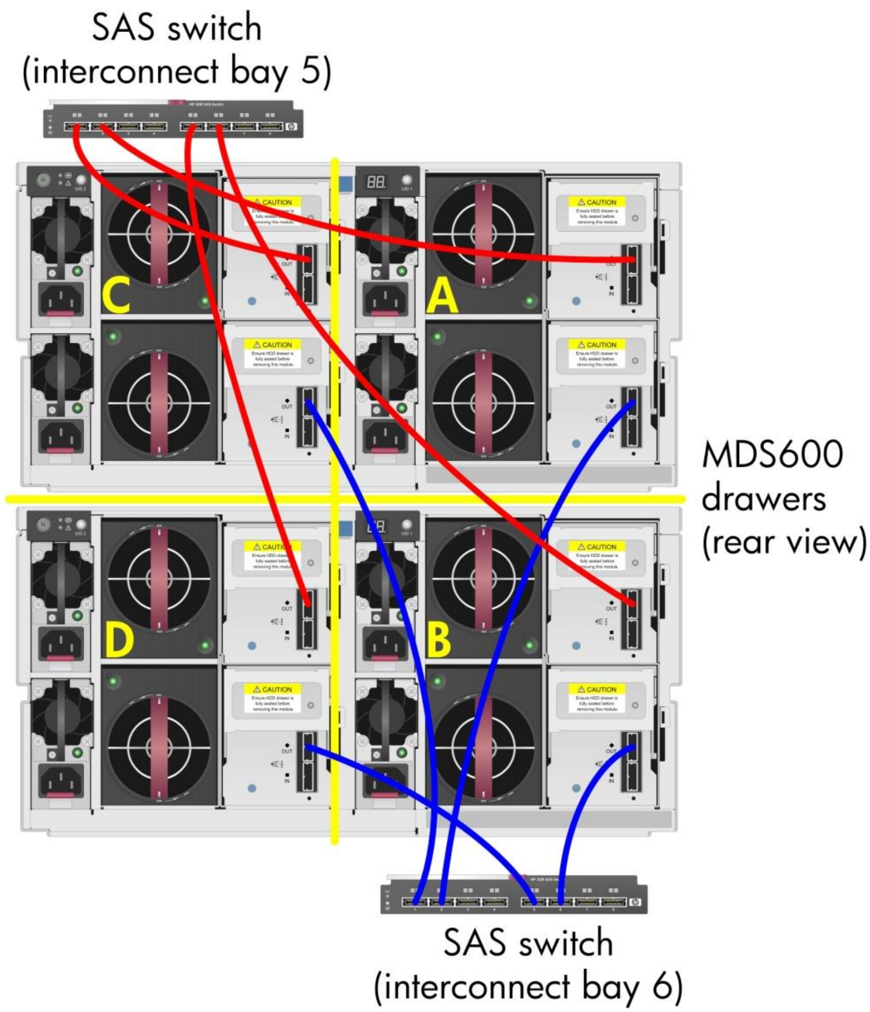

Powering up the SAN solution

Prior to configuring the SAN, make certain you follow the power-on sequence outlined in Table A-1. Following this sequence ensures that solution components communicate with each other as intended.

If the solution has been powered on before you start the configuration, use the power buttons within OA[31] to power down the SAS switches and P4000sb storage blades. Reapply power in the correct order.

Table A-1. Power-on sequence

Step	Component	Comments
1	MDS600 systems	Manually power-on each MDS600 system. The MDS600 systems must be online before SAN/iQ is invoked. Do not power-on the c7000 enclosure until the MDS600 systems have been powered on for at least three minutes, allowing all drives to spin up.
2	SAS switches	To ensure the SAS switches power-on next, it may help to temporarily unseat the Virtual Connect Flex-10 modules from Interconnect Bays 1 and 2. Set the power- on sequence via OA.
3	Virtual Connect Flex-10 modules	Ensure the SAS switches have been powered on. Set the power-on sequence via OA.
4	P4000sb storage blades	Ensure the SAS switches have been powered on for one – two minutes. Set the power-on sequence via OA.
5	Hyper-V Host 1	If you wish to isolate iSCSI traffic within the enclosure, you must install Hyper-V on Host 1 and create a VM on local storage to support the CMC. Set the power-on sequence via OA. If you are using an external CMC and do not wish to isolate your iSCSI network, include Host 1 in step 6.
6	Remaining Hyper-V hosts	Power-on Hyper-V hosts in any order.

Note
In the event of a catastrophic power failure, you can set the power-on order via OA (requires OA firmware 3.0 or later).

Configuring the enclosure

Instructions for configuring your enclosure are included with the SAN solution and are also available with other HP BladeSystem documentation at http://www.hp.com/go/bladesystem.

[31] Requires administrator access to Onboard Administrator

Appendix B: Designing FlexNICs

The strength of Virtual Connect Flex-10 is that it allows you to segment 10 Gb of bandwidth in order to isolate different types of network traffic. You can assign the appropriate bandwidth to each traffic type, as shown in Figure B-1.

Figure B-1. Configuring FlexNICs

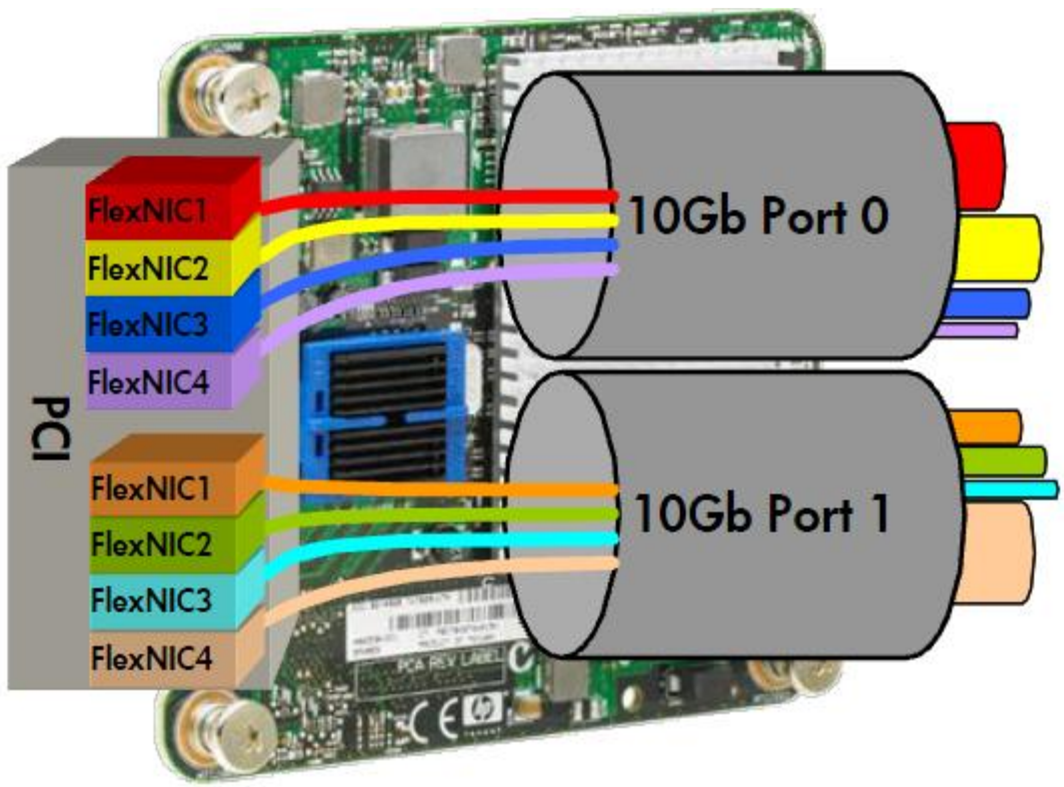

For more information

HP ProLiant servers	http://www.hp.com/go/proliant
HP data storage	http://www.hp.com/go/storage
HP Converged Infrastructure	http://www.hp.com/go/ci
HP BladeSystem	http://www.hp.com/go/bladesystem
HP FlexFabric	http://www.hp.com/go/flexfabric
Citrix XenDesktop	http://www.citrix.com/english/ps2/products/product.asp?contentid=163057
Citrix Provisioning Services	http://support.citrix.com/product/provsvr/psv5.0/
Windows Server 2008 R2 Hyper-V	http://www.microsoft.com/windowsserver2008/en/us/hyperv-main.aspx
Intel Xeon processors	http://www.intel.com/go/xeon

To help us improve our documents, please provide feedback at http://h20219.www2.hp.com/ActiveAnswers/us/en/solutions/technical_tools_feedback.html.

4AA3-3066ENW, Created February 2011