



Technical Report

Citrix XenDesktop on NetApp Storage Solution Guide

Rachel Zhu, Neil Glick, Will Strickland, NetApp
March 2013 | TR-4138

TABLE OF CONTENTS

1	Scope	5
2	Introduction to Clustered Data ONTAP	5
2.1	Benefits of Clustered Data ONTAP	5
2.2	NetApp Storage Cluster Components	7
2.3	Clustered Data ONTAP Networking Concepts	8
2.4	Cluster Management	8
3	Introduction to Citrix XenDesktop	8
3.1	Citrix FlexCast Delivery Technology	9
3.2	Desktop Types	9
3.3	Desktop Provisioning Method	10
3.4	Provisioning Decision	14
4	Solution Architecture and Performance Data	15
4.1	High-Level Architecture	15
4.2	XenDesktop on Clustered Data ONTAP Solution Components	18
5	Performance Test Results	20
5.1	Performance Test Tool	20
5.2	Performance Test Plan	21
5.3	Citrix PVS Workload Characteristics	21
5.4	Performance Effect of Flash Solutions	25
6	Storage Configuration Best Practice	25
6.1	Protocol Decision	25
6.2	Clustered Data ONTAP Failover Group	26
6.3	One LIF per Volume	26
6.4	LIF Migrate	27
6.5	10 Gigabit Ethernet	27
6.6	Jumbo Frames	27
6.7	Flow Control Overview	28
6.8	Aggregate and Volume	28
6.9	Secure Default Export Policy	28
6.10	Avoid Misalignment	29
7	Storage Operation Best Practice	30
7.1	Deduplication Decision	30
7.2	Thin Provisioning and Volume Autogrow Policies	31

7.3	Space Reclamation.....	31
7.4	Antivirus.....	31
7.5	Monitoring NetApp and XenDesktop Infrastructure.....	31
8	Backup and Recovery of the Virtual Desktops in a VMware vSphere Environment	32
9	User Data and Profile Management	34
9.1	User Data.....	34
9.2	User Profile Data.....	34
9.3	User-Installed Applications.....	35
10	Application Virtualization with XenApp.....	36
11	Citrix Components Scalability Consideration	37
11.1	Infrastructure Components Limit.....	37
11.2	Hypervisor Configuration Limits and Pod Concept.....	39
11.3	VM Resource Recommendation	40
12	Storage Sizing Guide.....	40
12.1	Solution Assessment	40
12.2	Capacity Considerations	41
12.3	Performance Considerations.....	42
12.4	Desktop Deployment Scenarios.....	42
12.5	IOPS and Capacity Considerations.....	43
12.6	Storage Scalability Guidelines	44
	Acknowledgement	45
	References.....	45
	Citrix References	45
	NetApp References	46
	VMware References	46

LIST OF TABLES

Table 1)	Management tools.....	8
Table 2)	Provisioning methods.....	14
Table 3)	Test cases.....	21
Table 4)	Storage protocols for XenDesktop.....	25
Table 5)	Dedupe recommendations	30
Table 6)	XenDesktop on vSphere backup methods	33

Table 7) XenDesktop component configuration and limits.....	38
Table 8) VM resource recommendations.....	40
Table 9) Recommendations for hosted shared desktops.	40
Table 10) Sizing guidance.....	42
Table 11) Citrix PVS-based deployment scenarios.....	42
Table 12) IOPS and capacity assumptions.....	43
Table 13) Storage scalability guidelines.....	44
Table 14) Scalability guidelines.....	45

LIST OF FIGURES

Figure 1) Scalability in clustered Data ONTAP.....	6
Figure 2) Multi-tenancy concept.....	7
Figure 3) Ports and LIFs example.....	8
Figure 4) Citrix FlexCast delivery technology (graphic provided by Citrix).....	9
Figure 5) vDisk and write cache.....	11
Figure 6) MCS file layout.....	12
Figure 7) Hypervisor snapshot I/O workflow.....	13
Figure 8) Virtual machine clone creation.....	14
Figure 9) Architecture for pooled desktops using Citrix PVS.....	16
Figure 10) Architecture for persistent desktops using Citrix PVS.....	18
Figure 11) XenDesktop on clustered Data ONTAP architecture.....	19
Figure 12) XenDesktop on clustered Data ONTAP volume layout example.....	20
Figure 13) I/O of the write workload.....	22
Figure 14) 500-seat boot, I/O comparison, PVS with personal vDisk.....	24
Figure 15) 500-seat login, I/O comparison, PVS with personal vDisk.....	24
Figure 16) Setting jumbo frames.....	28
Figure 17) Diskpart.....	30
Figure 18) Separation of volumes based on data types.....	36
Figure 19) VSC cloning.....	37
Figure 20) Desktop transformation accelerator.....	39
Figure 21) FAS2240-2 system (2U) with 24 internal SAS drives (600GB; 15K RPM).....	45

1 Scope

The new world of desktop virtualization is much more comprehensive than the traditional virtual desktop infrastructure (VDI) because it addresses the needs of various types of users in a customer environment. Citrix XenDesktop with the various [FlexCast](#) delivery technologies can effectively meet the performance, security, and flexibility requirements of various types of users. Some of the popular FlexCast delivery technologies that enable this are:

- [Hosted VDI](#): Targeted toward knowledge workers
- [Hosted shared](#): Targeted toward task workers
- [Apps on demand](#): Virtualizing applications

Customers require innovative end-to-end storage solutions that can simplify the architecture in a cost-effective manner and provide the maximum uptime. When designing the storage architecture for environments that require a mix of FlexCast delivery technologies, the best practices around profile management, user data management, monitoring, management, and so on should be carefully considered. Also, to architect the solution correctly, the design should factor in the best practices for managing user profiles so that the VM images can be stateless and administrators no longer have to individually patch and update thousands of virtual desktops. Additionally, the storage design should factor in the best practices for user data management and application virtualization to achieve lower cost per desktop.

NetApp® clustered Data ONTAP®, with its key capabilities such as nondisruptive operations, unified storage, multiprotocol architecture, secure multi-tenancy, storage efficiency, read and write performance, and cost-efficient data protection, is ideal for cost-efficiently designing and deploying an end-to-end storage solution for desktop virtualization based on a single XenDesktop FlexCast technology or a mix of several of these technologies.

This solution guide provides key storage design and architecture best practices for deploying a mix of Citrix XenDesktop FlexCast technologies on a NetApp clustered Data ONTAP storage array. It also includes the results from performance testing and solution scalability guidelines that prove that clustered Data ONTAP can cost-effectively scale to thousands of desktops without adding complexity.

2 Introduction to Clustered Data ONTAP

For the first time, NetApp brings enterprise-ready, unified scaleout storage to market with the release of Data ONTAP 8.1. For details on how to configure clustered Data ONTAP with VMware vSphere®, refer to [TR-4068: VMware vSphere 5 on NetApp Data ONTAP 8.1 Operating in Cluster-Mode](#).

2.1 Benefits of Clustered Data ONTAP

All clustering technologies follow a common set of guiding principles. These principles include the following:

- **Nondisruptive operation.** The key to efficiency and the lynchpin of clustering is the ability to make sure that the cluster does not fail—ever.
- **Virtualized access is the managed entity.** Direct interaction with the nodes that make up the cluster is in and of itself a violation of the term cluster. During the initial configuration of the cluster, direct node access is a necessity; however, steady-state operations are abstracted from the nodes because the user interacts with the cluster as a single entity.
- **Data mobility and container transparency.** The end result of clustering—that is, the nondisruptive collection of independent nodes working together and presented as one holistic solution—is the ability of data to move freely within the boundaries of the cluster.

- **Delegated management and ubiquitous access.** In large complex clusters, the ability to delegate or segment features and functions into containers that can be acted upon independently of the cluster means the workload can be isolated; it is important to note that the cluster architecture itself must not place these isolations. This should not be confused with security concerns around the content being accessed.

Scale-Out

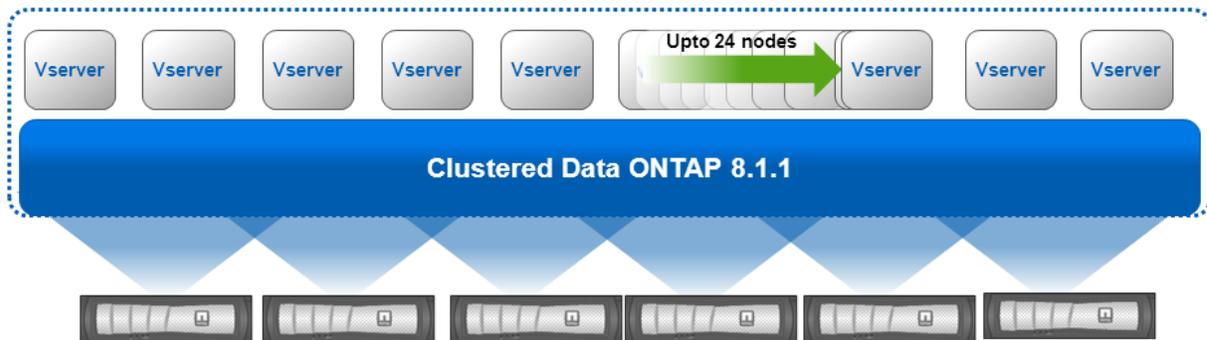
Data centers require agility. In a data center, each storage controller has limits to CPU, memory, and disk shelves. Scale-out means as storage environment grows, additional controllers can be added seamlessly to the resource pool residing on a shared storage infrastructure. Host and client connections as well as datastores can be moved seamlessly and nondisruptively anywhere within the resource pool.

Benefits of scale-out:

- Nondisruptive operations
- Keep adding thousands of users to virtual desktop environment without downtime
- Operational simplicity and flexibility

Clustered Data ONTAP is one way to solve the scalability requirements in a storage environment.

Figure 1) Scalability in clustered Data ONTAP.



Unified Storage

By supporting all common NAS and SAN protocols on a single platform, NetApp unified storage enables:

- Direct access to storage by each client
- Different platforms to share network files without using protocol emulation products such as Samba, NFS Maestro, or PC-NFS
- Simple and fast data storage and data access for all your client systems
- Fewer storage systems
- Greater efficiency from each system deployed

Clustered Data ONTAP has the ability to concurrently support several protocols in the same storage system. Data ONTAP 7G and 7-Mode versions also include support for multiple protocols.

Unified storage is important to XenDesktop solutions. Based on your workload, you can choose different protocol: for example, CIFS SMB 2.1 for user data, NFS for write cache, and SAN LUNs for Windows[®] applications.

The supported protocols are:

- NFS v3, v4, and v4.1, including pNFS
- iSCSI

- Fibre Channel
- FCoE
- SMB 1 and 2

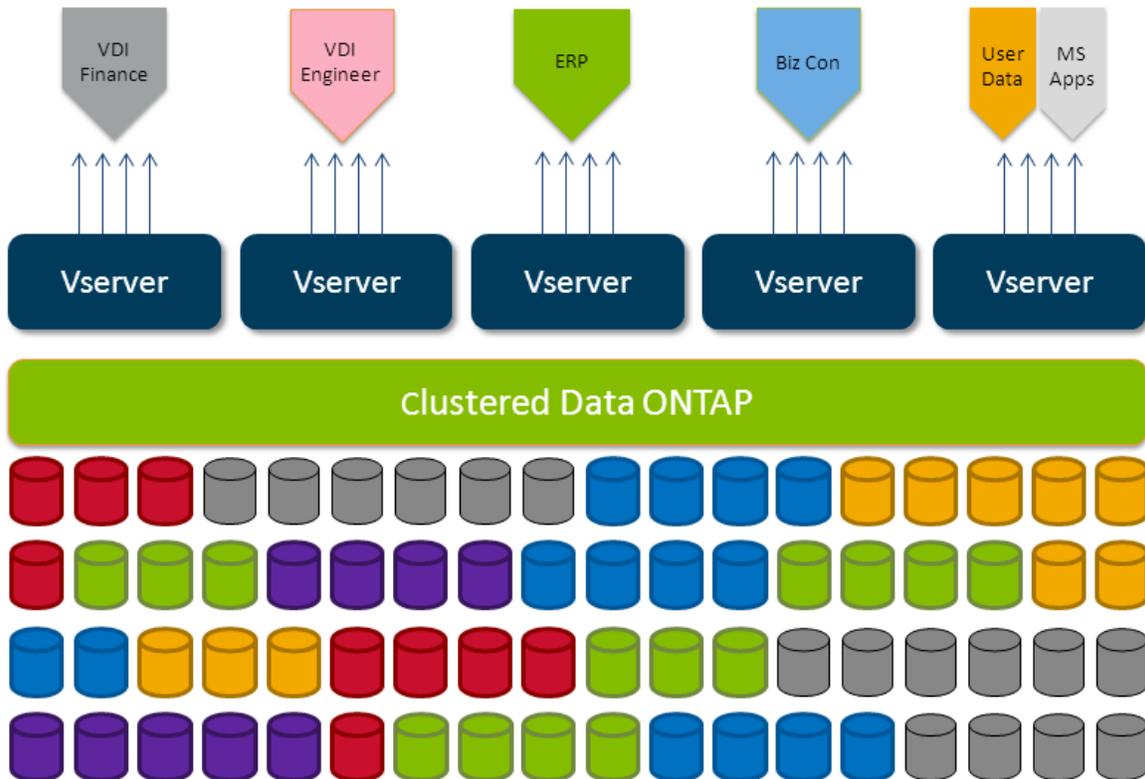
Multi-Tenancy

Isolated servers and data storage can result in low utilization, gross inefficiency, and inability to respond to changing business needs. Cloud architecture, delivering IT as a service (ITaaS), can overcome these limitations while reducing future IT expenditure.

The virtual server (Vserver) is the primary logical cluster component. Each Vserver can create volumes, logical interfaces, and protocol access. With clustered Data ONTAP, each department's virtual desktops and data can be separated to different Vservers. The administrator of each Vserver has the rights to provision volumes and other Vserver-specific operations. This is particularly advantageous for service providers or any multi-tenanted environments in which workload separation is desired.

Figure 2 shows the multi-tenancy concept in clustered Data ONTAP.

Figure 2) Multi-tenancy concept.



2.2 NetApp Storage Cluster Components

Although normally reserved for a glossary, it is important to address some key terms early in the text to establish a common knowledge baseline for the remainder of this publication.

- **Cluster.** The information boundary and domain within which information moves. The cluster is where high availability is defined between physical nodes and where Vservers operate.
- **Node.** A physical entity running Data ONTAP. This physical entity can be a traditional NetApp FAS controller; a supported third-party array front ended by a V-Series controller; and NetApp's virtual storage appliance (VSA), Data ONTAP-v™.

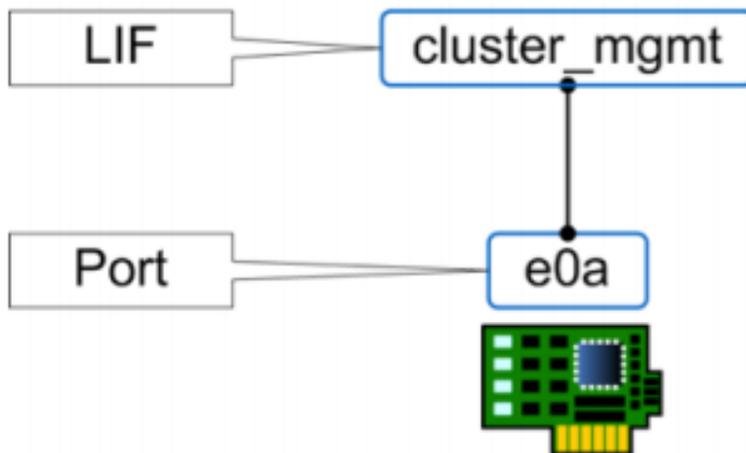
- **Vserver.** A secure virtualized storage controller that behaves and appears to the end user to be a physical entity (similar to a VM). It is connected to one or more nodes through internal networking relationships (covered later in this document). It is the highest visible element to an external consumer, abstracting the layer of interaction from the physical nodes. Based on these two statements, it is the entity used to provision cluster resources and can be compartmentalized in a secured manner to prevent access to other parts of the cluster.

2.3 Clustered Data ONTAP Networking Concepts

The physical interfaces on a node are referred to as ports. IP addresses are assigned to logical interfaces (LIFs). LIFs are logically connected to a port in much the same way that VM virtual network adapter and VMkernel ports connect to physical adapters, except without the constructs of virtual switches and port groups. Physical ports can be grouped into interface groups. VLANs can be created on top of physical ports or interface groups. LIFs can be associated with a port, interface group, or VLAN.

Figure 3 shows the clustered Data ONTAP network concept.

Figure 3) Ports and LIFs example.



2.4 Cluster Management

For complete and consistent management of storage and SAN infrastructure, NetApp recommends using the tools listed in Table 1, unless specified otherwise.

Table 1) Management tools.

Task	Management Tools
Vserver management	OnCommand® System Manager
Switch management and zoning switch vendor	GUI or CLI interfaces
Volume and LUN provisioning and management	NetApp Virtual Storage Console

3 Introduction to Citrix XenDesktop

Citrix XenDesktop transforms desktops as an on-demand service to any user, any device, anywhere. XenDesktop quickly and securely delivers any type of virtual desktop or Windows, web, and software as a service (SaaS) application to all the latest PCs, Apple® Mac® computers, tablets, smart phones, laptops, and thin clients—all with a high-definition Citrix HDX user experience. Also, with the full integration of

Citrix XenApp, you can deliver on-demand applications as a seamless part of your overall desktop management strategy, extending the benefits of virtualization throughout the enterprise.

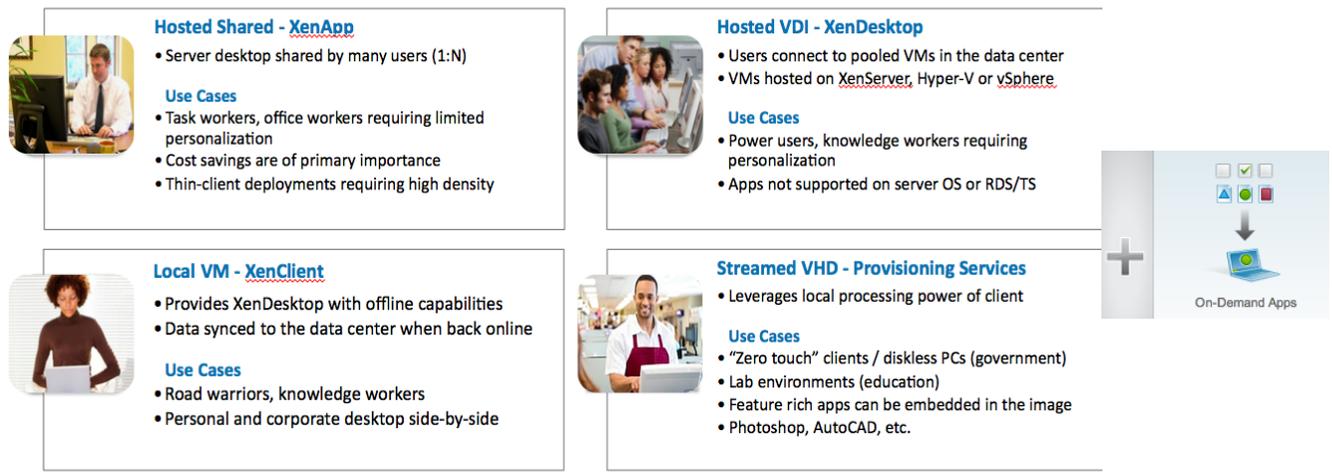
XenDesktop integrates with Citrix XenServer, Windows Server® 2008 Hyper-V™, and VMware vSphere and works out of the box with thin clients.

3.1 Citrix FlexCast Delivery Technology

Citrix FlexCast delivery technology enables IT to optimize the performance, security, and cost of virtual desktops for various types of users.

Figure 4 shows Citrix FlexCast delivery technology.

Figure 4) Citrix FlexCast delivery technology (graphic provided by Citrix).



Different groups of workers across an enterprise have varying performance and personalization requirements. Some require simplicity and standardization, while others need high performance or a fully personalized desktop. XenDesktop can meet all of these requirements in a single solution with Citrix FlexCast delivery technology. For more information on different desktop modules, go to <http://flexcast.citrix.com>.

3.2 Desktop Types

In XenDesktop, collections of machines of the same type are managed as a single entity called a catalog. The machine type determines the hosting infrastructure that will be used for desktops, that is, VMs or physical computers plus the associated storage.

Pooled

Pooled machines provide desktops that are allocated to users on a per-session, first-come, first-served basis. Pooled machines are direct copies of the master VM, although profile management can be used to apply the users' personal settings to their desktops and applications. Changes made to the pooled machines by users with local administrator permissions for their desktops are stored for the duration of the session. These customizations are discarded when the users log off. Maintaining a single master VM in the data center dramatically reduces the update and upgrade maintenance efforts, and it can reduce your storage costs by up to 90%.

Pooled desktops are best used for task workers who require standardized desktops, such as call center operators, retail workers, or students and faculty in educational institutions. They minimize desktop

management and storage costs by providing a locked-down standardized environment for your users. Citrix Provisioning Services (PVS) is a great method to provision pooled desktops.

Dedicated

Dedicated desktops are permanently assigned to individual users. Desktops can be manually assigned, or they can be automatically assigned to the first user who connects to them. Whenever users request a desktop, they are always connected to the same one. Therefore, users can personalize the desktop to suit their needs. Desktop customizations are stored separately and overlaid on copies of the master VM to recreate users' desktops whenever a dedicated machine is started. Dedicated machines provide significant savings for desktop management and storage costs by maintaining a single master VM in the data center for global updates, such as operating system updates and upgrades, while providing users with individual desktops that they can personalize.

Dedicated desktops are best for knowledge workers or mobile users who require individual desktops of which they can take ownership and personalize. NetApp Virtual Storage Console (VSC) and PVS with personal vDisk are two provisioning tools for dedicated desktops. VSC offers data deduplication for storage saving and a redeployment tool to reduce update and upgrade maintenance costs. PVS and personal vDisk are recommended for dedicated desktops because they simplify operations. Customers no longer have to patch or update individual VMs, which means that customers no longer have to size storage for write storms associated with patching/updating VMs.

Streamed

The streamed machine type enables you to provide desktops to repurposed PCs and thin clients that have been configured to load the operating system over the network from provisioning services. Target devices are managed in PVS as a device collection, and the desktops are delivered from a PVS vDisk imaged from a master target device. Using PVS to deliver desktops enables you to leverage the processing power of existing local hardware, while realizing all the benefits of centralized desktop management. This approach offers an inexpensive entry point to desktop virtualization by using existing local resources and reducing to a minimum the need for additional processing and storage capacity in the data center.

3.3 Desktop Provisioning Method

When you use Citrix XenDesktop on NetApp, you can choose three provisioning methods: PVS, machine create service (MCS), and NetApp clone.

Citrix Provisioning Services

The Citrix Provisioning Services (PVS) infrastructure is based on software-streaming technology. This technology allows computers to be provisioned and reprovisioned in real time from a single shared disk image. In doing so, administrators can completely eliminate the need to manage and patch individual systems. Instead, all image management is done on the master image: vDisk. The local hard disk drive of each system may be used for runtime data caching or, in some scenarios, removed from the system entirely, which reduces power usage, system failure rates, and security risks.

vDisk attributes:

- Read only during streaming
- Shared by many VMs, which makes update easy
- Recommended storage protocol: CIFS SMB for vDisk or SAN LUN

Write cache attributes:

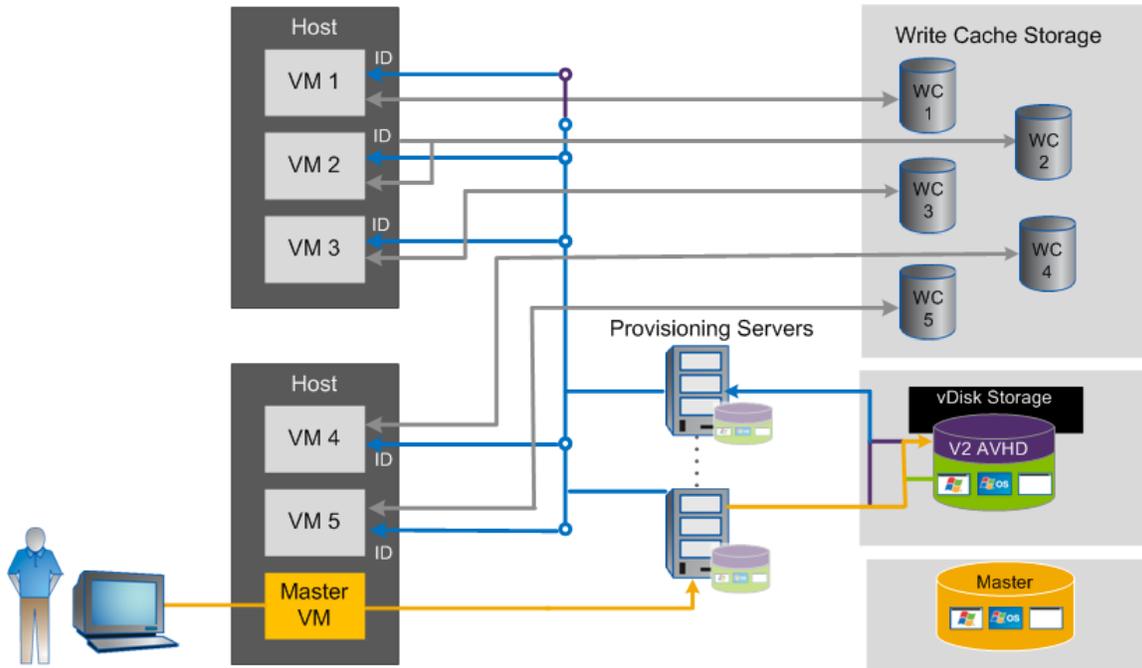
- One VM, one write cache file.
- Write cache file will be empty after VM reboots.

- Recommended storage protocol: NFS for write cache for vSphere and XenServer, SAN LUN for Hyper-V.
- The write cache file size is normally 4GB to 10GB. If PvDisk is used in the deployment, the write cache size will be smaller.

Note: Protocol decision for write cache and vDisk is explained in section 6.1.

Figure 5 shows that vDisk and write cache are the two main components with a PVS-deployed XenDesktop solution. vDisk and write cache need to be on shared storage for enterprise for data integrity, backup, scalability, and easy management.

Figure 5) vDisk and write cache.



Best Practices

Because PVS streams the vDisk (the master image) over the network, you require at least two 1GbE network interface cards (NICs) on your PVS servers. The recommended storage protocol is CIFS SMB for vDisk.

Citrix Provisioning Server supports 10GbE connection, and NetApp recommends using 10GbE for streaming desktops.

Machine Create Service

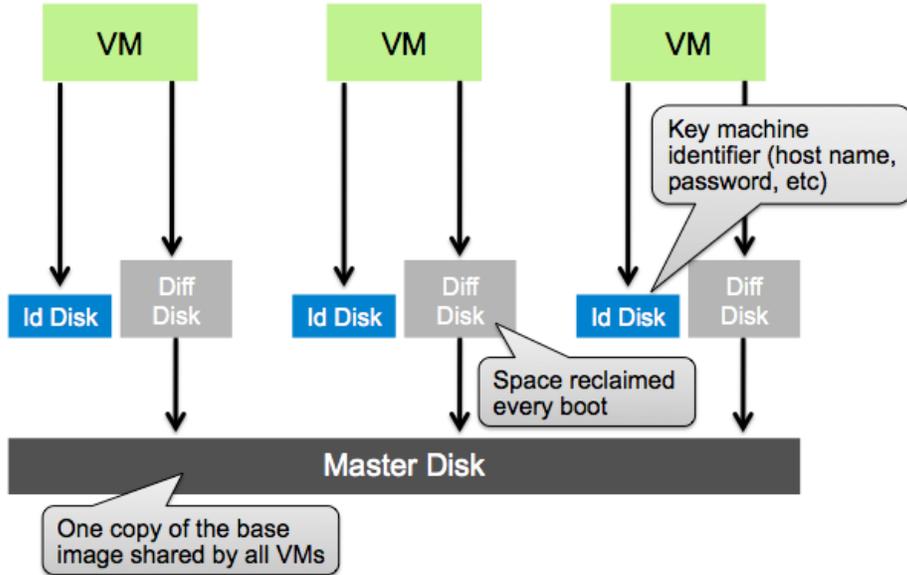
XenDesktop 5 introduced machine create service (MCS). MCS simplifies the tasks of creating, managing, and delivering virtual desktops to users.

When you provision desktops using MCS, a master image is copied to each storage volume. This master image copy uses the hypervisor snapshot clone. After a few minutes of the master image copy process, MCS creates a differential disk and an identity disk for each VM. The size of the created differential disk is the same as the master image to host the session data. The identity disk is normally 16MB and is hidden by default. The identity disk has the machine identity information such as host name and password.

Figure 6 shows the MCS file layout.

Figure 6) MCS file layout.

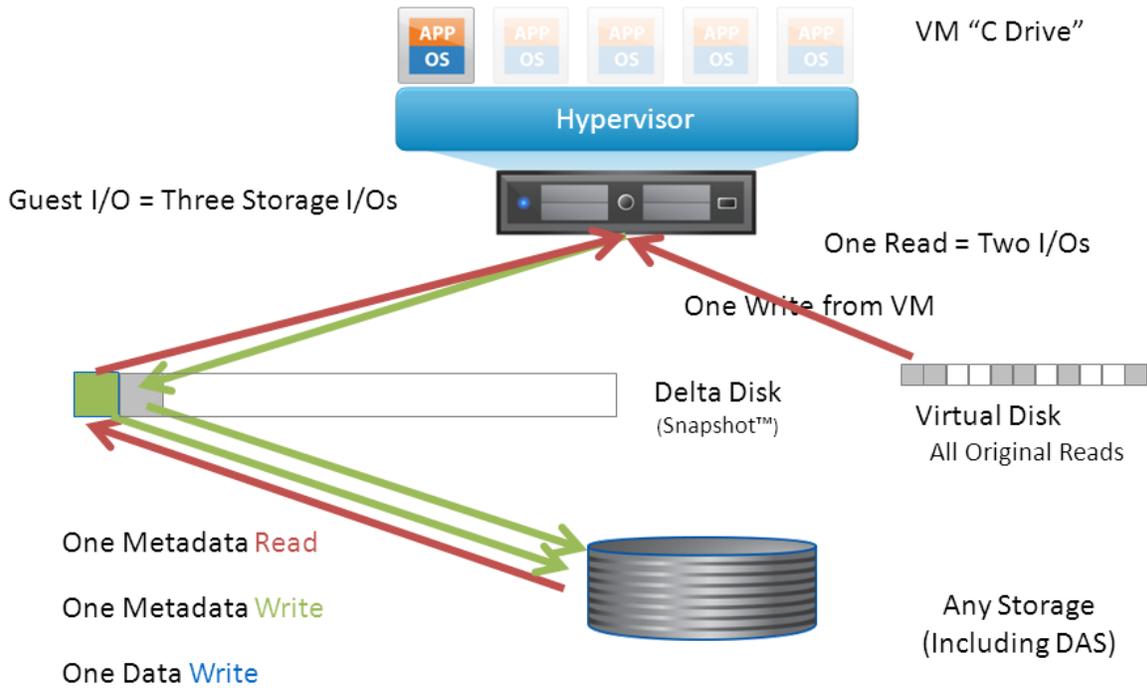
Machine Creation Service (MCS) – How it works



MCS uses hypervisor snapshots similar to VMware[®] linked clone technology. Hypervisor snapshots require one metadata read and one metadata write for each data write operation and one metadata read for each data read operation. Hypervisor snapshot is a costly solution without good scalability and performance.

Figure 7 shows you why MCS uses more I/O than PVS or NetApp clone.

Figure 7) Hypervisor snapshot I/O workflow.

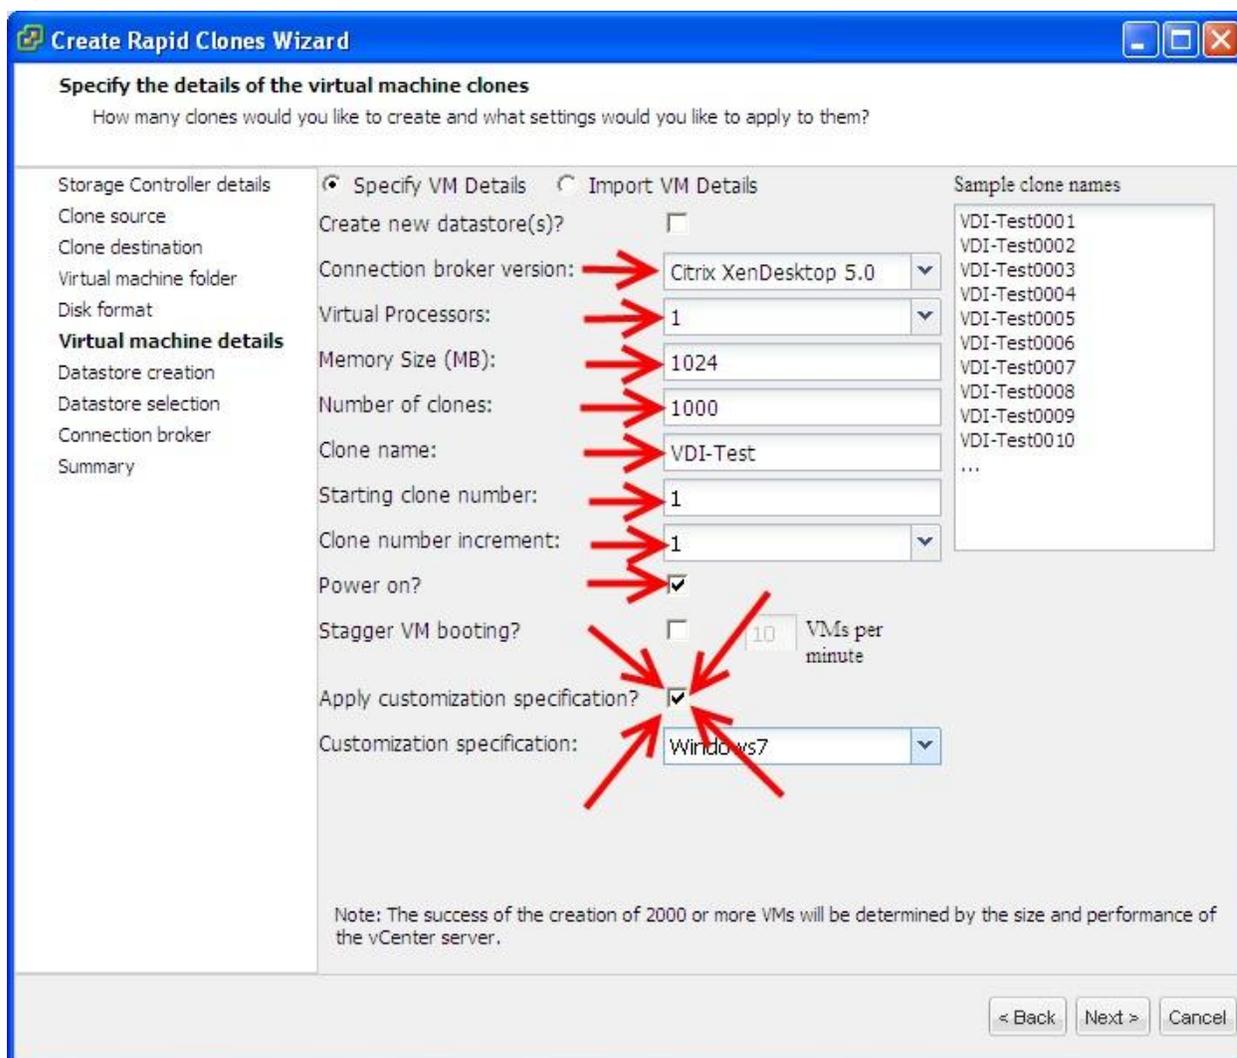


NetApp Clone Created by Virtual Storage Console

You can deploy desktops almost instantly without consuming additional storage by using NetApp FlexClone® technology for cloning individual files and then importing the desktops to XenDesktop.

The VM cloning capability is available with XenServer and VMware vSphere. After cloning a VM, VSC creates a CSV file that includes VM host name and Active Directory® (AD) accounts. XenDesktop can then import the CSV file to create a desktop group.

Figure 8) Virtual machine clone creation.



VSC also has storage management and VM backup and recovery capabilities.

Note: The VM backup and recovery are only available with VMware vSphere. NetApp Virtual Storage Console supports Citrix XenServer and VMware ESX® and VMware ESXi™.

3.4 Provisioning Decision

Table 2 introduces and discusses three provisioning methods in order to enable you to decide which provisioning methods you should use.

Table 2) Provisioning methods.

Provisioning Method	Vendor	Desktop Catalog in XenDesktop	VM Deployment	Supported Hypervisor
Virtual Storage Console	NetApp	Dedicated or pooled	Most storage efficient and fastest VM cloning. Requires NetApp FlexClone license.	ESX and XenServer
MCS machine	Citrix	Dedicated or pooled	VDI only. NFS is preferred protocol.	ESX,

Provisioning Method	Vendor	Desktop Catalog in XenDesktop	VM Deployment	Supported Hypervisor
creation service			Small-scale deployments (less than 2,500 seats based on Citrix consulting) that need to be quickly up and running, MCS needs almost three times more storage than PVS-based deployment. Quick and easy.	XenServer, and Hyper-V
PVS	Citrix	Pooled or dedicated with PvDisk	Proven scalability and storage efficiency. Easy redeployment.	ESX, XenServer, and Hyper-V

Best Practice

For pooled desktops, NetApp recommends PVS for easy redeployment and storage efficiency. For dedicated desktops, NetApp Virtual Storage Console is recommended for fast cloning and storage efficiency.

4 Solution Architecture and Performance Data

NetApp highly recommends implementing VDI layering technologies to separate the various components of a desktop (such as base OS image, user profiles and settings, corporate apps, user-installed apps, and user data) into manageable entities called layers. Layers help to achieve the lowest storage cost per desktop because the storage no longer has to be sized for peak IOPS and intelligent data management policies. For example, storage efficiency and snapshot-based backup and recovery can be applied to the different layers of the desktop.

Some of the key benefits of VDI layering are:

- **Ease of VDI image management.** Individual desktops no longer have to be patched or updated individually. This results in cost savings because the storage array no longer has to be sized for write I/O storms.
- **Efficient data management.** Separating the different desktop components into layers allows for the application of intelligent data management policies (such as deduplication, NetApp Snapshot™ backups, and so on) on different layers as required. For example, you can enable deduplication on storage volumes that host Citrix personal vDisks and user data.
- **Ease of application rollout and updates.** Allows the ease of managing the rollout of new applications and updates to existing applications.
- **Improved end-user experience.** Provides users the freedom to install applications and allow persistence of these applications upon updates to desktop OS or applications.

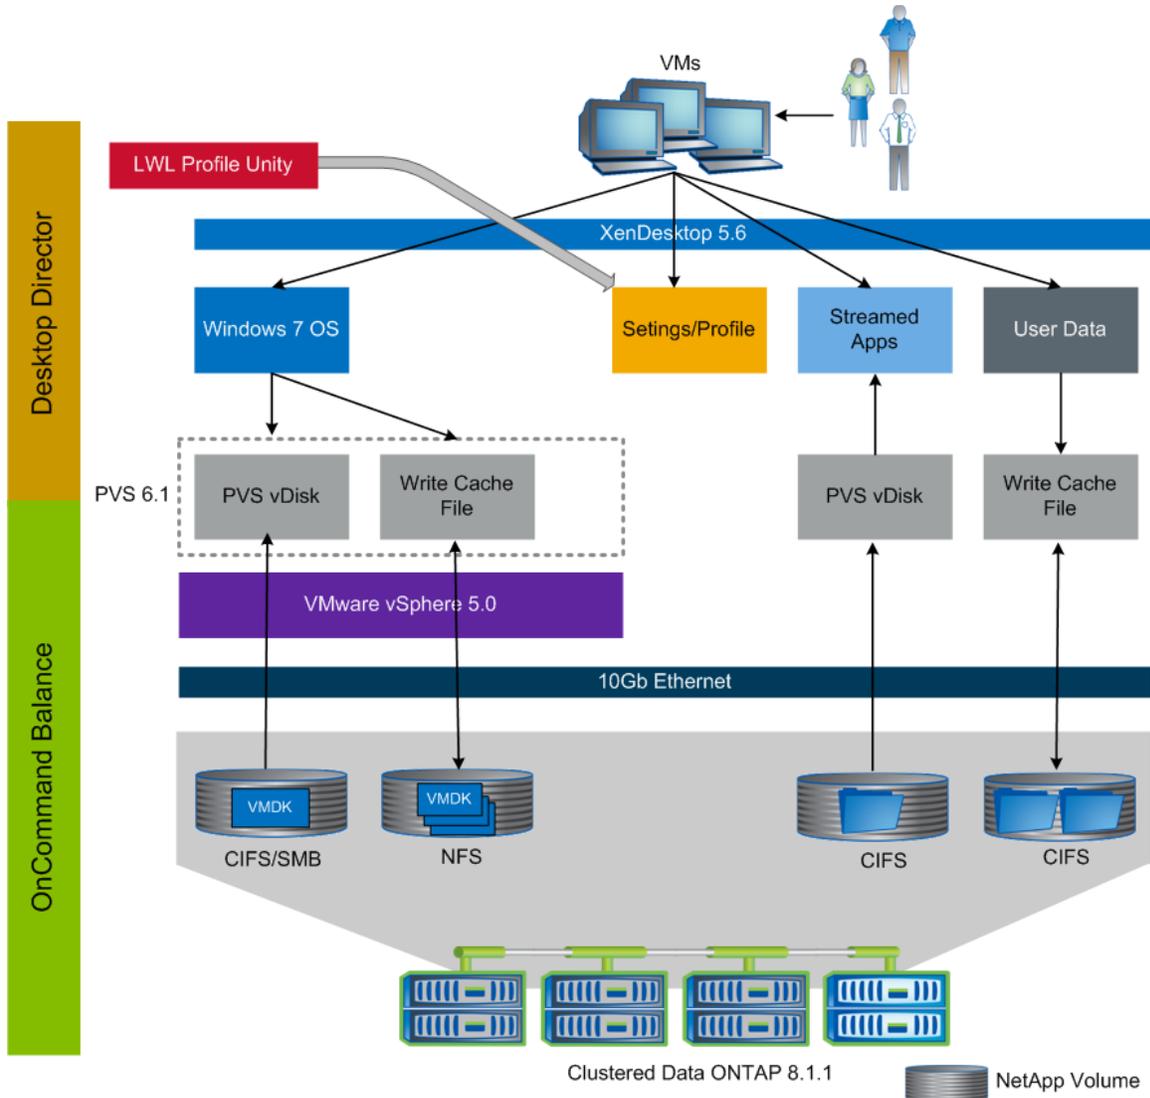
4.1 High-Level Architecture

This section outlines the recommended storage architecture for deploying a mix of various XenDesktop FlexCast delivery models such as hosted VDI and hosted shared desktops, along with intelligent VDI layering (such as profile management and user data management) on the same NetApp clustered Data ONTAP storage array.

Hosted VDI for Pooled Desktops

Figure 9 highlights the NetApp recommended storage architecture for deploying pooled desktops using the Citrix PVS provisioning method.

Figure 9) Architecture for pooled desktops using Citrix PVS.



Base OS image:

- PVS vDisk.** As highlighted in the section “Citrix Provisioning Services (PVS),” CIFS/SMB is the recommended protocol for hosting the PVS vDisk. CIFS/SMB allows the same vDisk to be shared among multiple PVS servers, unlike the hosting of vDisk on FC or iSCSI LUN, which requires one LUN per PVS server. Only one vDisk has to be patched or updated to roll out the updates to thousands of pooled desktops. This results in significant operational savings and architecture simplicity.
- PVS write cache file.** The PVS write cache file is hosted on NFS datastores for simplicity and scalability, as highlighted in the section “Citrix Provisioning Services (PVS).”

Profile management. To make sure that the user profiles and settings are preserved after the desktops are reprovisioned from the updated vDisk (as a result of desktop patch, update, and so on), leverage the

profile management software (Liquidware Labs ProfileUnity, Citrix UPM, and so on) to redirect the user profiles to the CIFS home directories. In our lab validation, Liquidware Labs ProfileUnity was used as the profile management software.

User data management. NetApp recommends hosting the user data on CIFS home directories to preserve data upon VM reboot or redeploy.

Application virtualization. Citrix XenApp can be used to stream baseline applications from a XenApp repository hosted on CIFS shares from NetApp storage so that applications do not have to be individually patch updated on each desktop.

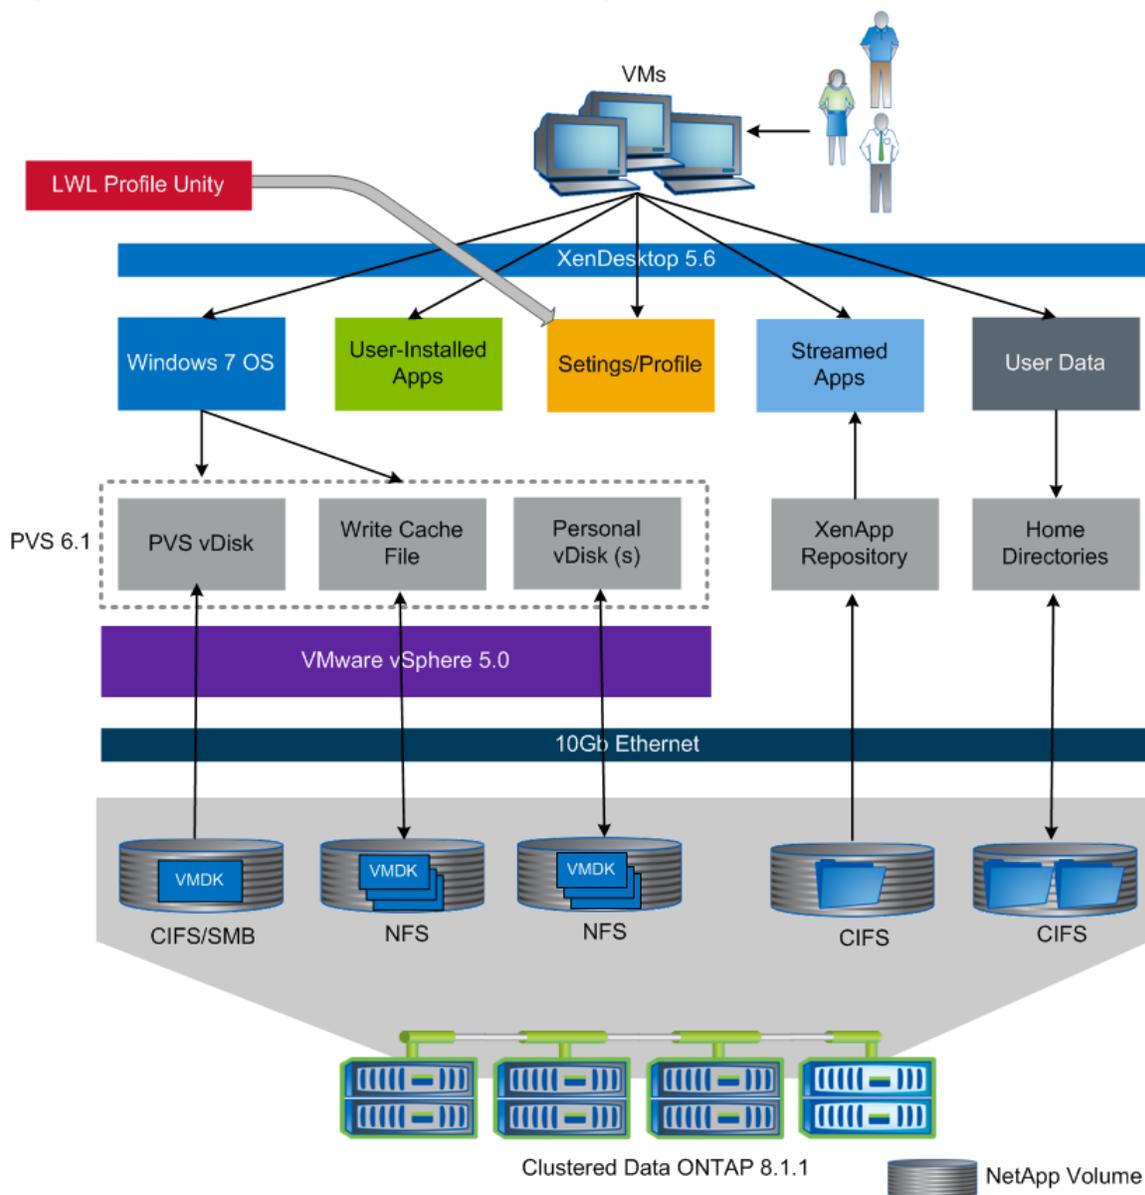
Monitoring and management. NetApp recommends using OnCommand Balance and Citrix Desktop Director to provide end-to-end monitoring and management of the solution.

As shown in Figure 9, all the required storage needs for various desktop data components in the layered VDI architecture (such as PVS vDisks, PVS write cache file, user profiles, user data, and so on) can be easily met by a single NetApp storage array running clustered Data ONTAP. For performance validation and detailed scalability guidelines for this deployment model, refer to section 12.1, “Solution Assessment.”

Hosted VDI for Persistent Desktops

Figure 10 highlights the NetApp recommended storage architecture for persistent desktops that leverages the Citrix PVS provisioning method and Citrix personal vDisk for layering the user-installed applications into a separate managed entity.

Figure 10) Architecture for persistent desktops using Citrix PVS.



For persistent desktop deployment models, the storage design guidelines for PVS vDisk, PVS write cache file, application virtualization, profile management, user data management, solution monitoring, and management are the same as highlighted for “Hosted VDI for Pooled Desktops.”

Personal vDisk: Citrix personal vDisk allows hosting the user-installed applications on a separate virtual disk. NetApp recommends hosting personal vDisks in an NFS datastore for vSphere and XenServer-based deployments. To achieve storage savings, deduplication should be enabled on the datastore that hosts the VMDK/VHD files associated with personal vDisk. Backup and recovery using NetApp Snapshot technology can be easily achieved by using the VSC plug-in.

4.2 XenDesktop on Clustered Data ONTAP Solution Components

Citrix PVS streams multiple desktops to the VMware ESXi 5.0 hypervisor. Citrix XenDesktop manages desktops and facilitates the connections between endpoint devices and desktops. XenApp is used as an application-delivery solution to enable any Windows application to be virtualized, centralized, and

managed. Also, hosted share desktops are delivered by XenApp. Liquidware Lab ProfileUnity and Citrix User Profile Manager manage the user profiles and folder redirection.

Four clustered Data ONTAP nodes are used as storage to host write cache on NFS shares, vDisks, user data, and profiles on CIFS SMB 2.1.

Figure 11 illustrates the architecture of XenDesktop on clustered Data ONTAP.

Figure 11) XenDesktop on clustered Data ONTAP architecture.

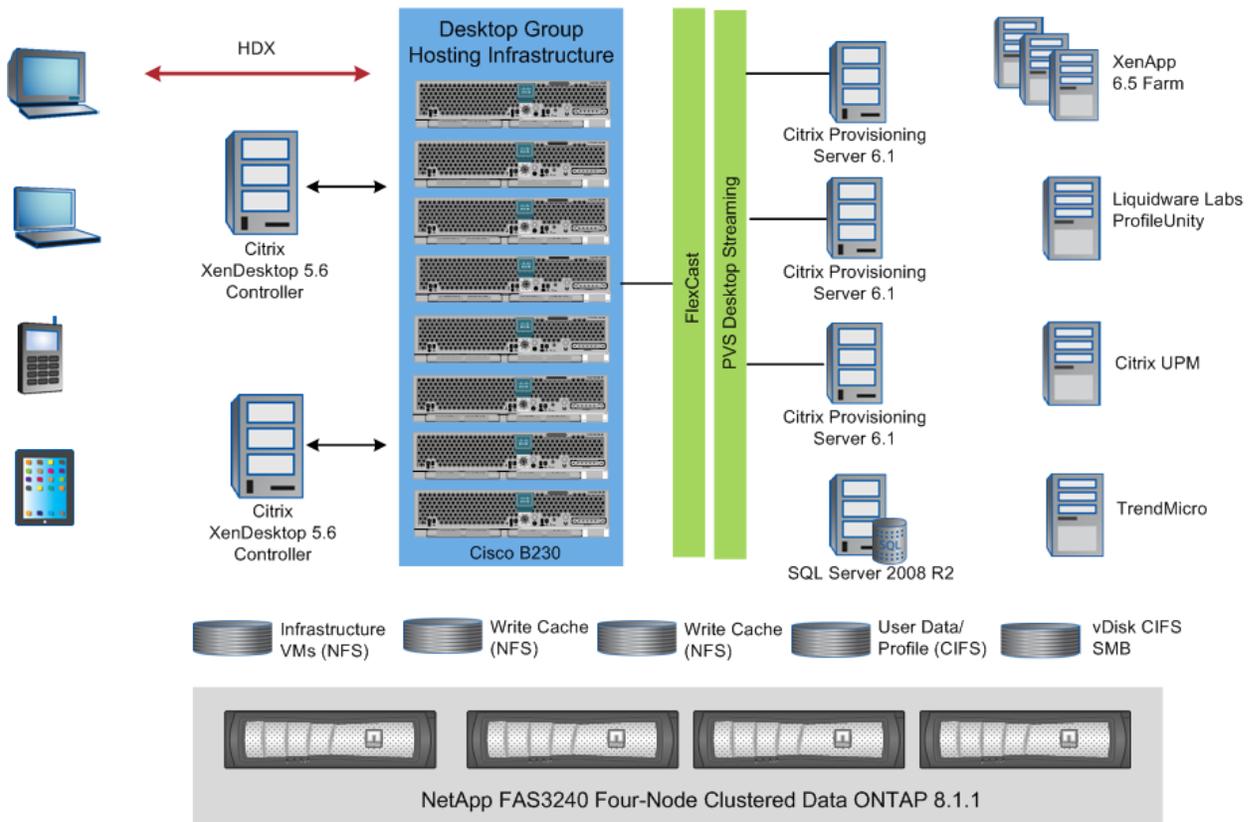


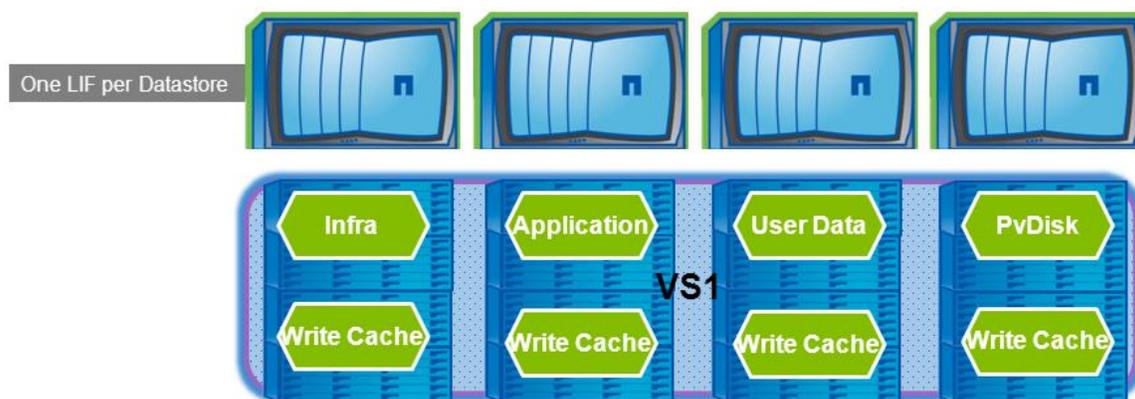
Figure 11 depicts the validated configuration and its components.

- 12 Cisco UCS[®] servers B230 M2 and 5 Fujitsu servers with VMware ESXi 5.0 installed.
- Cisco[®] switches and fabric interconnects and an integrated set of storage and software components.
- FAS3270 four-node cluster running 8.1.1 clustered Data ONTAP. Each node has two shelves of DS2246 configured with 450GB, 10,000 RPM SAS drives, and 256GB of intelligent flash memory cache.

One large aggregate was created on each node; for details, see section 6.8, “Aggregate and Volume.”

Figure 12 shows the various volumes created for this solution.

Figure 12) XenDesktop on clustered Data ONTAP volume layout example.



Best Practice

Perform load balancing on each node in the cluster. Spread the write cache across the controllers because the majority of the workload is from the write cache.

5 Performance Test Results

Extensive testing was conducted in an effort to better understand Citrix PVS and its OS streaming technology and to characterize the workload in detail.

The various NetApp configurations used for the tests included:

- FAS3270 four-node cluster running clustered Data ONTAP 8.1.1
- ESX 5.0 cluster that includes Cisco UCS and Fujitsu blade servers
- 10GbE for NFS and internal network

5.1 Performance Test Tool

Login VSI was used as a workload generator to simulate a multiuser workload and to size for server capacity in VDI environments. Login VSI is a collection of scripts and automation tools that utilizes real user-space applications to generate its workload. These applications include Microsoft® Word, Excel®, Outlook®, and web browsing along with various freeware tools, including a PDF printer and flash video. Random data is generated in each application, and reads, edits, and saves are all performed. For more information about Login VSI, go to www.loginvsi.com.

The Login VSI website defines a “heavy” workload as follows¹:

The heavy workload is higher on memory and CPU consumption because more applications are running in the background.

- *This workload simulated a power user.*
- *Difference in comparison to the medium workload:*
 - *Type rate is 130ms per character.*

¹ <http://www.loginvsi.com/documentation/v3/workloads>.

- Total Idle time is only 40 seconds.
- The heavy workload opens up to eight applications simultaneously.

5.2 Performance Test Plan

Scale testing of up to 500 seats for boot storm, login storm, and steady states was performed. For each test, Login VSI heavy workload was used.

Table 3) Test cases.

Workload	Test Cases
Boot	All 500 virtual machines were booted at once.
Login	The test assumed one user logging in and beginning work every four seconds until the maximum of 500 users was reached, at which point "steady state" was assumed.
Steady state	In the steady state workload all users performed various tasks using Microsoft Office, web browsing, PDF printing, playing flash videos, and using the freeware mind mapper application.

For the boot storm tests, all virtual machines booted within 10 minutes for all configurations.

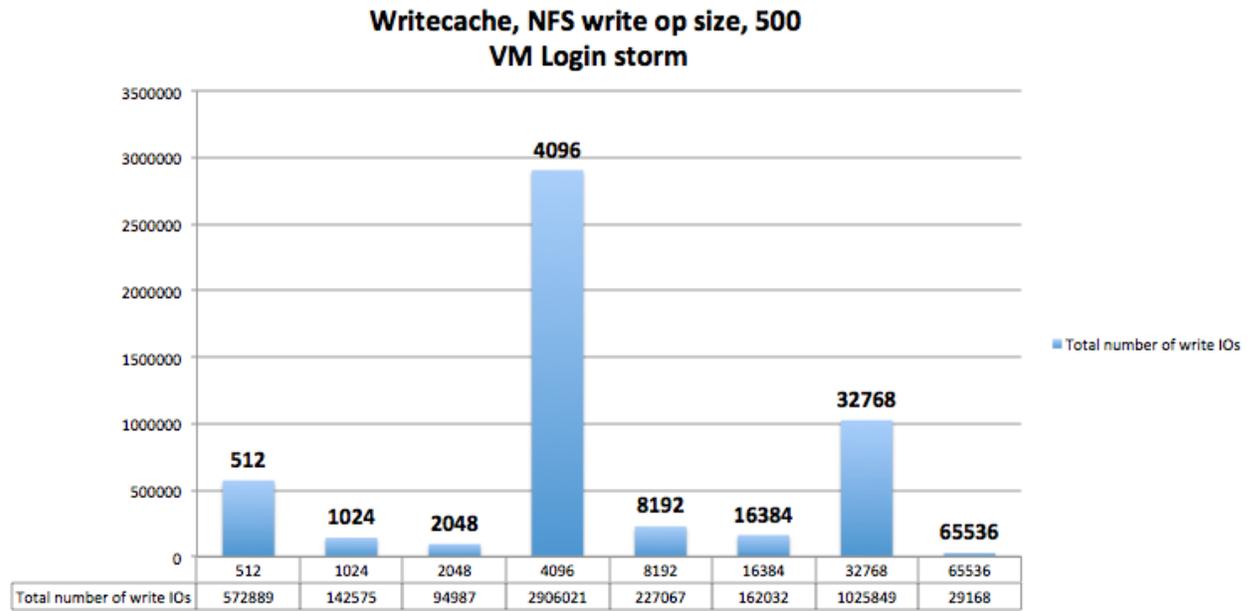
5.3 Citrix PVS Workload Characteristics

Nonpersistent Desktops

In a nonpersistent configuration, the vast majority of workload generated by streamed desktops is to the write cache storage. Comparatively, the read operations constitute very little of the total I/O except at the beginning of the VM boot process. After the initial boot, reads to the OS vDisk are mostly served from the PVS server's cache.

The write portion of the workload includes all the OS and application-level changes that the VMs incur. Figure 13 summarizes the I/O size breakdown of the write workload. The Y axis represents the I/O size range in bytes, and the X axis represents the I/O size.

Figure 13) I/O of the write workload.



Note that the I/O size breakdown in Figure 12 is (for the most part) similar across both persistent and nonpersistent use cases. The main difference in each use case is that the location of the I/O changes with each configuration. For example, eight IOPS per desktop move from the write cache to the personal vDisk when it is added to the configuration.

The entire PVS solution is approximately 90% write from the storage's perspective in all cases, with the size breakdown as shown in Figure 12. The total average I/O size to storage is from 8k through 10k.

The following chart details the IOPS per desktop observed during the Login VSI heavy workload testing of all possible use cases with pooled (nonpersistent) desktops with and without CIFS profile management and user data.

PVS use case I/O summary (IOPs per desktop)	Server	Storage	
	PVS	Writecache	CIFS
Pooled	5	10	N/A
Pooled + Profile mgmt	5	7	7
Pooled + Profile mgmt + home dir	5	7	7 + user sized
Pooled + home dir	5	10	user sized

1. writecache 8k avg, 90% write
2. CIFS 20% read, 25% write, 55% other

In all cases, the PVS server handled the majority of reads to the OS disk, averaging five read operations per desktop.

With nonpersistent pooled desktops, the PVS write cache includes the rest of the I/O. The storage that contains the read-only OS vDisk incurred almost no I/O activity after initial boot and averaged zero IOPS per desktop to the storage (due to the PVS server cache). The write cache showed a peak average of 10 IOPS per desktop during the login storm, with the steady state showing 15% to 20% fewer I/Os in all configurations. The write cache workload op size averaged 8k, with 90% of the workload being writes.

The addition of CIFS profile management had an effect of taking some workload off of the write cache. Login VSI heavy tests showed three IOPS per desktop were removed from write cache and served from

CIFS. The additional four IOPS per desktop seen on the CIFS side were composed of metadata operations (open/close, getattr, lock).

Sizing for CIFS home directories should be done as a separate workload from the virtual desktop workload. The storage resource needs for CIFS home directories will be highly variable and dependent on the needs of the users and applications in the environment.

Persistent Desktops

With the persistent (assigned) desktop configuration utilizing Citrix personal vDisk, the effect of I/O to the storage can be summarized as follows:

- Much of the workload that went to the write cache in a nonpersistent desktop now goes to the personal vDisk. Write cache workload is much lower in a persistent environment.
- Per Citrix, the personal vDisk driver on the guest machine contains a small read cache. This results in 10% to 20% decrease in total operations to the storage over nonpersistent desktops.

The amount of I/O that the personal vDisk incurs will be highly variable and dependent on the applications that are installed on the individual desktops. The I/O is affected by the type of application installed on the personal vDisks in the environment and where that data is stored. As an example, a desktop environment that uses highly I/O-intensive applications or performs copy or compression of large files will see much higher I/O to the personal vDisk storage than an environment that only runs word processing or spreadsheet applications.

In our Login VSI testing most of the applications were installed on the golden image; therefore, the majority of reads were serviced from the PVS server cache instead of the personal vDisk storage. This will naturally give best results in terms of storage resource utilization and efficiency. The user data written by the Login VSI applications was contained mostly on the CIFS user data storage. The personal vDisk handled mostly OS-level changes to the file system.

In a persistent configuration, most of the workload that was serviced by the write cache is now divided among the personal vDisk storage and the CIFS user data if profile management software is in use. One major difference that was observed when using CIFS for user data is that additional overhead of CIFS metadata operations is incurred. This is because of the client's data being contained on file-based storage rather than its block-based vDisk or VMDK.

The following is a breakdown of IOPS per desktop in each of the persistent (assigned) desktop use cases.

	Server	Storage	Pvdisk	CIFS
	PVS	Writecache		
Assigned	5	2	6	
Assigned + Profile mgmt	5	2	5	7
Assigned + Profile mgmt + home dir	5	2	5	7 + user sized
Assigned + home dir	5	2	6	user sized

1. writecache 8k avg, 90% write
2. CIFS 20% read, 25% write, 55% other

The total I/O seen in tests with only personal vDisk enabled shows less total I/O to storage than in the nonpersistent tests. A peak average of eight total IOPS was incurred, with two going to write cache and six being served from the personal vDisk storage. The reason for this is that personal vDisk adds an additional read cache at the guest driver layer, which slightly decreases the amount of I/O served from storage.

Adding CIFS profile management and home directories into the environment takes some I/O from personal vDisk and adds it to CIFS. The additional metadata operations necessary for SMB file access make up the remainder.

After the personal vDisk and CIFS user data is configured in a PVS environment, the write cache is mostly relegated to handling operations to the guest OS page file, and its I/O is greatly reduced. Most of the I/O to the OS and application files will now be served from the PVS server cache, the personal vDisk storage, and CIFS home directories, and sizing and planning should be adjusted accordingly.

Figures 13 and 14 summarize the breakdown of read and write operations to the personal vDisk and write cache storage during boot and login scenarios for persistent desktops. Again, as in all cases, the workload was approximately 90% writes.

Figure 14) 500-seat boot, I/O comparison, PVS with personal vDisk.

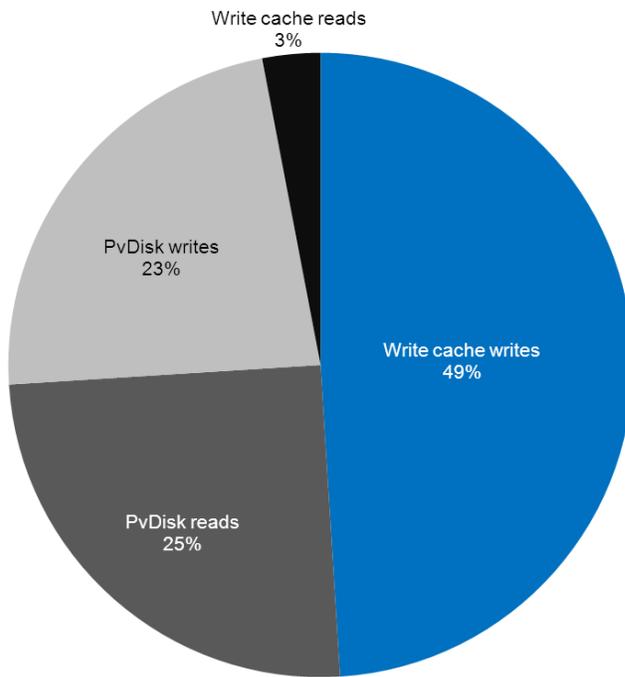
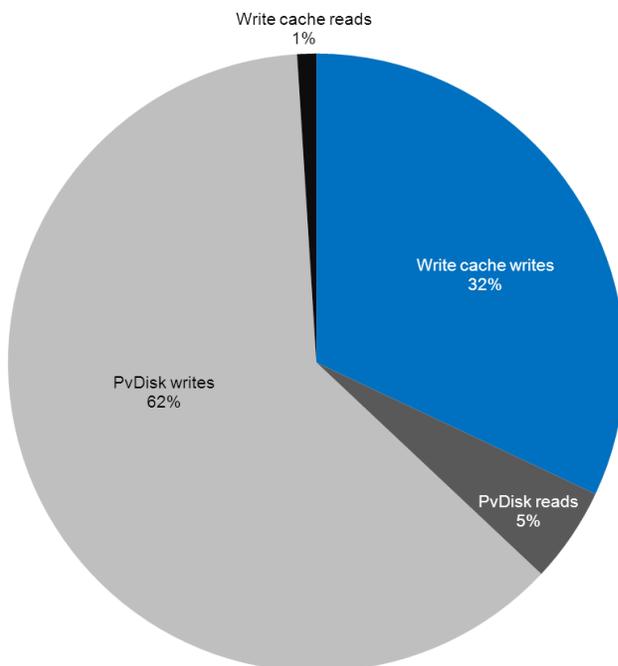


Figure 15) 500-seat login, I/O comparison, PVS with personal vDisk.



5.4 Performance Effect of Flash Solutions

NetApp Flash Cache™ was shown to decrease the total boot time of the environment by 14%. As a whole, Flash Cache improved the read latency on the infrastructure storage and improved the latency on the small amount of rereads to write cache and personal vDisk. Reads to the OS vDisk storage mostly occur only during boot, as again the OS disk is cached by the PVS server, and very few operations were performed to the read-only OS vDisk storage during login and steady-state work. Flash Cache did not significantly affect the write cache workload because it was composed of approximately 90% writes.

6 Storage Configuration Best Practice

NetApp provides a scalable, unified storage and data management solution. The unique benefits of the NetApp solution are:

- **Storage efficiency.** Significant cost savings with multiple levels of storage efficiency for all of the VMs.
- **Performance.** Enhanced user experience with virtual storage tiering (VST) and write I/O optimization that strongly complements NetApp storage efficiency capabilities.
- **Operational agility.** Enhanced Citrix XenDesktop solution management with tight partner integration.
- **Data protection.** Enhanced protection of both the virtual desktop OS data and the user data, with very low overhead for both cost and operations data components.

6.1 Protocol Decision

NetApp has a few recommendations for protocols in clustered Data ONTAP and 7-Mode. The choice of protocol is based on the customers' current infrastructure and best practices.

NFS, iSCSI, and FCP have very similar performance within 7% margin. For more information, refer to [TR-3697: Performance Report: Multiprotocol Performance Test of VMware ESX 3.5 on NetApp Storage Systems](#).

NetApp recommends these protocols primarily based on ease of management and cost efficiency.

Table 4 shows the preferred storage protocols for the deployment of XenDesktop with Provisioning Server. Virtual desktop solution covers delivering OS, personal applications (instant messages, Pandora for music), corporate apps (MS Office), and managing user profiles and user data.

Table 4) Storage protocols for XenDesktop.

	ESXi, ESX, XenServer	Hyper-V	Reason
vDisk	For PVS 6.1 and higher, use CIFS SMB 2.1 and higher. For PVS 6.0 and earlier, use SAN LUN.	CIFS SMB 2.1 and higher	PVS 6.1 changes the Windows registry settings for oplock and CIFS SMB enable; vDisk is read from storage once and cache in PVS RAM. CIFS SMB offers good performance and easy update on vDisk.
Write cache	NFS	SAN LUN	NFS is a space-efficient and easily managed protocol. NFS is default thin provisioning, which optimizes the utilization of available storage. For Hyper-V, spread the VMs across different LUNs. NetApp recommends 150 VMs to 600 VMs per LUN.
Personal vDisk	NFS	SAN LUN	Same reason as write cache.
User	CIFS SMB	CIFS SMB	Best practice is to separate user data from the VM

	ESXi, ESX, XenServer	Hyper-V	Reason
data/profile			OS. This end-user-centric focus allows administrators to deliver a higher quality “desktop-like” experience with data center benefits such as security, compliance, backup and recovery, and disaster recovery.

6.2 Clustered Data ONTAP Failover Group

LIFs and ports have roles; different ports are used for management, storage, data motion, and fault tolerance. Roles include cluster or node management, cluster (for traffic between nodes), intercluster (for SnapMirror® replication to a separate cluster), and data. From a solution perspective, data LIFs are further classified by how they are used by servers and applications and whether they are on private nonroutable networks, corporate internal routable networks, or a DMZ.

The NetApp cluster connects to these various networks by using data ports; the data LIFs must use specific sets of ports on each node for traffic to be routed properly. Some LIFs, such as cluster management and data LIFs for NFS and CIFS, can fail over between ports within the same node or between nodes so that if a cable is unplugged or a node fails, traffic will continue without interruption. Failover groups are used to control to which ports a LIF can fail over. If failover groups are not set up or are set up incorrectly, LIFs can fail over to a port on a wrong network and cause loss of connectivity.

Best Practices

- All data ports should be members of an appropriate failover group.
- All data LIFs should be associated with the appropriate failover group.
- To keep network connectivity as standardized as possible, use the same port on each node for the same purpose.

6.3 One LIF per Volume

NetApp OnCommand System Manager can be used to set up volumes and LIFs. Although LIFs can be created and managed through the command line, this document focuses on the NetApp OnCommand System Manager GUI. Note that System Manager 2.0R1 or later is required to perform these steps.

For good housekeeping purposes, NetApp recommends creating a new LIF whenever a new volume is created. A key feature in clustered Data ONTAP is its ability to move volumes in the same Vserver from one node to another. When you move a volume, make sure that you move the associated LIF as well. This will help keep the virtual cabling neat and prevent indirect I/O, which will occur if the migrated volume does not have an associated LIF to use. It is also a best practice to use the same port on each physical node for the same purpose. Due to the increased functionality in clustered Data ONTAP, more physical cables are necessary and can quickly become an administrative problem if care is not taken when labeling and placing cables. By using the same port on each cluster for the same purpose, you will always know what each port does.

Best Practices

- Each NFS datastore should have a data LIF for every node in the cluster.
- When a volume is moved to a different storage node, a data LIF must be present for that node, or the datastore will go offline.
- When you create a new Vserver, add one LIF per protocol, per node.

6.4 LIF Migrate

LIFs for NFS and CIFS can be migrated, but iSCSI LIFs cannot. If you use Flash Cache, you might need to get the IOXM to add another 10GbE card. If you get IOXM, you'll need to connect the HA pair with a fiber cable.

Best Practice

Use 10GbE for cluster interconnection and data networks. 1GbE is recommended for management networks. Make sure you have enough 10GbE cards and ports.

6.5 10 Gigabit Ethernet

Clustered Data ONTAP requires more cables because of its enhanced functionality. A minimum of three networks are required when creating a new clustered Data ONTAP system: cluster at 10GbE, management at 1GbE, and data at 10GbE. The physical space-saving functionality of the 32XX series can be a potential issue if Flash Cache cards are present in the controllers because 10GbE network cards are required, and 32XX series controllers do not have onboard 10GbE NICs. Check with NetApp sales engineers if this is a system will be migrated from 7-Mode to clustered Data ONTAP to make sure enough open slots are available on your existing controllers.

Best Practice

10GbE is not required for data, but is a best practice for optimal data traffic.

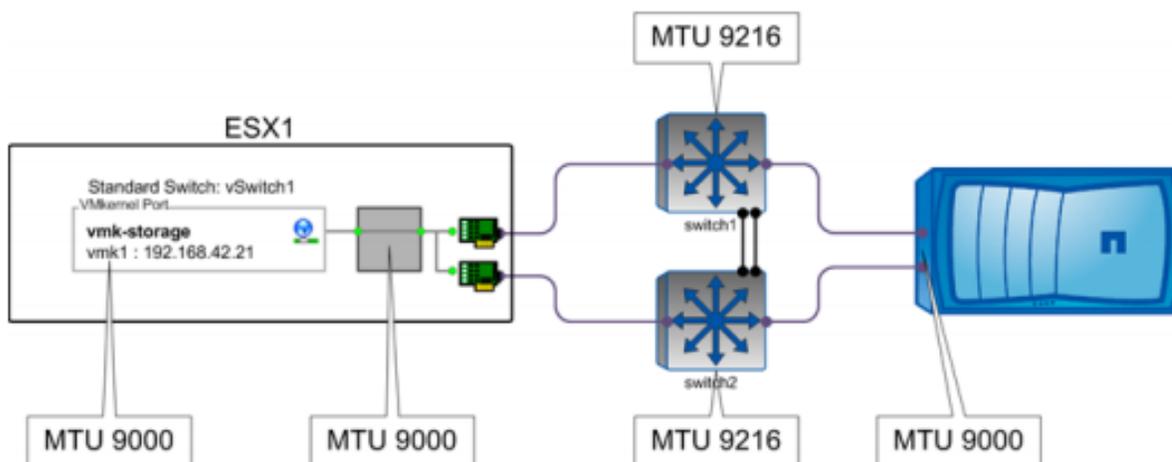
6.6 Jumbo Frames

Enable jumbo frames for data network and Link Aggregation Control Protocol (LACP) if the switch supports it.

NetApp recommends using jumbo frames or MTU 9000 for the data network. Enabling jumbo frames can help speed up data traffic and use less CPU because fewer frames are being sent over the network. Jumbo frames must be enabled on all physical devices and logical entities from end to end to avoid truncation or fragmentation of packets with the maximum size. Link aggregation type will depend largely on the current switching technology used in the environment. Not all link aggregation implementations are alike or offer the same features. Some only offer failover from one link to another in the event of failure of one link. More complete solutions offer actual aggregation in which traffic can flow on two or more links at the same time. There is also LACP, which allows devices to negotiate the configuration of ports into bundles. For more information regarding the setup of link aggregation with VMware ESXi and clustered Data ONTAP, refer to [TR-4068: VMware vSphere 5 on NetApp Data ONTAP 8.1 Operating in Cluster-Mode](#).

Figure 16 shows you how to set jumbo frames on each network component.

Figure 16) Setting jumbo frames.



Best Practice

NetApp recommends using jumbo frames or MTU 9000 for the data network.

6.7 Flow Control Overview

Modern network equipment and protocols generally handle port congestion better than in the past. Although NetApp had previously recommended flow control “send” on ESX hosts and NetApp storage controllers, the current recommendation, especially with 10GbE equipment, is to disable flow control on ESXi, NetApp FAS, and the switches in between.

With ESXi 5, flow control is not exposed in the vSphere client GUI. The `ethtool` command sets flow control on a per-interface basis. There are three options for flow control: `autoneg`, `tx`, and `rx`. The `tx` option is equivalent to “send” on other devices.

6.8 Aggregate and Volume

An aggregate is NetApp’s virtualization layer that abstracts physical disks from logical datasets, which are referred to as flexible volumes. Aggregates are the means by which the total IOPS available to all of the physical disks are pooled as a resource. This design is well suited to meet the needs of an unpredictable and mixed workload. NetApp recommends using a small aggregate as the root aggregate, whenever possible. In most cases, a default of three disks for the root aggregate was used. This aggregate stores the files that are required for running and providing GUI management tools for the storage system. The remaining storage should be placed into a smaller number of larger aggregates. The maximum size of aggregate depends on the system; for details, refer to [NetApp configuration limits](#) described in [TR-3838: Storage Subsystem Configuration Guide](#).

Clustered Data ONTAP supports a maximum volume size of 100TB.

For practical purposes, NFS does not have a limit to the number of VMs per datastore. Also, the NFS protocol does not limit the addressable size of an NFS datastore, which means it automatically supports the current 100TB maximum volume size of clustered Data ONTAP. With this in mind, VMs can be grouped according to business requirements, which can be organizational (department, tenant, or application) or service level based, such as the type of storage, replication requirements, or schedule.

6.9 Secure Default Export Policy

Clustered Data ONTAP changes the architecture of exports by defining export policies scoped within a Vserver. An export policy has a set of rules that includes which clients get what type of access. Each

volume has exactly one export policy applied to it. All volumes that are used by the vSphere environment, or at least for each VMware cluster, can use the same export policy so that all hosts see the set of volumes in the same manner. When a new host is added to the vSphere cluster, a rule for that host is added to the policy. The rule for the new host includes a client-match pattern (or simply an IP address) for the new host, protocol, and permissions and authentication methods for read/write, read-only, superuser, anonymous user, and some other options that are of lesser significance for vSphere. When a new volume is added and other volumes are already in use as datastores, the same export policy can be used, as for the previous volumes. For more information on clustered Data ONTAP export policies and junction paths, refer to [TR-4068: VMware vSphere 5 on NetApp Data ONTAP 8.1 Operating in Cluster-Mode](#).

6.10 Avoid Misalignment

Multiple layers of storage are involved. Each layer is organized into blocks or chunks for efficient accessibility of the storage. The size and the starting offset of each block can be different at each layer. Although a different block size across the storage layers doesn't require any special attention, the starting offset does. For optimal performance, the starting offset of a file system should align with the start of a block in the next lower layer of storage. For example, an NTFS file system that resides on a LUN should have an offset that is divisible by the block size of the storage array presenting the LUN. Misalignment of block boundaries at any one of these storage layers can result in performance degradation. Misalignment causes storage to read from or write to more blocks than necessary to perform logical I/O. For more detail on misalignment, check [Best Practices for File System Alignment in Virtual Environments](#).

In XenDesktop deployment, depending on vDisk OS type, Windows 7 and Windows 2008 do not have an OS alignment issue. Windows XP and Windows 2003 Server have OS alignment issues.

Because the VMFS layer is not involved with NFS, only the alignment of the guest VM file system within the VMDK to the NetApp storage array is required. There is no misalignment if the write cache file is on an NFS share. If a SAN LUN is chosen to use host write cache, use the following procedure to correct misalignment.

Windows 7 and Windows 2008 do not have an OS alignment issue.

When you create a PVS write cache file for the VM template, it is a best practice to mount the LUN on an existing Windows server and run DiskPart.

1. Mount the write cache file disk on Windows server.
2. Click Start > Run, and enter `diskpart`.
3. Enter `Select Disk#`.
4. Choose the write cache file disk number.
5. Enter:

```
create partition primary align=32.
```

6. Unmount the write cache disk from the Windows server.
7. Attach the disk to the PVS VM template.

Figure 17 shows how to use DiskPart.

Figure 17) Diskpart.

```
C:\Documents and Settings\Administrator>diskpart
Microsoft DiskPart version 5.2.3798.3959
Copyright (C) 1999-2001 Microsoft Corporation.
On computer: VIRTUAL_CENTER1

DISKPART> list disk

   Disk ###  Status         Size           Free           Dyn  Gpt
   -----  -
   Disk 0    Online          20 GB          8033 KB
   Disk 1    Online          5114 MB        5114 MB

DISKPART> select disk 1
Disk 1 is now the selected disk.

DISKPART> create partition primary align=32
DiskPart succeeded in creating the specified partition.
DISKPART>
```

7 Storage Operation Best Practice

7.1 Deduplication Decision

NetApp can help you use less storage space. Citrix XenDesktop environments can benefit from the cost savings associated with NetApp deduplication, as discussed previously. Each VM consumes storage as new writes occur. Scheduling and monitoring deduplication operations for the NetApp volumes hosting VMs are very important. It is important to schedule the deduplication operations to run during off-peak hours so that the end-user experience is not affected. Also, it is important to understand the number of simultaneous dedupe operations that can be performed on the storage controller. Planning for dedupe operations ultimately depends on your environment. The status and storage savings of dedupe operations can be monitored by using the NetApp System Manager, VSC deduplication Management tab, and Virtual Storage Console. For details on NetApp deduplication, refer to [NetApp TR-3505: NetApp Deduplication for FAS, Deployment and Implementation Guide](#).

Unlock the power of deduplication and optimize space utilization on the user data, personal vDisk, infrastructure, and vDisk volumes.

Table 5 summarizes NetApp’s recommendations for deduplication in the XenDesktop environment.

Table 5) Dedupe recommendations.

	Deduplication Recommendation	Reason
vDisk	Yes	OS data can be deduplicated.
Write cache	No	Log file and page file are transient data and are only temporary, and therefore, it is important to not enable deduplication on these volumes. Doing so would result in utilizing storage resources for no reason.
Personal vDisk	Yes	Use dedupe on the same applications between the users.
User data/profile	Yes	Use dedupe on user data and profile.

7.2 Thin Provisioning and Volume Autogrow Policies

NetApp recommends using thin provisioning for write cache and PvDisk.

NetApp thin provisioning is great way to overprovision storage on volumes that don't necessarily need all of their space immediately. To avoid potential problems if the volume is quickly filled up, configure warning thresholds and autogrow policies. When these two processes are initiated, volumes will approach, but never hit 100%, which can potentially bring production to a halt.

7.3 Space Reclamation

When customers deploy a virtual desktop infrastructure using NFS, they can maintain storage efficiency of thin-provisioned virtual machines by using Virtual Storage Console 2.1.1.

Space reclamation requires that the following conditions be met:

- NFS only
- NTFS on basic disks only (GPT or MBR partitions)
- Data ONTAP 7.3.4 or higher
- Data ONTAP 8.0 or higher operating in 7-Mode
- Virtual machine powered off
- No virtual machine VMware snapshots

7.4 Antivirus

A successful VDI implementation begins with change. Using traditional desktop methods to install applications, back up data, and protect systems from viruses will only lead to performance problems and administrative headache. Antivirus companies such as Trend Micro, McAfee, and Sophos allow you to protect your user desktops and avoid bottlenecks that can cripple the performance of your desktop environment. For more information, see:

- <http://www.trendmicro.com/us/enterprise/product-security/index.html>
- <http://www.sophos.com/en-us/products/endpoint/sophos-anti-virus.aspx>
- <http://www.mcafee.com/us/products/move-anti-virus.aspx>

7.5 Monitoring NetApp and XenDesktop Infrastructure

NetApp Operations Manager

Operations Manager delivers complete NetApp infrastructure monitoring for optimal storage efficiency and availability.

As data grows in volume and complexity, you need a complete, current depiction of your storage resources. The central console of NetApp Operations Manager software provides a comprehensive monitoring and management dashboard for instant analysis of your storage infrastructure.

Use the detailed visibility provided by NetApp Operations Manager, part of the OnCommand management software family, to:

- Get intelligent, comprehensive reports of storage utilization and trend information to support capacity planning, space usage, and backup allocation.
- See how physical resources are affected by data utilization and deduplication.
- Monitor system performance and health to resolve potential problems before they occur.
- Deploy, provision, and manage your complete enterprise storage network from a central location.

With its comprehensive array of alerts, reports, and performance and configuration tools, NetApp Operations Manager helps keep your storage infrastructure aligned with your business requirements to improve capacity planning and utilization and increase IT efficiency.

For more details on Operations Manager, visit the [Operations Manager solutions page](#).

NetApp eSupport

The NetApp proactive eSupport suite provides an early warning system that can reduce the number and severity of technical support cases. Automation tools identify issues early, before they have an effect, and can initiate a fix without customer burden, before people even know there is a potential problem. Support automation works 24/7 to benchmark system status, collect diagnostic information behind the scenes, and issue proactive alerts. You can view the full scope of your NetApp environment on demand at the company or device level.

The NetApp eSupport suite of support automation tools includes:

- NetApp Remote Support Diagnostics Tool
- NetApp Premium AutoSupport™
- NetApp AutoSupport

For more information about NetApp eSupport, visit [NetApp Global Support](#).

NetApp OnCommand Balance

NetApp OnCommand Balance provides directions to optimize the performance and capacity of your virtual and physical data center infrastructure, stepping you closer to the virtual cloud.

Use OnCommand Balance to:

- Manage virtual infrastructure end to end from virtual and physical servers to storage disks.
- Reduce troubleshooting time by 80%.
- Optimize VM densities.
- Improve server and storage utilization.
- Manage service levels.

For more information on OnCommand Insight Balance, visit [Insight Balance Library](#).

8 Backup and Recovery of the Virtual Desktops in a VMware vSphere Environment

As virtualized environments scale and data increase, having one standardized backup, recovery, and disaster recovery solution becomes a necessity. Solutions such as VSC for vSphere provide administrators with the ability to manage their entire data protection environment with the same procedures and tools independently of the applications to be backed up. As new nodes, Vservers, VMs, and applications are added, lightweight agents are all that are needed to enable backup, recovery, and disaster recovery.

The PVS write cache files are transient; therefore, they are discarded every time the VM reboots. VSC backup and recovery integrate VMware snapshots with the NetApp array-based block-level Snapshot copies to provide consistent backups for the virtual desktops. It is NetApp primary storage data deduplication aware and also integrates with NetApp SnapMirror replication technology, which preserves the storage savings across the source and destination storage arrays. It is not required to rerun dedupe on the destination storage array. The backup and recovery plug-in also provides a user-friendly GUI for the data protection schemes.

Snapshot technology is built into fabric of clustered Data ONTAP, which allows for instantaneous backups of data in each storage volume. Although this technology is great and should be used for volumes such as user data, personal vDisk, infrastructure, and vDisk snapshots, Snapshot reserve should be disabled on the write cache volumes. Because these volumes contain only temporary transient data, there is no need to reserve space for backups or back up these volumes. By doing this, the maximum amount of space can be used that is assigned to these volumes. In contrast, Snapshot reserve should be enabled for vDisk, Pvdisk, and user data volumes.

For XenDesktop disaster and recovery, refer to [TR-3931: DR Solution for XenDesktop on NetApp](#).

Table 6 demonstrates the various pieces of a Citrix XenDesktop PVS implementation and the backup methods recommended for each piece.

Table 6) XenDesktop on vSphere backup methods.

	Snapshot Reserve	Backup Method
Write cache	Off	Temporary data; does not require backup
vDisk	On	NetApp Snapshot
Personal vDisk	On	VSC for vSphere
User data	On	NetApp Snapshot

Information regarding the installation and basic usage of NetApp VSC can be found in [Virtual Storage Console 4.1 for VMware vSphere Installation and Administration Guide](#).

VSC Backup Best Practices for XenDesktop Pvdisk

The method employed to back up and recover your data depends on how you architected your virtual machine storage in XenDesktop at the host level. See the [XenDesktop Deployment Guide](#) for setup and deployment instructions.

Although various methods can be used to back up the VDI environment, the easiest way is to do so is by using a VSC tool. Multiple machines can be selected for immediate scheduled backups. Because the base image of a PVS desktop is read only, a backup is not required, but it critical to back up the Pvdisk. When creating the backup policy, make sure that the “Include datastores with independent disks” option is selected. Because Pvdisks are normally on separate datastores, if this option is not selected, it will not be properly backed up.

If SnapMirror between volumes has been set up, VSC can trigger an update.

Backup schedules will vary greatly and can be created per customer requirements.

VSC Recovery Best Practices for XenDesktop Pvdisk

A restore of a virtual machine can occur only if a backup has been created and the NetApp Snapshot copy still exists. Multiple restore options are available within VSC and should be selected based on your needs. It is important to note that while restoring, to verify individual Pvdisk to make sure it was properly backed up. If Pvdisk was stored on a separate datastore and not backed up, the user’s applications and data cannot be restored if it was not selectively backed up. A virtual machine can be restored entirely over the original, or a Snapshot copy of the datastore can be mounted to an ESXi server and the virtual machine’s drives mounted to another machine for individual file restore.

For more information, see [the Backup and Recovery modules for VSC](#).

9 User Data and Profile Management

9.1 User Data

The first type of user data is the end-user data or home directory data. This data is the intellectual property of each company and is directly end user generated. In a virtual desktop environment, the home directory data located in a NetApp volume is shared through the CIFS protocol and is mapped as a drive letter in the virtual desktop. This data often requires backup and recovery as well as disaster recovery. Using CIFS home directory brings more efficiency in the management and protection of user data. End-user data files should be deduplicated and compressed to achieve storage efficiency to reduce the overall solution cost.

Best Practice

NetApp recommends using deduplication and compressing end-user data files stored in home directories to obtain storage efficiency. NetApp strongly recommends storing user data on the CIFS home directory.

9.2 User Profile Data

The second type of user data is the user profile, otherwise known as the personal data. This data allows the user to have a customized virtual desktop environment when using a nonpersistent virtual desktop. The user profile is normally stored in C:\users in a Microsoft Windows physical machine and can be redirected to a CIFS share for use in a nonpersistent virtual desktop environment. There are many profile management solutions in the market that simplify management, improve reliability, and reduce network requirement over standard Windows folder redirection. Using a profile management solution will help speed the migration process from a physical desktop or laptop by first virtualizing the profile and then virtualizing the desktop. This will improve login times over folder redirection alone and centralize the end user's profile for better backup, recovery, and disaster recovery of this data.

Best Practice

NetApp recommends using a profile management solution such as Citrix User Profile Management (UPM) or Liquidware Labs ProfileUnity to allow end users to customize their experience in a nonpersistent desktop environment.

Profile Management

The ability to save user desktop customization is paramount when transitioning traditional desktop users to a VDI environment. For cases in which read-only desktop technologies such as PVS are used, the ability becomes even greater to allow successful adoption of the new desktop.

Various profile redirection software products are available. Liquidware Labs ProfileUnity and Citrix User Profile Management were used to test profile redirection in functionality and in performance testing. Overall, the I/O characteristics to storage are roughly similar to Liquidware Labs and Citrix UPM. Because both applications synchronize user profiles and data folders to CIFS and perform roughly equivalent tasks, it would be expected that I/O characteristics would be similar.

It should be noted that the Login VSI user profiles were not very large, so the results might vary depending on user profile size. Tests performed by Citrix show the effect of the number of files in a profile on user login time. This data is for a single user with varying number and size of files.

More information about the setup and deployment can be found on the [Liquidware Labs](#) homepage. For more information about Citrix UPM, go to <http://support.citrix.com/product/upm/v2.1/ - tab-doc>.

For further investigation regarding many of the existing products as well as their capabilities, see the UEM Smackdown located at

<http://www.pqr.com/images/stories/Downloads/whitepapers/user%20environment%20management%20uem%20smackdown.pdf>.

Best Practice

Create a CIFS share for user data and profile redirection. Take care when choosing the drive types if profile redirection software is used because a slower drive might slow down the user login time during login storms.

9.3 User-Installed Applications

The third type of user data is user-installed application data. Some companies allow end users to install applications on their physical desktops and laptops. In a nonpersistent virtual desktop environment, this can now be done with the use of a user-installed application solution. This allows the end user to install applications in a nonpersistent virtual desktop environment by redirecting the data to a CIFS share or persistent virtual disk.

Separating User Data from the VM

Because each data type of a virtual desktop has a different purpose, each data type also has different requirements. Separating the different components of a virtual desktop allows the administrator to create a stateless or nonpersistent virtual desktop environment. This deemphasizes the importance of the operating system and corporate installed applications within the virtual machine and allows administrators to focus on the end user, their experience, and their data. This end-user-centric focus allows administrators to deliver a higher quality “desktop-like” experience with data center benefits such as security, compliance, backup and recovery, and disaster recovery.

Figure 18 shows the data types organized by volume.

Figure 18) Separation of volumes based on data types.



10 Application Virtualization with XenApp

Different types of workers across the enterprise have various performance and personalization requirements. Some require simplicity and standardization, while others require high performance or a fully personalized desktop. IT administrators can deliver every type of virtual desktop, hosted or local, physical or virtual—each specifically tailored to meet the performance, security, and flexibility requirements of each individual user—with NetApp storage and Citrix desktop virtualization products. This section discusses the application virtualization capabilities of XenApp to deliver on-demand applications to virtual desktops and how to leverage the power of NetApp technology to power and simplify this enterprise application delivery software.

Application delivery consists of the three following use cases:

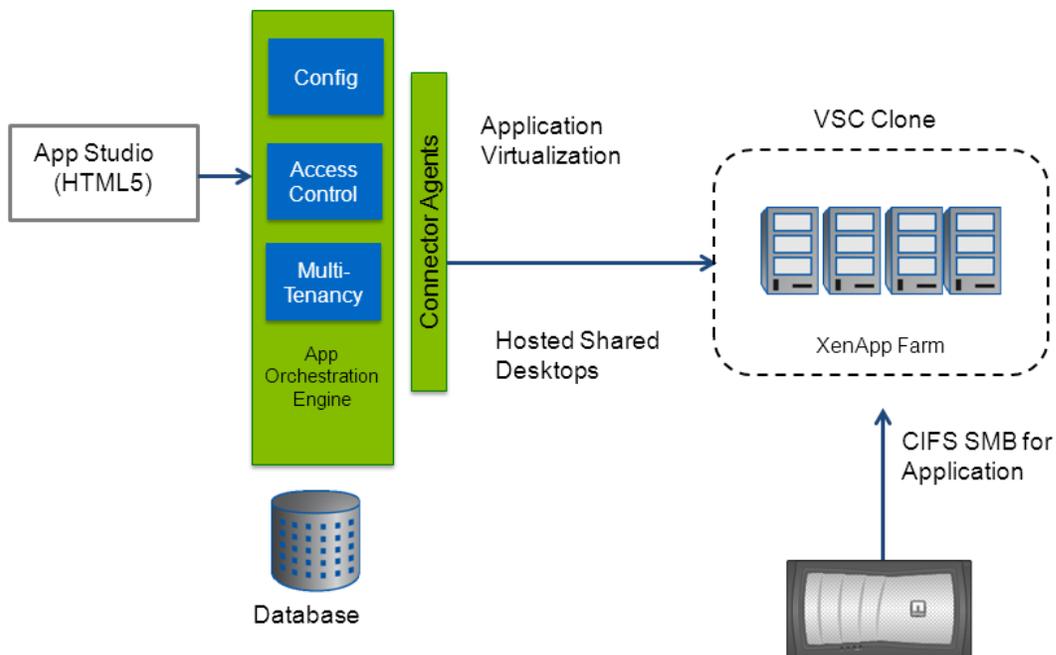
- **Server-side application virtualization.** These applications run inside the data center. XenApp presents each application interface on the user device and relays user actions from the device, such as keystrokes and mouse actions, back to the application.
- **Client-side application virtualization.** XenApp streams applications on demand to the user device from the data center and runs the application on the user device.

- **VM-hosted application virtualization.** Problematic applications or those requiring specific operating systems run inside a desktop on the data center. XenApp presents each application interface on the user device and relays user actions from the device, such as keystrokes and mouse actions, back to the application.

Server-side application virtualization was chosen as the preferred method for these tests, but other methods can be used as well. A single virtual XenApp server was created to run the Citrix AppCenter, and a virtual golden image XenApp server was created to create clones. In this architecture, the additional XenApp servers are identical; they are used to power the user demand for applications, and are housed in an NFS datastore. NetApp VSC is used to rapidly clone VMs and utilizes only small amounts of storage because of FlexClone technology. A customization specification file is used to add the new XenApp servers to Active Directory, and upon boot the new servers are added to the XenApp farm. After selecting an application, the Citrix Streaming Profiler is run, which in turn runs a virtual install and captures the application in a profile. It is recommended to place this profile file in a NetApp CIFS share so that it can be shared among the multiple virtual XenApp servers. When patches or changes need to be applied to the XenApp servers, the original and golden image are modified. The Redeploy function in VSC is then used to quickly update the XenApp clones to apply the new updates. For more information on XenApp, visit <http://support.citrix.com/proddocs/topic/xenapp/ps-library-wrapper.html>.

Figure 19 illustrates VSC's cloning capability in the XenApp solution.

Figure 19) VSC cloning.



11 Citrix Components Scalability Consideration

11.1 Infrastructure Components Limit

While scaling a large virtual desktop deployment, you must adhere to each XenDesktop component's configuration best practices and limits.

Table 7 shows each XenDesktop component's recommended configuration and limits. The data is from Citrix Synergy Citrix consulting group slides.

Table 7) XenDesktop component configuration and limits.

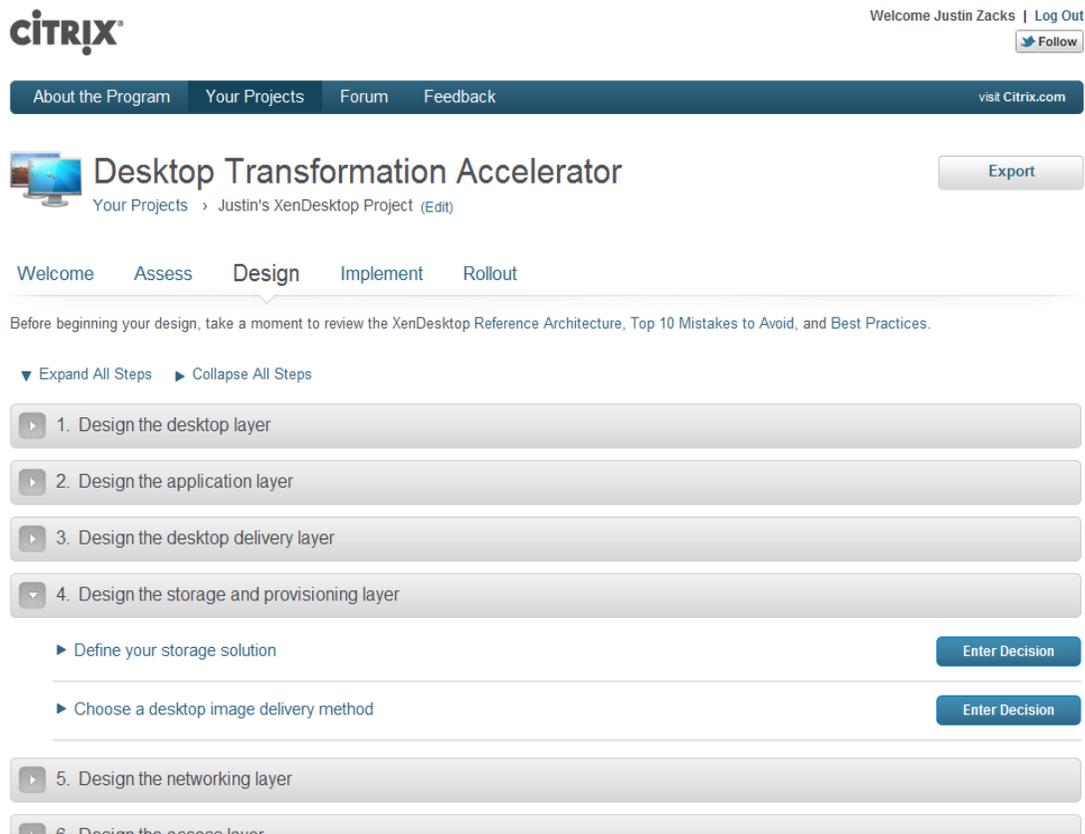
	Limit	Configuration
XenDesktop	<ul style="list-style-type: none"> One XenDesktop 5.6 supports 10,000 virtual desktops. For high-availability purposes, you need N+1 XenDesktop servers. NetApp recommends using VMs for XenDesktop server. 	A Windows Server 2008 R2 VM, 4 vCPUs, 4GB RAM, 40GB storage
PVS	<ul style="list-style-type: none"> One virtual Citrix PVS server can stream approximately 300–500 VMs N+1 	4 vCPUs, 8GB RAM RAM depends on the number of vDisks. 4GB RAM plus 2GB RAM multiplied with the vDisk number.
XenApp controller	<ul style="list-style-type: none"> One XenApp controller serves 1,000 users; for high availability, you require N+1 One-zone data collector, one XML broker server, 2,000 users Two-zone data collectors, two XML broker server, 2,000–5,000 users Two-zone data collectors and three XML broker servers, 5,000–10,000 users 	A Windows 2008 R2 VM, 4 vCPUs, 4GB RAM, 40GB storage
Licensing server	<ul style="list-style-type: none"> Single-license server is generally sufficient Limit of 170 license checkouts per second or 600,000 every hour 	2 vCPUs, 4GB RAM
Web interface	<ul style="list-style-type: none"> 30,000 connections per hour N+1 NetScaler for load balancing 	2 vCPUs, 4GB RAM
SQL Server®	<p>Two vCPUs, 4GB RAM support 2,500–10,000 users Four vCPUs, 16GB RAM support 10,000 users NetApp recommends using SQL Server mirror with witness.</p> <ul style="list-style-type: none"> Database and transaction log sizing is a significant factor Database is relatively static, but transaction logs grow rapidly for hosted VDI Transaction logs should be backed up and cleared regularly 	2 vCPUs, 4GB RAM

Desktop Transformation Accelerator

This free, customized, step-by-step implementation guidance online tool is developed by Citrix. For more information, go to <http://project.citrix.com/>.

- **Guided assess and design.** Leverage a detailed checklist of decisions, options, and tools
- **Peer benchmarks.** Compare your design against thousands of others
- **Best practices.** Use Citrix best practices to make informed decisions
- **Documentation.** Jump start your project by downloading a project plan and other documents
- **Step-by-step installation.** Get custom-generated installation instructions and resources

Figure 20) Desktop transformation accelerator.



11.2 Hypervisor Configuration Limits and Pod Concept

Carefully review the VMware documentation on configuration maximums associated with the various storage-related parameters critical to the system design. For vSphere, review the VMware configuration [vSphere 5.1 configuration maximum](#); for XenServer, review [XenServer 6 Configuration limits](#).

These configuration parameters should help determine the following design parameters:

- Proposed number of VMs per hypervisor host
- Proposed number of hypervisor hosts per hypervisor cluster
- Proposed number of datastores (storage repository) per hypervisor cluster
- Proposed number of VMs per hypervisor cluster
- Number of hypervisor clusters managed by a vCenter™ or XenCenter instance
- Proposed number of VMs per datastore (storage repository)
- Total number of datastores required for the project

Provisioning fewer, denser datastores provides key advantages such as the ease of system administration, solution scalability, ease of managing data protection schemes, and effectiveness of NetApp deduplication.

Pod Sizing

VDI pod is analogous to XenDesktop site. A hosted shared pod is analogous to a XenApp farm. Larger pods are easier to manage.

Users per pool = (single server density) X (resource pool size)

Pools per pod = Target users per pod/users per pool

11.3 VM Resource Recommendation

Refer to [XenDesktop Planning Guide: Hosted VM-Based Resource Allocation \(Doc ID CTX127277\)](#) for general guidelines. The best practice is to use Liquidware Labs or any data collector software to find the real requirement.

Table 8) VM resource recommendations.

User Types	vCPU	RAM	Estimated Users per Core
Normal	1	2GB	5–7
Light	1	1.5GB	8–12
Heavy	2	4GB	2–4

For hosted shared desktops, various users are simultaneously hosted on a XenApp server. As with the hosted VDI model, the number of users that can be hosted on a single server depends upon the user load, CPU, and memory. However, with hosted shared desktops, CPU overcommit is not recommended for XenApp servers. Table 9 outlines the recommended values for CPU and memory based on the user load for hosted shared desktops.

If your servers are hyperthreaded, double the vCPU count, but don't double your density. The hyper-threaded counts shouldn't be used in calculations or estimates, because they can only add up to ~15% additional capacity, depending on the workload.

Table 9) Recommendations for hosted shared desktops.

User Types	vCPU	RAM	Estimated Users per VM	Estimated Users per Core
Light	2 or 4 HT	16GB	60	30
Normal	2 or 4 HT	16GB	30	15
Heavy	2 or 4 HT	16GB	14	7

12 Storage Sizing Guide

Storage sizing has three steps:

1. Gathering solution requirements.
2. Estimating storage capacity and performance.
3. Obtaining recommendations for storage configuration.

12.1 Solution Assessment

Assessment is an important first step. [Liquidware Labs](#) Stratusphere FIT and [Lakeside VDI Assessment](#) are recommended to collect network, server, and storage requirements. NetApp has contracted with Liquidware Labs to provide free licenses to NetApp employees and channel partners. For information on how to obtain software and licenses, refer to this [FAQ](#). Liquidware Labs also provides a storage template that fits NetApp system performance modeler. For guidelines on how to use Stratusphere FIT and the NetApp custom report template, refer to [TR-3902: Guidelines for Virtual Desktop Storage Profiling](#).

Virtual desktop sizing varies depends on:

- Number of the seats
- VM workload (applications, VM size, and VM OS)
- Connection broker (VMware View® or Citrix XenDesktop)
- Hypervisor type (vSphere, XenServer, or Hyper-V)
- Provisioning method (NetApp clone, linked clone, PVS, or MCS)
- Storage future growth
- Disaster recovery requirements
- User home directories

There are many factors that affect storage sizing. NetApp has developed a sizing tool system performance modeler (SPM) to simplify the process of performance sizing for NetApp systems. It has a step-by-step wizard to support varied workload requirements and provide recommendations to meet the customers' performance needs.

Best Practice

NetApp recommends using the NetApp SPM tool to size the virtual desktop solution. Contact NetApp partners and NetApp sales engineers who have the access to SPM. When using the NetApp SPM to size a solution, it is recommended to separately size the VDI workload (including write cache and personal vDisk, if used) and the CIFS profile/home directory workload. When sizing CIFS, NetApp recommends sizing with CIFS heavy user workload. 80% concurrency was assumed. It was also assumed that 10GB per user for home directory space with 35% deduplication space savings were used. Each VM used 2GB of RAM. PVS write cache is sized at 5GB per desktop for nonpersistent/pooled and 2GB for persistent desktops with personal vDisk.

Storage sizing has two factors: capacity and performance.

12.2 Capacity Considerations

Deploying XenDesktop with PVS has the following capacity considerations:

- **vDisk.** The size of the vDisk depends greatly on the operating system and the number of applications to be installed on the vDisk. It is a best practice to create vDisks larger than necessary in order to leave room for any additional application installations or patches. Each organization should determine the space requirements for its vDisk images.
20GB vDisk with a Windows 7 image is used as an example. NetApp deduplication can be used for space saving.
- **Write cache file.** NetApp recommends the size range for each user to be 4–18GB. Write cache size is based on what type of workload and how often the VM is rebooted. 4GB is used in this example for the write-back cache. Because NFS is thin provisioned by default, only the space currently used by the virtual machine will be consumed on the NetApp storage. If iSCSI or FCP is used, N x 4GB would be consumed as soon as a new virtual machine is created.
- **PvDisk.** Normally 5–10GB depending on the application and the size of the profile. Use 20% of the master image as the starting point. It is recommended to run NetApp deduplication.
- **CIFS home directory.** Various factors must be considered for each home directory deployment. The key considerations for architecting and sizing a CIFS home directory solution include the number of users, the number of concurrent users, the space requirement for each user, and the network load. Run deduplication and obtain space saving.
- **vSwap.** VMware ESXi and Hyper-V each require 1GB per VM.
- **Infrastructure.** Host XenDesktop, PVS, SQL Server, DNS, and DHCP.

The space calculation formula for a 2,000-seat deployment:

Number of vDisk x 20GB + 2,000 x 4GB write cache + 2,000 x 10GB PvDisk + 2,000 x 5GB user home directory x 70% + 2,000 x 1GB vSwap + 500GB infrastructure

12.3 Performance Considerations

Performance requirement collection is a critical step. After using [Liquidware Labs Stratusphere FIT](#) and [Lakeside VDI Assessment](#) to gather I/O requirements, contact NetApp's account team to obtain recommended software and hardware configuration.

I/O has a few factors: size, read/write ratio, and random/sequential. We use 90% write and 10% read for PVS workload. Storage CPU utilization also needs to be considered. Table 10 can be used as guidance for your sizing on PVS workload when using Login VSI heavy workload.

Table 10) Sizing guidance.

	Boot IOPS	Login IOPS	Steady IOPS
Write cache (NFS)	8–10	9	7.5
vDisk (CIFS SMB 2.1)	0.5	0	0
Infrastructure (NFS)	2	1.5	0

12.4 Desktop Deployment Scenarios

To provide storage scalability guidelines, Citrix and NetApp have categorized the Citrix PVS-based deployments into six scenarios based on the use cases that we see deployed in customer environments. The first four deployment scenarios are more common, and the last two are relatively uncommon.

Table 11) Citrix PVS-based deployment scenarios.

Desktop Deployment Scenario	PVS-Based Desktop Deployment Models	Use Cases	User State (with Regard to OS)	Profile	Folders	User-Installed Apps
Common Deployment Scenarios						
1	Pooled streamed desktop	Task workers: training labs, call centers, admin workstations, and so on. Typically deployed when hosted shared desktops cannot be used because of application compatibility issues.	Nonpersistent	Roaming*	Redirected	Nonpersistent
2	Hosted shared desktops (streamed XenApp server)	Task workers: kiosks, training labs, call centers, and so on. No application compatibility issues with multiuser sessions on same Windows machine. Users that have minimal customization needs beyond roaming profile or folder redirection.	Nonpersistent	Roaming*	Redirected	Nonpersistent
3	Personal streamed desktop (roaming between desktops)	Knowledge workers: developers, contractors, and so on. Users who tend to need a lot of supporting apps that IT does not want to manage.	Persistent	Roaming*	Redirected	Persistent

Desktop Deployment Scenario	PVS-Based Desktop Deployment Models	Use Cases	User State (with Regard to OS)	Profile	Folders	User-Installed Apps
4	Personal streamed desktop (no roaming between desktops)	Variation of scenario 3 in which the user does not roam between desktops.	Persistent	Local	Local	Persistent
Less Common Deployment Scenarios						
5	Hosted shared desktop (streamed XenApp servers)	Task workers with very controlled and limited desktop environment.	Nonpersistent	Mandatory*	Redirected	Nonpersistent
6	Pooled streamed desktop	Variation of scenario 4 in which users sometimes log in from other desktops and require access to their data, but don't care about application settings.	Nonpersistent	Local	Redirected	Nonpersistent

* Using profile management software, for example, Citrix UPM.

12.5 IOPS and Capacity Considerations

Table 12 shows the IOPS and capacity assumptions (per desktop) used to generate the sample storage array scalability guidelines later in this section.

The IOPS requirements were determined based on the performance testing using Login VSI heavy user workload in conjunction with profile management tools such as Citrix UPM and Liquidware Labs ProfileUnity redirecting user profile and folders to CIFS shares.

The capacity requirements are assumptions, based on what is typically observed in customer deployments and during lab testing.

Note that these IOPS and capacity numbers might vary for your deployment based on your work profile and applications in your environment.

Table 12) IOPS and capacity assumptions.

Desktop Deployment Scenario		PVS vDisk (Shared by 100s of Users)	PVS Write Cache File	CIFS Share	Personal vDisk	IOPS Served by Shared Storage	IOPS Served by Host Cache	Capacity Requirement on Shared Storage
1	Capacity	25GB	5GB	5GB	N/A	-	-	10GB
	IOPS	5 (100%R; local)	7 (95%W)	7 (2:W, 1:R, 4:other)	N/A	14	5	-
2	Capacity	25GB	1GB	5GB	N/A	-	-	6GB
	IOPS	5 (100%R; local)	4 (66%R)	7 (2:W, 1:R, 4:other)	N/A	11	5	-
3	Capacity	25GB	2GB	5GB	10GB	-	-	17GB
	IOPS	5 (100%R; local)	2 (95% W)	7 (2:W, 1:R, 4:other)	5 (95% W)	14	5	-
4	Capacity	25GB	2GB	N/A	10GB	-	-	12GB

Desktop Deployment Scenario		PVS vDisk (Shared by 100s of Users)	PVS Write Cache File	CIFS Share	Personal vDisk	IOPS Served by Shared Storage	IOPS Served by Host Cache	Capacity Requirement on Shared Storage
	IOPS	5 (100%R; local)	2 (95% W)	N/A	6 (90% W)	8	5	-
5	Capacity	-	1GB	5GB	N/A	-	-	6GB
	IOPS	5 (100%R; local)	4 (66%R)	7 (2:W, 1:R, 4:other)	N/A	11	5	-
6	Capacity	-	5GB	5GB	N/A	-	-	10GB
	IOPS	5 (100%R; local)	7 (95%W)	7 (2:W, 1:R, 4:other)	N/A	14	5	-

12.6 Storage Scalability Guidelines

Based on the IOPS and capacity requirements per desktop listed in Table 12, Table 13 provides the NetApp unified storage scalability guidelines for each desktop deployment scenario on FAS2240HA and FS3250HA systems. Similar scalability guidelines can be obtained for other FAS models by using the NetApp SPM sizing tool.

Table 13) Storage scalability guidelines.

Scenario	Sample Configurations			
	FAS2240HA Configuration	Max. # of Desktops	FAS3250HA with Flash Cache	Max. # of Desktops
1	FAS2240HA w/ 54 disks (10RU)	1,400	FAS3250HA w/ 84 disks (22RU)	3,000
2	FAS2240HA w/ 112 disks (18RU)	3,200	FAS3250HA w/ 67 disks (18RU)	4,500
3	FAS2240HA w/ 60 disks (10RU)	1,300	FAS3250HA w/ 109 disks (26RU)	2,900
4	FAS2240HA w/ 50 disks (10RU)	1,600	FAS3250HA w/ 100 disks (26RU) –No Flash Cache	3,500
5	FAS2240HA w/ 112 disks (18RU)	3,200	FAS3250HA w/ 67 disks (18RU)	4,500
6	FAS2240HA w/ 54 disks (10RU)	1,400	FAS3250HA w/ 84 disks (22RU)	3,000

The following key assumptions were made while sizing these storage configurations:

1. All VMs in each configuration boot under 10 minutes.
2. Login VSI application latency is within the acceptable limits; that is, the VSI_{max} number less than 1,500.
3. NetApp storage controller utilization no more than 60% during steady state operations.
4. No performance degradation occurs when one NetApp controller fails.
5. Personal vDisk dedupe savings: 35% assumption (savings might be higher or lower for your environment depending on the type and amount of data in the personal vDisk).
6. 2GB storage per VM as VM memory vSwap space on storage.
7. CIFS sizing:
 - a. Workload type: heavy
 - b. User concurrency: 80%
 - c. Dedupe savings: 35% assumption (savings might be higher or lower for your environment depending on the type and quantity of data in the CIFS share)

8. NetApp Flash Cache: Flash Cache provided significant disk savings on the FAS3250 platform for all the desktop deployment scenarios that require CIFS share (except scenario #4) and for the hosted shared desktops use case, in which 66% read traffic to the PVS write cache file was observed.
9. Disk type used in this sizing is 600GB SAS (15K RPM).

Table 14 shows the scalability guidelines on a highly dense two-rack unit FAS2240-2 system with 24 internal drives (600GB; 10K RPM SAS). The sizing assumptions discussed previously were used to generate these guidelines.

Figure 21) FAS2240-2 system (2U) with 24 internal SAS drives (600GB; 15K RPM).



Table 14) Scalability guidelines.

Desktop Deployment Scenario	Maximum # of Desktops
1	600
2	550
3	500
4	800
5	550
6	600

Note: The same two-rack unit FAS2240-2 system with 24 internal drives (600GB, 10K RPM SAS) can host more than 1,000 desktops for very basic VDI deployment:

- Citrix PVS with 5GB write cache file and 2GB RAM per VM
- 10–12 IOPS per desktop to PVS write cache file (95% writes)
- No profile management or folder redirection

Acknowledgement

Troy Mangum, Senior Manager, Platform and Integrations Engineering

Abhinav Joshi, Senior Product Manager, Desktop Virtualization

Brad Flanary, Manager, Systems and Architecture

References

Citrix References

- [How to Optimize XenDesktop Machines](#)
- [XenDesktop Modular Reference Architecture](#)
- [XenDesktop Design Handbook](#)

Other articles of note:

- [Hosted VM Based Resource Allocation](#)
- [Memory and Storage for Provisioning Server](#)

- [XenDesktop Scalability](#)
- [XenDesktop DB Sizing Best Practices](#)

NetApp References

- [TR-3795: Deployment Guide for XenDesktop 3.0 and VMware ESX Server on NetApp](#)
- [TR-3505: NetApp Deduplication for FAS, Deployment and Implementation Guide](#)
- [TR-3705: NetApp and VMware View Solution Guide](#)
- [TR-3732: Citrix XenServer and NetApp Storage Best Practices](#)
- [TR-4068: VMware vSphere 5 on NetApp Data ONTAP 8.1 Operating in Cluster-Mode](#)

VMware References

- [Sysprep file locations and versions](#)
- [vSphere Virtual Machine Administration Guide](#)

Refer to the [Interoperability Matrix Tool](#) (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

[Go further, faster®](#)

© 2013 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, AutoSupport, Data ONTAP, Data ONTAP-v, Flash Cache, FlexClone, OnCommand, SnapMirror, and Snapshot are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Apple and Mac are registered trademarks of Apple, Inc. ESX, View, VMware, and VMware vSphere are registered trademarks and ESXi and vCenter are trademarks of VMware, Inc. Active Directory, Excel, Microsoft, Outlook, SQL Server, Windows, and Windows Server are registered trademarks and Hyper-V is a trademark of Microsoft Corporation. Cisco and Cisco UCS are registered trademarks of Cisco Systems, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4138-0313