



Technical Report

VMware Horizon View 5 Solutions Guide

Chris Gebhardt, NetApp
November 2013 | TR-4181

Abstract

This solution guide outlines VMware® Horizon View™ 5 on NetApp storage. The guide covers pool design, cloning choices, and guest optimization, as well as end-user data layout. New technologies that are available in View 5.1, such as View Storage Accelerator, are also covered. This solution guide explains the different design choices and why each choice affects the outcome of the virtual desktop deployment.

TABLE OF CONTENTS

1	Executive Summary.....	6
1.1	Implementing Best Practices.....	6
1.2	Audience.....	6
2	Scope	6
3	VMware View 5	7
3.1	Overview	7
3.2	VMware View with Clustered Data ONTAP	9
3.3	Storage Efficiency for VMware View	10
3.4	VMware View Performance Considerations.....	11
4	VMware View 5 Microsoft Windows 7 Guest OS	12
4.1	Overview	12
4.2	How to Optimize the Guest OS	12
4.3	High-Level Guest Optimization Workflow	13
4.4	Setting Guest OS Timeout Value	14
5	VMware View 5 Desktop Pools	14
5.1	Overview	14
5.2	Automated Desktop Pools.....	15
5.3	Manual Desktop Pools	17
5.4	User Assignment.....	18
5.5	Desktop Pool Comparison	18
6	VMware View 5 Composer	19
6.1	Overview	19
7	VMware View 5 Full Clones	20
7.1	Overview	20
7.2	VMware Full-Clone Architecture	20
7.3	Virtual Desktop Machine Architecture	21
7.4	Datastore Design for Full Clones	21
7.5	Number of Virtual Desktops per Datastore.....	22
7.6	Performance Considerations.....	24
7.7	Known Limitations.....	25
7.8	Management Considerations	25
8	VMware View 5 Native Linked Clones	25

8.1	Overview	25
8.2	VMware Native Linked-Clone Architecture.....	26
8.3	Native Linked-Clone Architecture.....	26
8.4	Additional Native Linked-Clone Disks	27
8.5	Common Maintenance Activities.....	28
8.6	Datastore Design for Native Linked Clones	28
8.7	Performance Considerations.....	29
8.8	Recommendations to Reduce Additional I/Os.....	29
8.9	Sizing Considerations	29
9	VMware View 5.1 View Storage Accelerator	30
9.1	Overview	30
9.2	Storage Considerations	33
9.3	Additional Documentation	33
10	VMware View 5 VSC Clones	33
10.1	Overview	33
10.2	VMware VSC-Clone Architecture.....	34
10.3	Virtual Desktop Machine Architecture	34
10.4	Datastore Design for VSC Clones.....	35
10.5	Number of Virtual Desktops per Datastore.....	36
10.6	Performance Considerations.....	38
10.7	Known Limitations	39
10.8	Management Considerations	39
11	VDI User Data Layout	39
11.1	Overview	39
11.2	What Is User Data?.....	40
11.3	Different Types of User Data.....	40
11.4	Separate User Data from the Virtual Machine.....	41
11.5	Profile Management Considerations	42
12	System Performance Modeler for VMware View	42
13	Feedback	43
	References.....	43
	NetApp Documents	43
	VMware Documents	44

Version History	44
About the Author	44
Acknowledgements	45

LIST OF TABLES

Table 1) Deduplication recommendations for different data types.....	11
Table 2) Automated pool maintenance activities.....	16
Table 3) User assignment by clone type and additional technologies.....	16
Table 4) Desktop pool comparison.....	18
Table 5) Datastore settings for deduplication, backup and recovery, and replication.....	22
Table 6) Recommended number of VMs per datastore.....	23
Table 7) Storage consumption based on initial deployment after deduplication.....	24
Table 8) Best practices for deduplication, backup and recovery, and replication.....	29
Table 9) Datastore settings for deduplication, backup and recovery, and replication.....	36
Table 10) Recommended number of VMs per datastore.....	37
Table 11) Storage consumption based on initial deployment after deduplication.....	38

LIST OF FIGURES

Figure 1) VMware View architecture.....	7
Figure 2) Scaling a VMware View infrastructure.....	9
Figure 3) Moving volumes nondisruptively between storage nodes.....	10
Figure 4) Linear scalability with clustered Data ONTAP.....	10
Figure 5) Command files included in the VMware Horizon View Optimization Guide for Windows 7.....	13
Figure 6) VMware View pool design.....	15
Figure 7) VMware View linked clone using View Composer.....	19
Figure 8) Full-clone solution architecture.....	20
Figure 9) VMware full-clone architecture.....	21
Figure 10) Unique copy equation and example for Data ONTAP 8.0.x and earlier.....	23
Figure 11) Unique copy equation and example for Data ONTAP 8.1 and later.....	23
Figure 12) Native linked-clone solution architecture.....	26
Figure 13) Native linked-clone architecture.....	27
Figure 14) Selecting the hypervisor, broker, and cloning method dialog box.....	30
Figure 15) Host cache settings.....	31
Figure 16) Enabling View Storage Accelerator.....	32
Figure 17) Setting blackout days.....	33
Figure 18) VSC-clone solution architecture.....	34

Figure 19) VMware VSC-clone architecture.	35
Figure 20) Unique copy equation and example for Data ONTAP 8.0.x and earlier.	37
Figure 21) Unique copy equation and example for Data ONTAP 8.1 and later.	37
Figure 22) Different data types of a virtual desktop.	39
Figure 23) Data type organized by volume.	41

1 Executive Summary

NetApp solutions enable companies to optimize their virtual desktop infrastructures by providing advanced storage and data management capabilities. NetApp offers industry-leading storage solutions that simplify virtual desktop provisioning; enable mass VM cloning and redeployment; handle typical input/output bursts such as boot storm, antivirus storms, efficient operating system, applications, user data management, and so on; provide individual VM backup and restores; deliver simple and flexible business continuance; and help reduce virtual desktop storage.

This solution guide offers guidelines and best practices for architecting, deploying, and managing VMware Horizon View virtual desktop infrastructure (VDI) solutions on NetApp® storage systems. NetApp has been providing advanced storage features to VMware ESX®-based solutions since the product began shipping in 2001. During that time, NetApp has continuously enhanced the design, deployment, and operational guidelines for the storage systems and ESX server-based VDI solutions. These techniques have been documented and are referred to as best practices. This guide describes them in detail.

1.1 Implementing Best Practices

The recommendations and practices presented in this document should be considered deployment requirements unless otherwise stated. Although choosing not to implement all of the best practices contained in this guide does not affect your ability to obtain support from NetApp and VMware, disregarding any of these practices commonly results in the need to implement them at a later date, on a much larger environment, and often with the requirement of application downtime. For this reason, NetApp recommends that you implement all of the best practices as defined in this document as a part of initial deployment or migration.

All recommendations in this document apply specifically to deploying vSphere® on NetApp. Therefore, this document supersedes all recommendations and best practices expressed in other versions of this document and all others.

NetApp Data ONTAP® 8.1 or later is required to implement the NetApp vSphere plug-in. However, many features discussed in this paper may be available only in versions of Data ONTAP later than 8.1. NetApp recommends using clustered Data ONTAP 8.2 or later for new Horizon View deployments whenever possible.

NetApp and our partners offer professional services to architect and deploy the designs described in this document. These services can help you to enable optimal virtual storage architecture for your virtual data center.

This document refers to current software versions from NetApp, VMware, and other software vendors. The versions listed in this document are supported, but previous versions may no longer be supported. For official supported versions, consult a NetApp systems engineer.

1.2 Audience

The target audience for this paper is familiar with concepts pertaining to VMware vSphere, including VMware ESX, VMware vCenter Server™, and NetApp Data ONTAP 8.0 or later. For high-level information and an overview of the benefits that are available when creating a virtual infrastructure on NetApp storage, see [Comprehensive Virtual Desktop Deployment with VMware and NetApp](#).

2 Scope

This document offers architectural, deployment, and management guidelines for customers who are thinking about or have already decided to implement VMware View on NetApp virtualized storage. It provides an overview of the VMware View technology concepts; key solution architecture considerations

for implementing VMware Horizon View on NetApp; storage estimation and data layout recommendations; and solution, deployment, and management guidelines. It does not include features available in later releases of VMware Horizon™ 5.2, such as View Composer API Integration (VCAI), the SE sparse disk format, and best practices for View Storage Accelerator. Future updates to this document will cover these areas.

3 VMware View 5

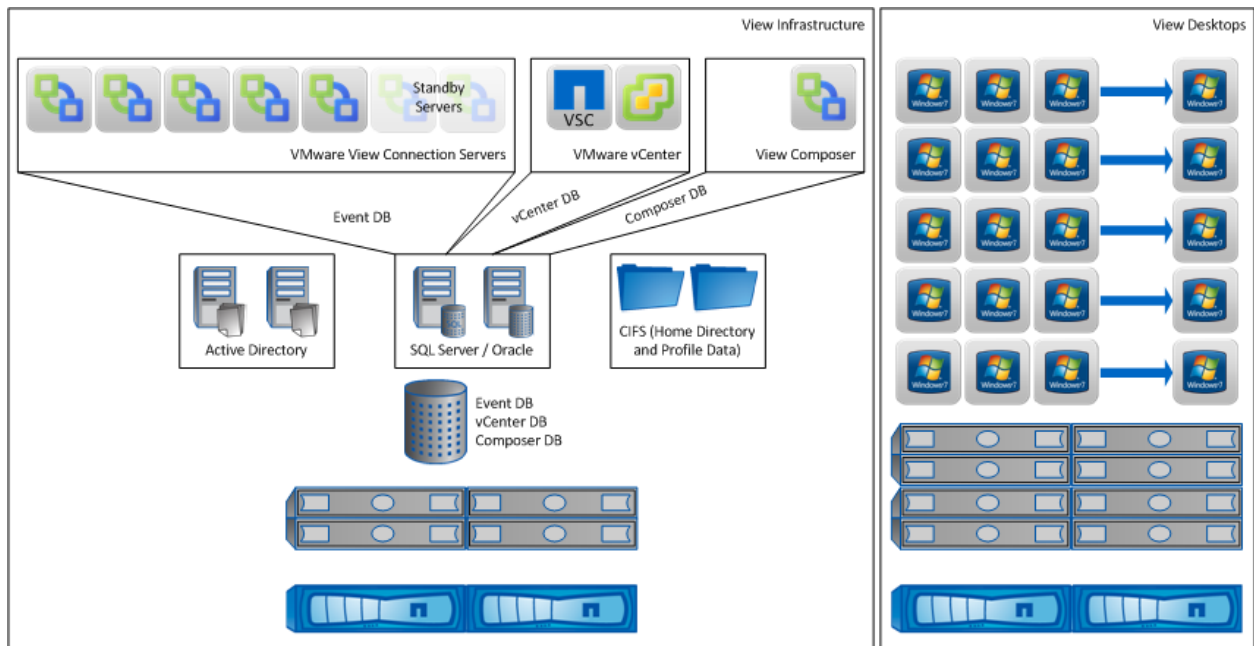
3.1 Overview

VMware View is an enterprise-class solution that delivers corporate PC services to end users. It enables organizations to increase corporate IT control and the manageability and flexibility of desktop resources while providing end users with a familiar desktop experience.

VMware View enables IT to deliver personalized and efficient Windows® virtualized desktops to end users. The combination of View with NetApp storage improves the end-user experience and provides IT with a cost-effective solution. NetApp enhances the end-user experience and improves virtual desktop performance with Virtual Storage Tier (VST), FlexClone® technology, write optimizations, and deduplication. By using NetApp storage efficiency technology, such as deduplication and compression, IT departments can significantly reduce capital costs.

A well-architected VMware View environment is made up of two distinct infrastructures: the View infrastructure (shown on the left in Figure 1) and the virtual desktop infrastructure (shown on the right in Figure 1). These infrastructures are divided into two discrete environments for fault isolation and performance assurance. If these environments were placed on the same infrastructure, one or more of the virtual desktop machines could potentially take more than its fair share of the resources and tax the core infrastructure, resulting in performance issues for all desktops.

Figure 1) VMware View architecture.



Best Practice

The VMware View infrastructure virtual machines (VMs) and the virtual desktops should operate on separate hardware. This separation provides the highest level of fault isolation and performance assurance.

VMware View Infrastructure

The VMware View infrastructure includes View Connection Server, View Composer, and NetApp storage controllers.

VMware View Connection Server is responsible for provisioning and managing virtual desktops and for brokering the connections between clients and the virtual desktop machines. A single Connection Server instance can support up to 2,000 simultaneous connections. In addition, five Connection Server instances can work together to support up to 10,000 virtual desktops. For increased availability, View supports using two additional Connection Server instances as standby servers. The Connection Server can optionally log events to a centralized database that is running either Oracle® Database or Microsoft® SQL Server®.

Best Practice

- Use the full version of SQL Server or Oracle for the View Composer database.
- Use the NetApp SnapManager® solution for database backup and recovery.
- Use NetApp Virtual Storage Console (VSC) to set the recommended timeout values and to provision datastores when deploying a vSphere infrastructure on NetApp storage.

Note: For detailed information about database selection, refer to <http://pubs.vmware.com/view-52/index.jsp>.

Best Practice

NetApp recommends architecting at least one spare Connection Server instance in a View solution to provide service in the event of a host failure. In addition, use anti-affinity rules in vSphere to keep the Connection Server instances on different physical hosts.

The VMware View Composer server can be a standalone VM, or it can be co-installed on any server, except for a VMware View connection, security, or transfer server. The View Composer server is responsible for creating View linked clones and for performing maintenance operations on these clones. The View Composer server works in conjunction with the VMware View Connection Server to rapidly provision storage-efficient virtual desktops for use in the View desktop environment.

The final component of the VMware View infrastructure is the location in which the end-user data (home directory or profile data) resides. This data is stored and served by the NetApp storage controller through the CIFS protocol.

Best Practice

As of VMware View 5.1, VMware View Composer can be installed as a standalone server; it supports both Windows based vCenter Servers and the vCenter Server appliance.

Virtual Desktop Infrastructure

The virtual desktop infrastructure is where the virtual desktop machines reside. The architecture of this infrastructure comprises virtual desktop machines running Windows desktop operating systems (OSs) and physical servers running VMware ESXi™.

Note: If Windows XP or Windows Vista® is used as the desktop OS, refer to [TR-3747: Best Practices for File System Alignment in Virtual Environment](#).

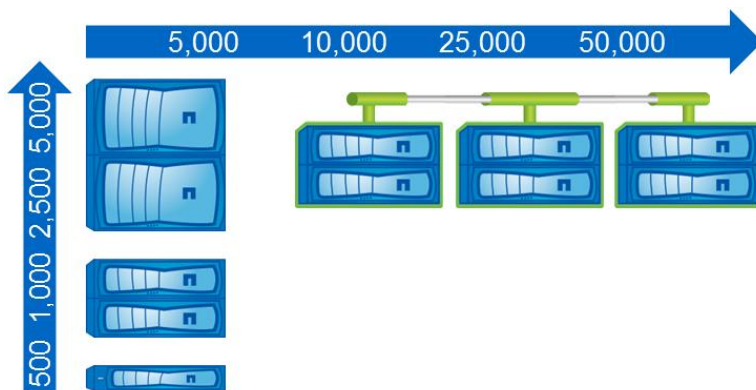
VMware View Pool of Desktops

In a typical large-scale virtual desktop deployment, the maximum limits of the VMware View Connection Server can be reached when each Connection Server instance supports up to 2,000 simultaneous connections. When this occurs, it is necessary to add more Connection Server instances and to build additional VMware View desktop infrastructures to support additional virtual desktops. Each such desktop infrastructure is referred to as a *pool of desktops (POD)*.

A POD is a building-block approach to architecting a solution. The size of the POD is defined by the VMware View desktop infrastructure (the desktop VMs) plus any additional VMware View infrastructure resources that are necessary to support the desktop infrastructure PODs. In some cases, it might be best to design PODs that are smaller than the maximum size to allow for growth in each POD or to reduce the size of the fault domain.

Using a POD-based design gives IT a simplified management model and a standardized way to scale linearly and predictably. Figure 2 shows vertical scaling with different NetApp FAS controllers, and clustered Data ONTAP allows horizontal scalability. By using clustered Data ONTAP, customers can have smaller fault domains that result in higher availability.

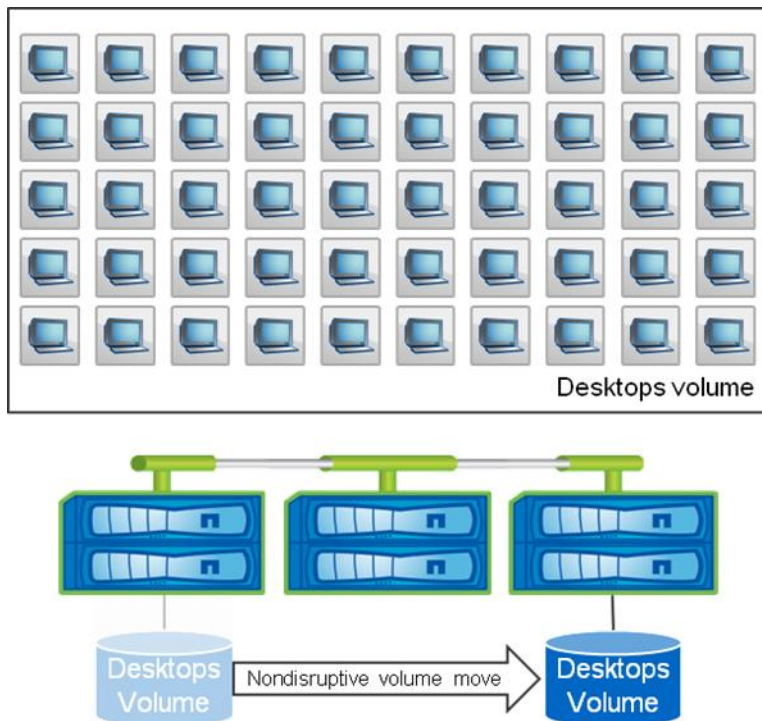
Figure 2) Scaling a VMware View infrastructure.



3.2 VMware View with Clustered Data ONTAP


VMware View environments benefit greatly from using clustered Data ONTAP to increase business agility, reduce scheduled downtime, and simplify storage management. For example, by using the `vol move` operation, IT can seamlessly and nondisruptively move entire pools of users or entire PODs from one node to another within the storage cluster. This operation can be used for maintenance activities, hardware upgrades, or even load balancing. Figure 3 illustrates the movement of a volume that contains virtual desktops from one node in the storage controller to another.

Figure 3) Moving volumes nondisruptively between storage nodes.



Clustered Data ONTAP also enables linear scalability. It allows IT to seamlessly add, upgrade, or remove storage infrastructure to meet the needs of the virtual desktops, as illustrated in Figure 4. Clustered Data ONTAP is the perfect infrastructure solution for leveraging POD-based designs.

Figure 4) Linear scalability with clustered Data ONTAP.



3.3 Storage Efficiency for VMware View

One of the critical barriers to virtual desktop infrastructure adoption is the increased cost of using shared storage to obtain a highly available enterprise-quality infrastructure. A virtual desktop deployment creates a high level of data redundancy, especially for the VM OS data. When traditional storage is used, this means that the amount of storage needed is equal to the sum of the storage required by each VM. For

example, if each VM were 20GB in size, and 1,000 VMs were planned in the solution, then at least 20TB of usable data would be required on the shared storage.

Thin provisioning, data deduplication, compression, space reclamation, and FlexClone technology are the critical components of a NetApp solution that offers multiple levels of storage efficiency across the virtual desktop OS data, installed applications, and user data. This efficiency helps customers save 50% to 90% of the cost associated with shared storage (estimates based on existing customer deployments and lab validation of NetApp solutions).

In a VMware View environment, deduplication provides significant space savings, given that each VM is an identical copy of the OS, applications, and patches. Deduplication is also ideal for user and personal (profile) data stored in CIFS home directories.

Note: Not all data in a virtual desktop infrastructure environment is ideal for deduplication. Data such as swap files and other transient data should not be deduplicated.

Table 1 presents deduplication recommendations for different data types.

Table 1) Deduplication recommendations for different data types.

Datastore Type	Enable Deduplication?	Clone Type
Template datastore	Yes	All
Linked-clone replica datastore	Yes	Native linked clones
Linked-clone OS disk datastore Δ (delta)	No	Native linked clones
Full clones	Yes*	Full clones
NetApp Virtual Storage Console clones	Yes*	FlexClone clones
View Composer Array Integration linked clones (NFS only)	Yes*	FlexClone clones
vSwap	No	All
Disposable file redirection disk	No	Native linked clones
View Composer data disk	Yes	Native linked clones

* In Table 1, the schedule for deduplicating full clones regardless of provisioning method can vary depending on the amount of change in the VM. More frequent deduplication might be required to shorten the deduplication process. NetApp recommends monitoring the deduplication process and adjusting it to fit the replication and backup requirements.

3.4 VMware View Performance Considerations

Another critical barrier to VMware View adoption is the performance issue associated with hosting thousands of VMs on shared storage—specifically, performance associated with events that produce a large influx of simultaneous I/O, such as login storms, boot storms, and antivirus operations.

With physical desktops, these events are not a problem because each machine has its own disks, and I/O is contained in a single desktop. However, because VMware View uses a shared storage infrastructure, significant performance issues might arise during these critical operations. This essentially means that the solution would require a large number of additional spindles to meet the performance requirements, resulting in an increase in the overall solution cost.

To solve this problem, the NetApp solution includes the Virtual Storage Tier. VST is a core component of the Data ONTAP OS and is extended with Flash Cache™ intelligent caching. These components reduce the cost of the solution by offering the following benefits:

- Require far fewer disks and much less cache
- Do not require tiers of solid-state drive disks to alleviate boot and login storms
- Serve read data from cache and free up disk I/O to perform writes
- Provide better throughput and system utilization
- Provide faster response times and a better overall end-user experience

Using this NetApp solution translates into higher cache deduplication and high cache hit rates. VST is especially effective at addressing the simultaneous system boot (boot storm) and login of hundreds or thousands of virtual desktop systems that can overload a traditional legacy storage system.

With the NetApp space-efficient cloning technologies, VMware View implementations can realize as much as 99% initial space savings.

Note: For information about how the savings estimates were validated, refer to [WP-7108: 50,000-Seat VMware View Deployment](#).

4 VMware View 5 Microsoft Windows 7 Guest OS

4.1 Overview

Many features, services, and maintenance tasks in the Microsoft Windows 7 operating system were designed to be used specifically with desktops and laptops, but not with virtual desktops. The services that should be disabled on the virtual desktop template include, but are not limited to, Disk Defragmentation, Microsoft Update, System Restore, and Windows Defender. These unnecessary services provide little or no value in a virtual desktop environment, but they still consume CPU, memory, and storage resources from both a capacity and an I/O perspective.

To achieve the most scalable, highest performing, and most cost effective VMware View deployment, it is necessary to optimize the guest OS.

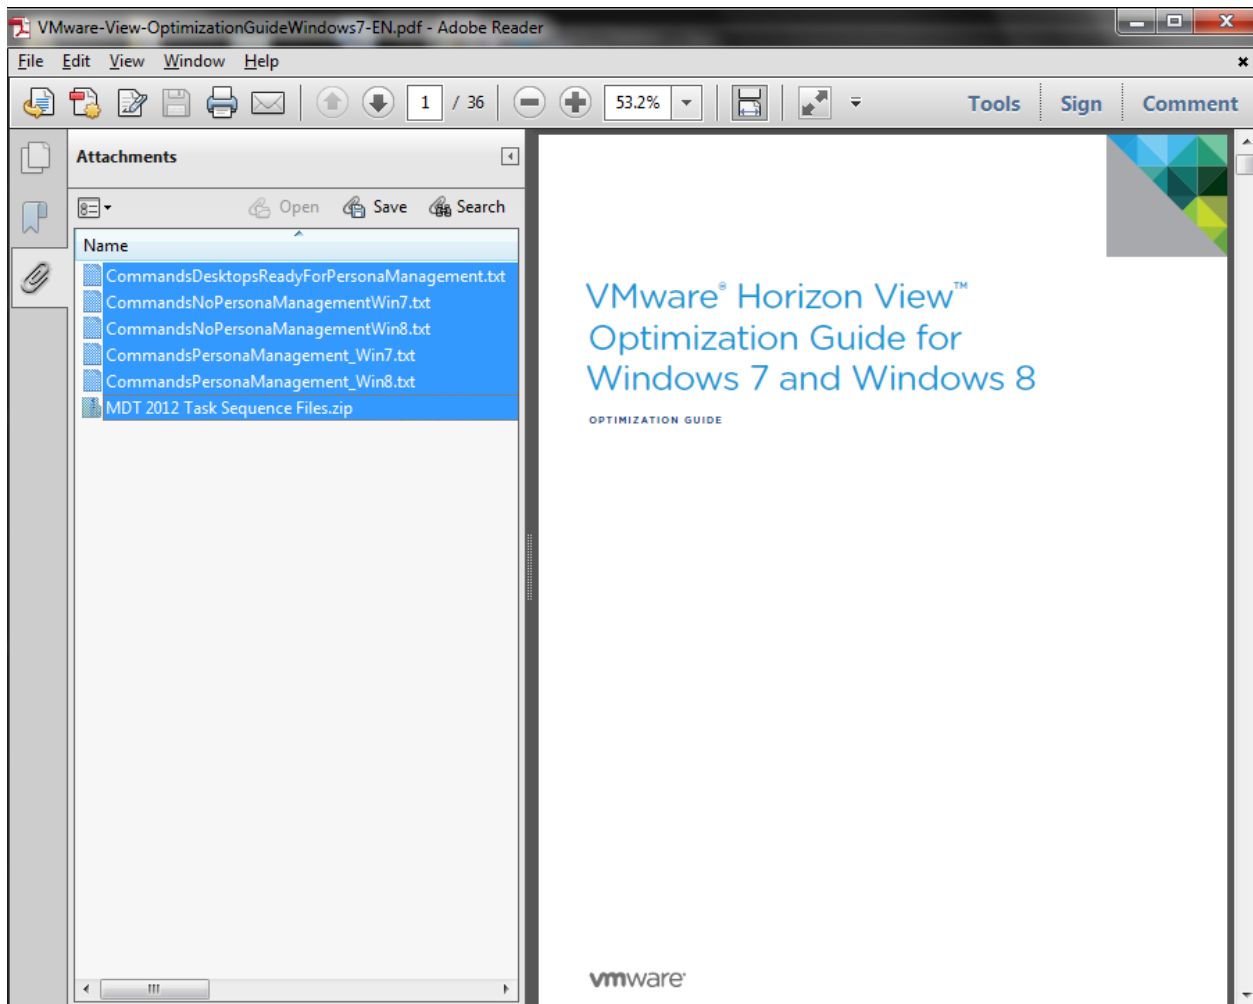
4.2 How to Optimize the Guest OS

The [VMware Horizon View Optimization Guide for Windows 7](#) describes the guest OS optimization process, from how to install Windows 7 to how to prepare the virtual machine for deployment.

The most important part of the process is to run one of the command files that are attached to the guide. These files are attached to the PDF document and can be extracted, renamed, and executed on the template VM. The command files simplify the manual process of setting registry values, running PowerShell® scripts, and changing the default settings.

Figure 5 shows the location of the command files in the PDF.

Figure 5) Command files included in the VMware Horizon View Optimization Guide for Windows 7.



4.3 High-Level Guest Optimization Workflow

Here is the high-level workflow for optimizing the guest OS for the template VM:

1. Create a template VM.
2. Install the Windows OS and any required patches.
3. Install VMware Tools.
4. Apply the guest OS optimizations with the scripts that are documented in the VMware optimization guide.
5. Install VMware View Agent.
6. Set the guest OS timeout value.
7. Power off the template VM.

Best Practices

VMware recommends applying the optimizations listed in the optimization guide to a group policy to prevent end users from modifying the values in the guest VM. Alternatively, do not give end users access privileges that allow them to modify the desktop settings.

NetApp recommends creating a new template VM when deploying virtual desktops. Using an existing laptop or desktop image can cause performance issues.

4.4 Setting Guest OS Timeout Value

It is important to set the Windows guest OS timeout value to help improve recoverability after a storage failover event. The NetApp Knowledge Base article [KB3013622](#) describes the recommended settings and the rationale for deciding which timeout value to use.

Best Practice

NetApp recommends setting the Windows guest OS timeout value to either 60 seconds (3c in hex) or 190 seconds. The timeout value can be set in one of two ways:

- By running the guest OS timeout script for Windows as described in the [Virtual Storage Console 4.1 for VMware vSphere Installation and Administration Guide](#).
- By manually editing the registry with `regedit.exe` and modifying the following key (this example sets the timeout value to 60 seconds):

```
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk]
```

```
"TimeoutValue"=dword:0000003c
```

5 VMware View 5 Desktop Pools

5.1 Overview

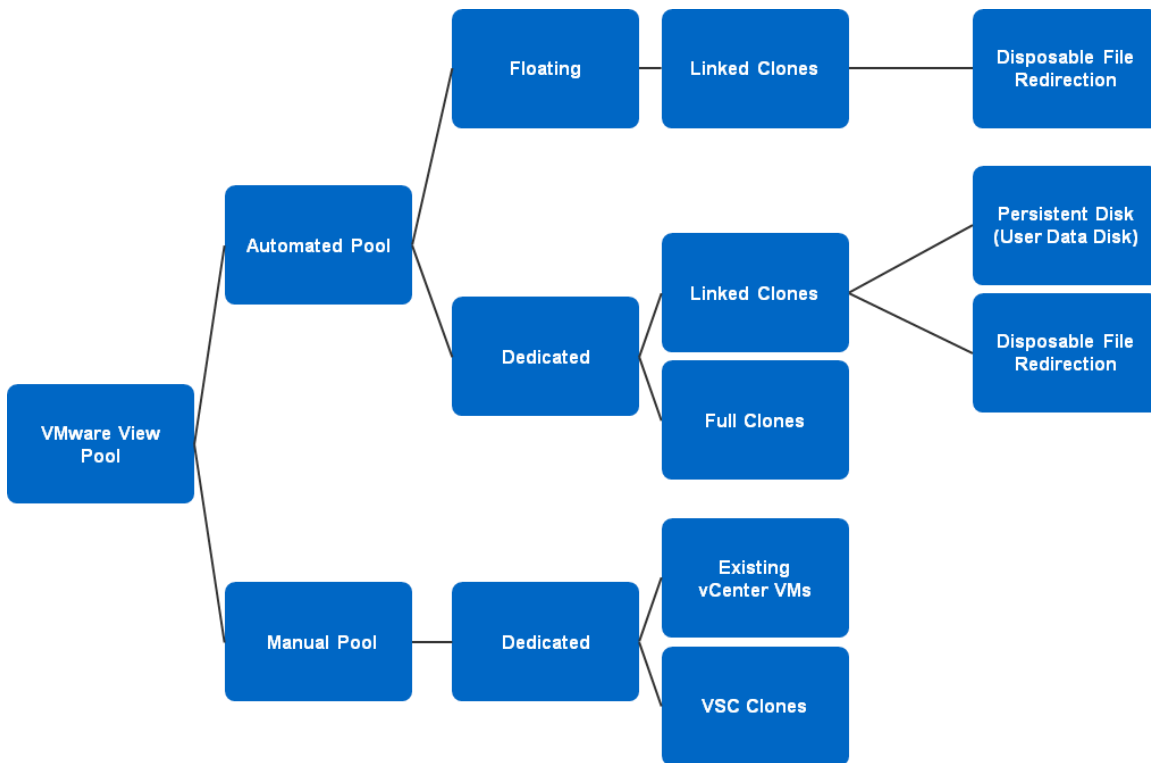
VMware View groups desktops into discrete management units called *pools*. Policies and entitlements can be set for each pool so that all desktops in a pool have the same provisioning methods, user assignment policies, logoff actions, display settings, data redirection settings, data persistence rules, and so on. There are three types of desktops pools, but this document discusses only two types, manual pools and automatic pools.

Each desktop pool type can be configured with a different user assignment type. The user assignment to a pool can be either dedicated or floating:

- **Dedicated user assignment.** The user always gets the same desktop on login.
- **Floating user assignment.** The user may or may not get the same desktop on login.

Figure 6 shows the pool types, user assignment types, and clone types that can be used in combination with each other.

Figure 6) VMware View pool design.



Note: The pool types (automated and manual) and user assignment types (floating and dedicated) shown in Figure 6 are the types most commonly deployed. Some combinations of pool type and user assignment are not included because NetApp does not recommend them.

5.2 Automated Desktop Pools

An automated desktop pool dynamically provisions virtual desktops. With this pool type, VMware View creates a portion of the desktops immediately and then, based on demand, provisions additional desktops to the limits that were set for the pool. An automated pool can contain dedicated or floating desktops. These desktops can be full clones or linked clones.

One of the major benefits of using VMware View with automated pools is that additional desktops are created dynamically on demand. This automation greatly simplifies the repetitive administrative tasks associated with provisioning desktops.

For example, suppose that an administrator sets up a pool to create 500 desktops. The administrator wants 100 desktops to be provisioned immediately and to have 25 desktops immediately available in case demand grows. If the provisioned desktops begin to be consumed and the available machines are consumed quickly, VMware View can provision additional desktops so that 25 desktops are readily available until all 500 desktops have been provisioned.

Full-Clone Desktop Pools

Full-clone desktops are created by cloning a virtual machine (VM) from a template VM by using vCenter cloning. Because each virtual desktop is an identical copy of the template (without VAAI) it uses the same amount of storage as the template when it is provisioned. As a result, full-clone desktops require a considerable amount of storage. Using NetApp deduplication technology can significantly reduce the amount of storage required by dedicated full clones in automated desktop pools. Alternatively, VAAI can be used to provision dedicated full clones in a pre-deduplicated manner.

Full-clone desktops in automated pools are useful when dedicated (persistent) desktops are required to be provisioned on demand.

Best Practice

NetApp recommends using VAAI to create dedicated full clones when provisioning new persistent desktop environments.

Note: NetApp does not recommend using automated pools with floating user assignment and full clones when the logoff action is set to Delete VM on Logoff. If all the users log off, a deletion and recreation storm for all the desktops in the pool would occur, and because deduplication is a batch process, the provisioning of the desktops might initially exceed the capacity of the storage controller until the deduplication process is able to run. If the free storage capacity is adequate to meet the demand of desktop deletion and recreation before deduplication runs, then using the Delete VM on Logoff option might be possible.

Linked-Clone Desktop Pools

To the end user, a linked-clone desktop looks and feels like a normal desktop, but it is storage efficient, consuming a fraction of the storage required for a full desktop. Because of the architecture of linked clones, three unique maintenance operations can be performed to improve the storage efficiency, the performance, and the security and compliance of the virtual desktop environment: refresh, recompose, and rebalance.

Table 2 describes the maintenance activities that are supported by each combination of user assignment and clone type in an automated desktop pool.

Table 2) Automated pool maintenance activities.

User Assignment	Clone Type	Refresh	Recompose	Rebalance	Possible Logoff Action
Floating	Linked clone	Yes	Yes	Yes	Refresh on logoff
Floating	Full clone	N/A	N/A	N/A	Delete on logoff
Dedicated	Linked clone	Yes	Yes	Yes	
Dedicated	Full clone	N/A	N/A	N/A	

Linked-clone desktops have the ability to separate the C: \ drive from transient data through disposable file redirection (DFR). With dedicated user assignments, linked clones can use a View Composer persistent disk or user data disk (UDD) to store end-user profile data. Table 3 describes the availability of these technologies for each combination of user assignment and clone type.

Table 3) User assignment by clone type and additional technologies.

User Assignment	Clone Type	DFR	Persistent Disk (User Data Disk)
Floating	Linked clone	Yes	N/A
Floating	Full clone	N/A	N/A
Dedicated	Linked clone	Yes	Yes
Dedicated	Full clone	N/A	N/A

Best Practice

NetApp recommends using DFR with linked clones in automated pools to help maintain performance and storage efficiency in the desktop environment.

5.3 Manual Desktop Pools

A manual desktop pool uses virtual desktops that exist in VMware vCenter. These desktops could have been created by using standard vCenter cloning, the NetApp Virtual Storage Console (VSC), or some other means. The most common use for manual desktops is when the solution requires persistent desktops that already exist within vCenter.

Desktops in manual pools are managed much like traditional desktops and laptops, in that patching must occur on each individual desktop. However, some situations allow a template to be patched and redeployed, depending on the cloning method that was used (vCenter or VSC).

Best Practice

NetApp recommends using dedicated user assignments with manual desktop pools. Although floating assignments can be chosen during pool creation, this architecture has limited use cases.

Full-Clone Desktop Pools

Full-clone desktops are created by copying a VM from a template or through some other method and then creating a pool in VMware View that uses these desktops. Because each virtual desktop is an identical copy of the template, it uses the same amount of storage as the template when it is provisioned. As a result, full-clone desktops require a considerable amount of storage. Using NetApp deduplication technology or leveraging VAAI to perform the cloning operation can significantly reduce the amount of storage required by full clones in manual pools.

Some points to consider when deciding whether to deploy full clones instead of VSC clones in manual pools are the amount of time and manual intervention that it takes to create each clone, the space that is required, and the amount of time it takes to run deduplication after provisioning.

Best Practice

NetApp does not recommend the use of P2V as a method to deploy users' existing laptops or desktops into VMware View. The user profile and user data should be migrated before the user migrates off a physical device. Some profile management tools that are designed to migrate user profiles can operate in both physical and virtual infrastructures.

VSC-Clone Desktop Pools

NetApp recommends the use of VSC for provisioning desktops in manual pools. VSC can create clones much faster because it uses NetApp FlexClone technology to rapidly provision the VMs. VSC clones are storage efficient because FlexClone technology does not copy data; instead, it creates the clone in a preduplicated state. By contrast, using vCenter to create the clones in a manual pool is time consuming and requires the use of NetApp deduplication after provisioning to improve storage efficiency.

Best Practice

NetApp recommends using VSC to create manual pools because VSC makes sure that best practices are followed when datastores and VMs are provisioned. VSC also reduces the time it takes to provision VMs, and it creates VMs that are storage efficient.

5.4 User Assignment

Dedicated Assignment

The dedicated assignment of desktops makes sure that users log in to the same virtual desktop each time they log in. Dedicated assignment allows users to store data either in a persistent disk (when using linked clones) or locally (when using full clones).

User-to-desktop entitlement can be a manual process or an automatic process. The administrator can entitle a given desktop to a user or can opt to allow VMware View to automatically entitle the user to a desktop when the user logs in for the first time.

Floating Assignment

With floating user assignment, users are not assigned to particular desktops and could get connected to a different desktop from the pool each time they connect.

Best Practice

NetApp recommends using a profile management solution instead of VMware View Composer persistent disks. Profile management solutions are more robust than persistent disks and make management, backup and recovery, and disaster recovery of the user data that is contained in them easier.

5.5 Desktop Pool Comparison

Table 4 compares the technologies and the benefits of using each of the pool types, user assignment types, and cloning methods. The lines in bold are recommended.

Table 4) Desktop pool comparison.

Pool Type	User Assignment	Clone Type	Storage Efficiency	Provisioning Speed	Performance
Automated	Floating	Linked clone with DFR	Yes	Fast	Low
Automated	Dedicated	Linked clone with DFR and UDD	Yes	Fast	Low
Automated	Dedicated	Full clone	With deduplication	Slow	Medium/high
Automated	Dedicated	Full Clone w/ VAAI	Yes	Fast	High
Manual	Dedicated	VSC	Yes	Fast	High
Manual	Dedicated	Existing vCenter VMs	With deduplication	N/A	Medium/high

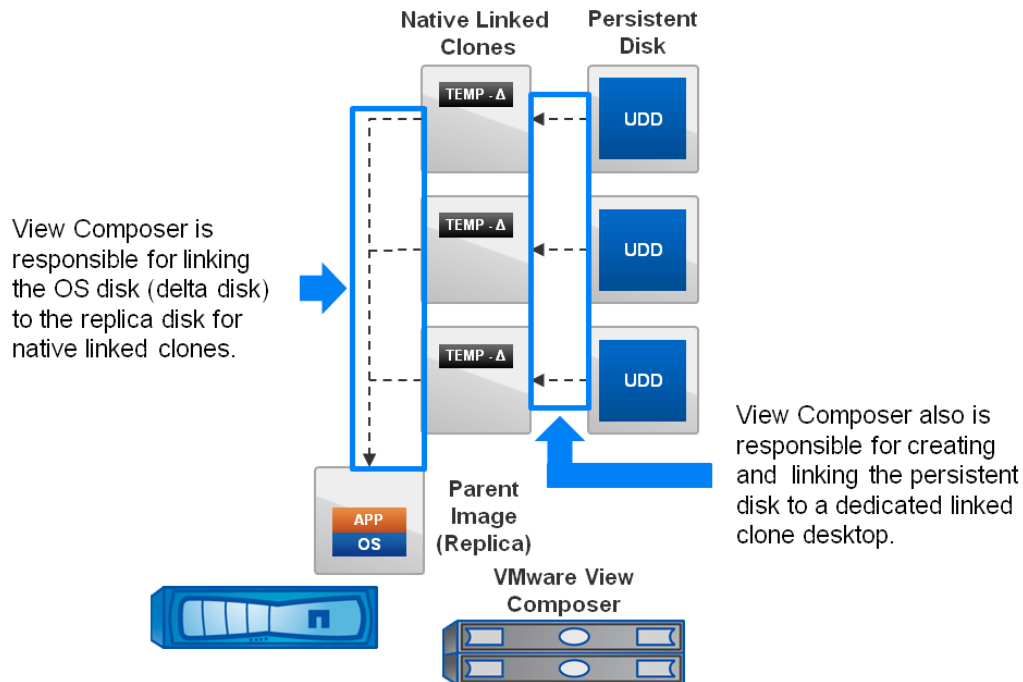
NetApp recommends a solution that requires nonpersistent desktops, automated pools used with floating user assignments, and linked clones. When dedicated, persistent desktops are a solution requirement, automated pools with dedicated user assignments and VAAI created clones offer the best overall performance and storage efficiency.

6 VMware View 5 Composer

6.1 Overview

The VMware View 5 Composer server is a critical component of solutions that use VMware View linked clones. This server is responsible for the creation and maintenance operations of VMware View linked clones. It works along with the View Connection Server to rapidly provision storage-efficient virtual desktops for use in the VMware View desktop environment. These linked-clone desktops created by the Composer can be either dedicated or floating virtual desktops in an automated pool. The Composer server is also involved during maintenance operations, such as refresh, recompose, and rebalance. These operations improve the storage efficiency, performance, security, and compliance of the virtual desktop environment. The Composer server is also involved in the creation of persistent disks (user data disks), which are used to store user preferences and profiles. Figure 7 shows a VMware View linked clone using VMware View Composer.

Figure 7) VMware View linked clone using View Composer.



The Composer server can be installed on a vCenter Server or as a standalone server, excluding any servers participating in the VMware View environment, such as the Connection Server, the transfer server, the security server, and so on.

7 VMware View 5 Full Clones

7.1 Overview

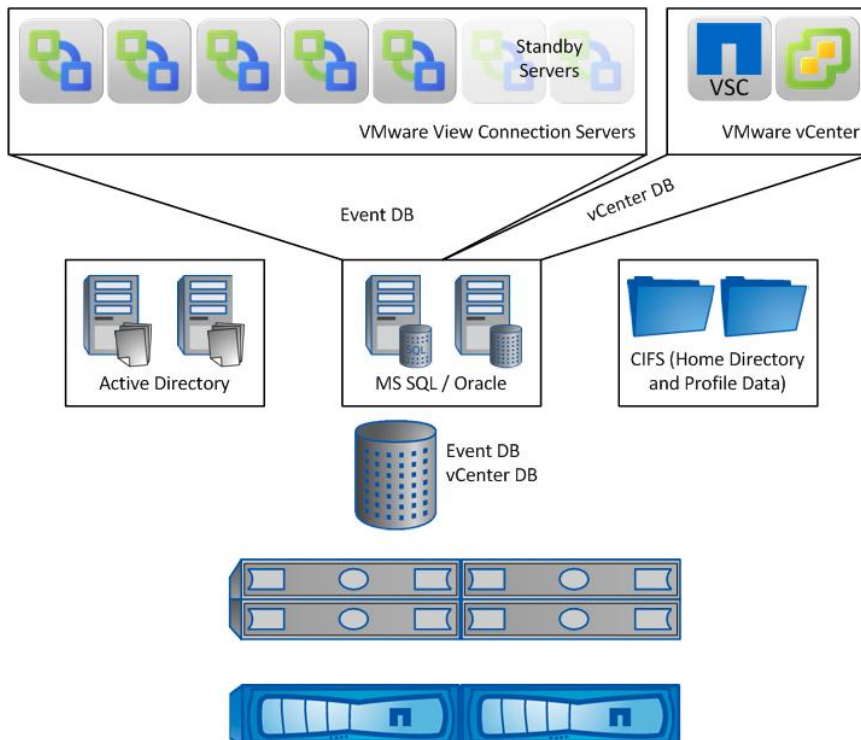
Desktops can be provisioned in many different ways for VMware View solutions that use NetApp storage. Each provisioning method has a different effect on storage consumption, performance, security, and management. Here are some of the benefits provided by full-clone desktops:

- **Storage efficiency.** Using NetApp deduplication or VAAI with full clones can significantly reduce storage requirements.
- **Performance.** Performance can be significantly improved by deduplicating datastores that contain full clones because much of the operating system is identical. The concept of deduplication and acceleration is called Virtual Storage Tier (VST). VST can help improve boot, login, and steady state workloads by caching data in a virtual tier (which is commonly the controller memory), in a caching PCIe card, or in a solid-state-drive-based controller cache.
- **Ease of management.** Two functionalities make full clones easy to manage:
 - The VMware View administrative interface is the only management interface required to create full clones.
 - Desktops can be automatically provisioned based on thresholds or demand.

7.2 VMware Full-Clone Architecture

Figure 8 shows the server components of a full-clone solution. The solution has seven VMware View Connection Server instances: Five are designated as active servers and two are designated as standby servers in the event of a failure. Full clones are created in the VMware View servers, and clients access individual virtual machines (VMs) from the VMware View servers. The VMware vCenter Server manages the VMs and provides core services, such as VMware high availability and the creation of the full-clone VMs.

Figure 8) Full-clone solution architecture.



7.3 Virtual Desktop Machine Architecture

A VMware full clone created by VMware View contains at least one virtual disk. This virtual disk is a host-based clone or VAAI clone (copy) of the VM template that contains the operating system and any corporate-installed patches and applications. The host-based clone by itself is not storage efficient, but by using NetApp deduplication, a 99% reduction in initial storage consumption can be achieved after provisioning. The VAAI clone is a storage-efficient clone. Over time, it is necessary to continue to deduplicate the datastore in which the full clones are created to maintain storage efficiency. Typical deduplication savings of 50% to 85% have been seen in real-world deployments, but the actual storage savings vary by implementation.

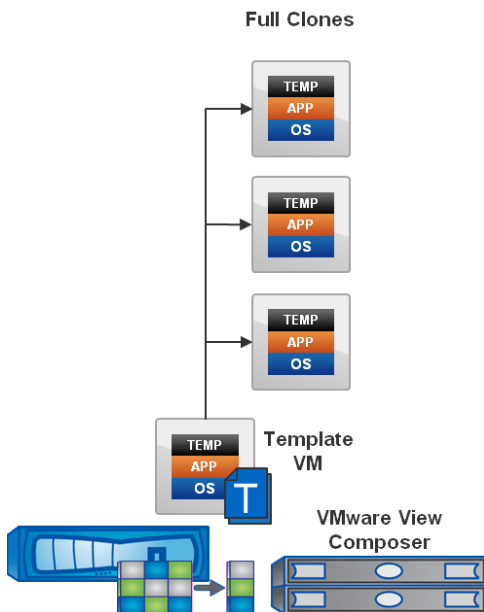
Figure 9 illustrates three full-cloned VMs, each of which is a distinct copy created from the template VM. In this case, the vCenter Server clones the entire virtual desktop from the template VM, including the OS, corporate-installed patches and applications, and any temporary data that resided in the template.

Note: It is important to remove all transient data, such as installers and temporary files, from the template VM. Removing transient data reduces the size of the template, the time it takes to clone it, and the time it takes to perform the initial deduplication of the VM datastore.

Best Practices

- Remove all transient data from the template VM before deploying virtual desktops.
- When using Network File System (NFS) with VSC, perform a space-reclamation process on the template to make the VM as small as possible.

Figure 9) VMware full-clone architecture.



7.4 Datastore Design for Full Clones

The design of datastore architectures for VMware View full-clone environments should include at least two datastores to simplify management, improve storage efficiency, and allow different backup, recovery, and disaster recovery policies for each data type.

The first datastore should contain the template VMs. In this datastore, patching and maintenance of the virtual-desktop templates can take place. Using the NetApp VSC, the administrator can create a clone of the original template, make changes to the template, and create new virtual desktops from the newly

changed template. The first datastore can then be replicated to remote sites or other environments by using the VSC datastore remote replication feature. This feature allows the administrator to update the templates in the source datastore and automatically replicate them to one or many destinations by using NetApp SnapMirror® technology. When SnapMirror is used, deduplication should be applied to the template datastore to achieve storage efficiency and thin replication.

The second datastore should contain the actual virtual desktops. Table 5 summarizes the datastore settings for deduplication, backup and recovery, and replication.

Table 5) Datastore settings for deduplication, backup and recovery, and replication.

Datastore Name	Recommended or Optional	Deduplication	Backup and Recovery	Replication
Template datastore	Recommended	Yes	Yes	Yes
VM datastore	Recommended	Yes	Yes, if dedicated and persistent	Yes, if dedicated and persistent
vSwap	Optional	No	No	No

Note: Deduplication and replication of VM datastores for dedicated and persistent desktops might take a significant amount of time to complete, or they might not complete at all. Transient data (such as `pagefile` and other temporary files) that is modified daily can result in gigabytes of data changing per VM every day, significantly lengthening the time that deduplication and replication take to complete.

Note: If using VAAI for cloning of full desktops, the virtual machine template must reside in the destination datastore. It is recommended to copy the templates to each of the datastores. NetApp also recommends that there be a one-pool-to-datastore relationship in this situation.

7.5 Number of Virtual Desktops per Datastore

Two guiding principles should be observed for deciding how many virtual desktops to use per datastore: link speed and deduplication limits.

Link Speed

The faster the link, the more virtual desktops can be added to a datastore, because the current NFS clients in ESX do not distribute packets evenly across the links. A single ESX server connects only to a single datastore over a single Ethernet path, regardless of how many links are aggregated.

When using 1GbE, it is important to have multiple links, use link aggregation, and set load balancing to the Route Based on IP Hash option on the ESX server. The storage administrator can create as many IP aliases as the number of physical links participating in the link aggregation, so that each datastore can be mounted from a different IP address to better use all of the potential paths.

As a general rule, the number of datastores on a single controller should be equal to or be a multiple of the number of links available on the storage controller. The probability for equal use across each link increases with this configuration. For more information, refer to [TR-3749: NetApp Storage Best Practices for VMware vSphere](#).

Deduplication Limits

Deduplication is the process of detecting duplicate data blocks, updating the reference to the blocks, updating the reference counts of the blocks, and deleting the duplicate blocks. In Data ONTAP 8.1, the maximum number of block references increased from 255 to 32,768.

In Data ONTAP releases earlier than 8.1, if a system had 255 identical blocks, they would deduplicate down to one block. If the system had 256 blocks, they would deduplicate to two blocks. As the range of

255 blocks was exceeded, an additional block was referenced for each group of 255 blocks. Figure 10 shows the copy equation for Data ONTAP 8.0.x and earlier releases and an example of the application of this equation.

Figure 10) Unique copy equation and example for Data ONTAP 8.0.x and earlier.

$$\text{Round up } \left[\frac{\# \text{ Desktops}}{255} \right] = \# \text{ Unique VM Copies per Datastore}$$

$$\text{Round up } \left[\frac{1000}{255} \right] = 4$$

In Data ONTAP 8.1, if a system has 32,768 identical blocks, they deduplicate down to one block. If the system has 32,769 blocks, they deduplicate to two blocks. As the range of 32,768 blocks is exceeded, an additional block is referenced for each group of 32,768 blocks. Figure 11 shows the copy equation for Data ONTAP 8.1 and later releases and an example of the application of this equation.

Figure 11) Unique copy equation and example for Data ONTAP 8.1 and later.

$$\text{Round up } \left[\frac{\# \text{ Desktops}}{32,768} \right] = \# \text{ Unique VM Copies per Datastore}$$

$$\text{Round up } \left[\frac{1000}{32,768} \right] = 1$$

Table 6 presents the recommended number of virtual desktops to use per datastore, depending on the version of Data ONTAP involved and on the link speed.

Table 6) Recommended number of VMs per datastore.

Data ONTAP Version	Network Type	Recommended VMs per Datastore
8.0.x or earlier	1GbE	250
8.0.x or earlier	10GbE	250
8.1 or later	1GbE	250
8.1 or later	10GbE	1,250

Note: In Data ONTAP 8.0.x and earlier, the maximum number of deduplicated block references is 255. The number 250 was chosen as a recommendation in Table 6 because it allows a round number of VMs per pool of desktops (POD) when architecting a solution.

In Data ONTAP 8.1 and later, the maximum number of deduplicated block references is 32,768. However, because of the limited bandwidth over a 1GbE link, the number selected as a

recommendation for this network type was 250 VMs per datastore. These are recommendations, not requirements.

7.6 Performance Considerations

Provisioning Time

An important consideration when selecting VMware full clones (without VAAI) for virtual desktop deployments is the time that the initial deployment may take. Because each full clone is a copy of the template VM, significant time must be allocated to provision the desktops before they are required.

The speed of the cloning operation is determined by multiple factors, such as the original template size, the current storage system utilization, the controller model, the link speed, and many other factors. Because NetApp can deduplicate virtual desktops so that they do not consume additional space on the storage controller, it might make sense to preprovision all of the desktops up front to help speed up the provisioning process when new users request desktops.

Reducing Storage Consumption

Table 7 shows the initial storage required to deploy 250, 5,000, and 50,000 virtual desktops when a 25GB template VM is used. It also demonstrates how deduplication can reduce the storage required for large and small virtual desktop environments.

Table 7) Storage consumption based on initial deployment after deduplication.

Users at 25GB per Desktop	250	5,000	50,000
Number of template datastores	1	2	20
Number of desktop datastores	1	20	200
Total storage requirement at 25GB per desktop	6TB	115TB	1.1PB
NetApp storage needed	74GB	148GB	2.9TB
NetApp storage savings	6TB	115TB	1.1PB
NetApp storage savings percentage	99%	99%	99%

Note: The data in Table 7 is based on a 50,000-seat virtual desktop deployment on NetApp storage. For more information about this architecture, refer to [WP-7108: 50,000-Seat VMware View Deployment](#) and [TR-3949: NetApp and VMware View 5,000-Seat Performance Report](#).

Initial Deduplication

After the virtual desktops are provisioned, NetApp recommends running deduplication on the system to reduce the amount of consumed storage. It might be necessary to provision a portion of the desktops, run deduplication, and then provision additional desktops so that not all of the storage is consumed in the datastore. This process should be repeated as many times as necessary.

Note: Depending on the number of virtual desktops and the commonality of the data in the VM datastores, the initial deduplication process might take a considerable amount of time to complete. In lab environments, the process has been observed to take 16 hours or longer to complete on datastores that contain 250 VMs.

Virtual Storage Tier

VST allows customers to benefit from NetApp storage efficiency and to significantly increase I/O performance. VST is natively built into the Data ONTAP OS and works by leveraging block-sharing

technologies (such as NetApp primary storage deduplication and file or volume FlexClone technology) to reduce the amount of required cache and eliminate duplicate disk reads. Because only one instance of any duplicate block is read into the cache, NetApp storage requires less cache than traditional storage solutions.

VMware View implementations can see up to 99% initial space savings when NetApp space-efficient cloning technologies are used. These savings translate into higher cache deduplication and high cache hit rates. VST is especially effective in addressing the simultaneous system boot (or boot storm) and login of hundreds to thousands of virtual desktop systems, which can overload a traditional legacy storage system.

I/O Profile of Full Clone

Unlike virtual desktop clones that rely on hypervisor snapshots, VMware View full clones do not have additional overhead associated with reading and writing metadata to the OS disk (delta file). There is a one-to-one relationship between the guest I/O and the storage I/O, and no special sizing considerations need to be addressed when full clones are used.

7.7 Known Limitations

One of the main benefits of using automated desktop pools is the ability to automatically provision virtual desktops on demand and to use a single management interface to both to create and manage the virtual desktops.

However, unlike linked clones, full clones cannot be refreshed, recomposed, or rebalanced. Full clones also require the administrator to manually or automatically patch each VM individually instead of patching the parent and recomposing the pool.

7.8 Management Considerations

Because full clones can take advantage of NetApp deduplication, caution must be taken when using advanced hypervisor features such as VMware Storage DRS™ and Storage vMotion®.

Technologies that move virtual desktops from one datastore to another, such as Storage DRS, require the VM to consume the total amount of space on the destination datastore that it required before deduplication. If Storage vMotion is used on a VM that is deduplicated, the VM is rehydrated; that is, the VM consumes the same amount of storage on the destination that it required before deduplication.

8 VMware View 5 Native Linked Clones

8.1 Overview

Desktops can be provisioned in many different ways for VMware View solutions that use NetApp storage. Each provisioning method has a different effect on storage consumption, performance, security, and management. VMware View can use View Composer to create storage-efficient virtual desktops called *native linked clones*. The benefits of native linked clones include:

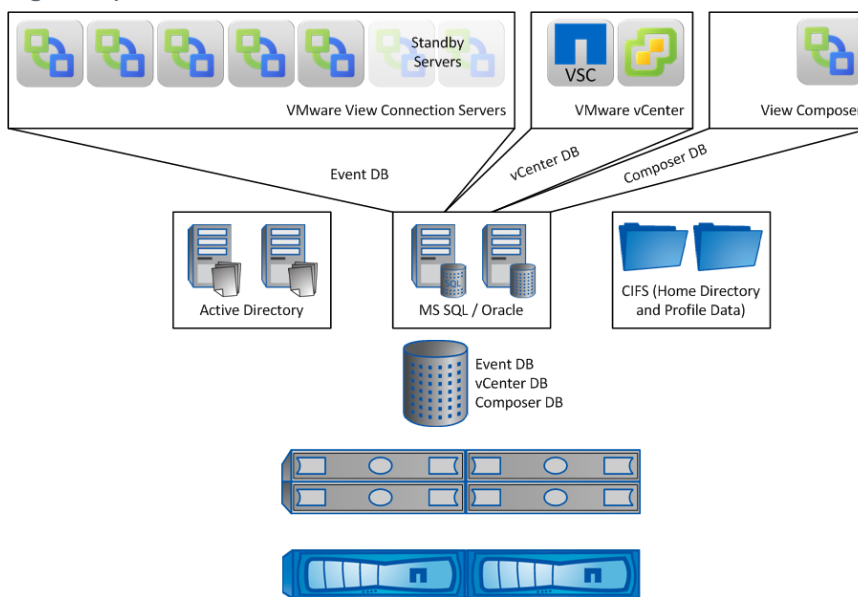
- **Storage efficiency.** Native linked clones use a common C: \ drive for multiple desktops, thus reducing the amount of required storage in comparison to fully provisioned desktops (full clones).
- **Ease of management.** Several functionalities make native linked clones easier to manage:
 - The VMware View administrative interface is the only management interface required to create native linked clones.
 - Desktops can be automatically provisioned based on thresholds or demand.
 - Desktops can be refreshed or reset to their original state at a scheduled time, or they can be manually initiated by an administrator.

- The virtual machine operating system or software can be patched by patching the master VM and recomposing the desktop pool.
- **Increased security and compliance.** Policies can be set on a pool of desktops to reset each desktop to its original state after the user logs off.

8.2 VMware Native Linked-Clone Architecture

Figure 12 shows the server components of a native linked-clone solution. The solution has seven Connection Server instances: Five are designated as active servers and two are designated as standby servers in the event of a failure. The VMware View servers are where the native linked clones are created and where the clients access the individual VMs. The vCenter Server manages the VMs and provides core services, such as VMware high availability. The View Composer server, which can be coinstalled on the vCenter Server, creates the native linked-clone VMs.

Figure 12) Native linked-clone solution architecture.



8.3 Native Linked-Clone Architecture

Native linked clones are composed of at least two disks:

- **Replica disk:**
 - Contains the guest OS and preinstalled applications
 - Is a thin-provisioned copy of the template VM
 - Is shared by all VMs that are in the same pool and therefore are read-only
- **OS disk:**
 - Is commonly referred to as the delta disk (Δ)
 - Is the linked clone in which unique data for each VM is written
 - Contains unique data consisting of temporary files, `pagefile.sys`, web browser caches, temporary Microsoft Office files, and any other data that the user has written to the `C:\` drive of the VM

Having two disks allows a native linked clone to be created in a storage-efficient manner by using the shared replica. A second benefit of having these two disks is that they enable fast reprovisioning or refreshing.

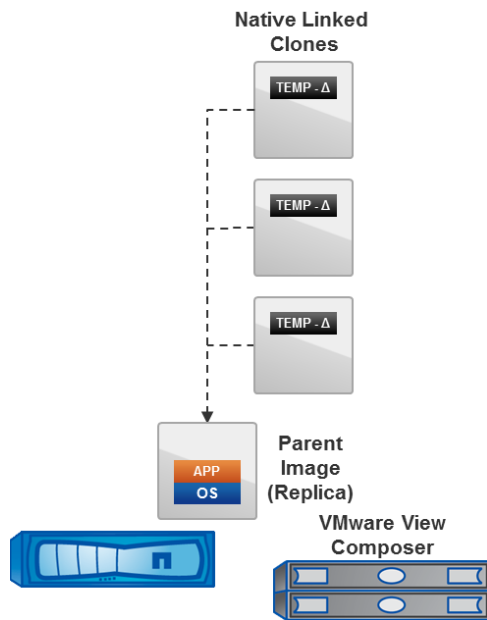
Best Practice

NetApp recommends using the space-efficient sparse virtual disk format (SE Sparse). To use this format of OS disk, the following configuration must be used:

- VMware View 5.2 or later
- vSphere 5.1 patch 1 or later
- VMX virtual machine version 9

Figure 13 shows three native linked-clone OS disks (VMs) represented by $\text{TEMP} - \Delta$ and a common parent image (replica). Figure 13 also refers to View Composer, a critical part of the VMware View native linked-clone architecture. The Composer server creates the native linked-clone VMs.

Figure 13) Native linked-clone architecture.



8.4 Additional Native Linked-Clone Disks

QuickPrep Configuration-Data Disk

The QuickPrep configuration-data disk contains the QuickPrep data and OS-specific configuration information that is generated during the VM provisioning process. This data is used when VMs are refreshed or when a recompose activity is performed. This disk is 20MB and is stored in the same datastore as the OS disk.

View Composer Persistent Disk (Optional)

VMware View Composer persistent disks are used to store the Windows user profile and user data for dedicated desktops. They are persistent across refresh, recompose, and rebalance activities.

Disposable-Data Disk (Optional)

The disposable-data disk allows the redirection of the OS page file and temporary files to a nonpersistent disk. This disk is replaced each time the VM is powered off, to maintain storage efficiency. This disk must be larger than the memory size of the guest VM.

Best Practice

NetApp recommends using a disposable-data disk when deploying native linked clones.

Note: For more information on the additional native linked-clone disks, refer to [VMware View Administration](#).

8.5 Common Maintenance Activities

Three common maintenance activities can be performed on native linked clones: refresh, recompose, and rebalance. Each of these maintenance activities has an effect on both the capacity of the storage controller and the performance of the guest VM.

Refresh

In a VM refresh operation, the VM's OS disk (the delta disk) is reverted back to the state it was in when it was first provisioned. A refresh operation maintains the VM's storage efficiency, security, and compliance by purging any sensitive data that has been written to the C: \ drive.

Best Practice

VMware recommends conducting VM refresh operations during nonpeak hours. For more information about VM refresh operations, refer to [VMware View Administration](#).

Recompose

In a VM recompose operation, the VM's replica disk is updated or replaced, and the VM is then refreshed and relinked to the new or updated replica. This operation is used when the template machine is patched or software updates are applied.

Best Practice

VMware recommends conducting VM recompose operations during nonpeak hours. For more information, refer to [VMware View Administration](#).

Rebalance

In a VM rebalance operation, the VMware View desktops are distributed evenly across datastores.

Best Practice

NetApp recommends not using this option, because it might temporarily increase the storage requirements due to the rehydration of deduplicated data.

8.6 Datastore Design for Native Linked Clones

When designing datastore architectures for VMware VMware View native linked-clone environments, some simple guidelines should be followed:

- Separate data that is transient; that is, data that should not be backed up, replicated, or deduplicated.
- Give VMware View Composer persistent disks their own datastore.

Table 8 summarizes the best practices for deduplication, backup and recovery, and replication in native linked-clone environments.

Table 8) Best practices for deduplication, backup and recovery, and replication.

Datastore Name	Recommended or Optional	Deduplication	Backup and Recovery	Replication
Template datastore	Recommended	Yes	Yes	Yes
Replica datastore	Recommended	Yes	No	No
OS disk datastore (delta disk)	Recommended	No	No	No
VMware View Composer persistent disk (user data disk)	Optional if persistent disks are used	Yes	Yes	Yes
vSwap	Optional	No	No	No

8.7 Performance Considerations

Although native linked clones are efficient from a storage consumption perspective, that efficiency comes with a cost. Native linked clones are based on hypervisor snapshots. Because a native linked clone consists of two disks, the replica and the OS disk (delta disk), when a guest performs a read operation from the `C:\` drive, the hypervisor must first check whether the data resides on the OS disk or on the replica. Therefore, each `C:\` drive guest read I/O equals two storage I/Os.

When a guest performs a write to the `C:\` drive, because the replica is read-only, the hypervisor must first read the metadata, potentially write the metadata if it does not already exist, and then write the actual data block. Therefore, each `C:\` drive guest write (from `pagefile.sys`, temporary files, Office temporary files, browser caches, and so forth) can generate three I/Os to the storage controller.

8.8 Recommendations to Reduce Additional I/Os

When architecting a VMware View environment that uses native linked clones, it is best to look at the environment holistically. Several strategies can be applied to help reduce the performance effects caused by the use of native linked clones. NetApp recommends the following best practices to help VMware View solutions scale much further:

Best Practices

- Use a profile management solution. When I/O is written to `C:\users` data, `C:\temp`, the browser cache, and so on, a profile management solution prevents the I/O from going to the OS disk, where the number of I/Os might potentially be three times the number of I/Os written by the guest.
- Use a CIFS share to store the user data (for example, `H:\` drive).
- Perform a user assessment to understand the guest VM resource requirements. Undersizing guest memory can have a significant performance effect on the page file usage.
- Understand the overcommit limits of the ESXi infrastructure. Overcommitment can cause significant storage performance issues.

8.9 Sizing Considerations

Special sizing issues must be considered when native linked clones are used. To appropriately size the storage for the I/O overhead that is caused by hypervisor snapshots, the NetApp System Performance Modeler has been updated to allow the user to size for different cloning methods. After completing the

general sections of the sizer and selecting the virtual desktop infrastructure workload type, a dialog box appears, as shown in Figure 14.

Figure 14) Selecting the hypervisor, broker, and cloning method dialog box.

System Performance Modeler (SPM) Wizard - Basic Sizing

Step 5

FAQ Links

Virtualization Infrastructure

Hypervisor Type: VMware vSphere

Connection Broker: VMware View

Cloning Method: Linked Clones

Disposable File Disk & Persona Management: ☒ Yes ☐ No

% of Total Writes to DFD & PM Solution: 90%

Workload Details

Workload Description: Helpdesk User

Protocol Type: NFS v3

Free Aggregate Space (%): 10.0

Free SnapReserve Space (%): 0.0

Number of IOPs: 12

Back Finish

In this dialog box, the user selects the hypervisor, broker, and cloning method for the virtualization infrastructure. After the cloning method is selected, the user can select whether to use either a VMware persistent disk (also known as a disposable file disk [DFD]) or a profile management solution for the user data.

If the option Disposable File Disk & Persona Management is selected, a drop-down list becomes available for the user to specify the percentage of total writes that go to the C : \ drive (% of Total Writes to DFD and PM Solution). This percentage is used to address the overhead caused by the hypervisor snapshots. In the example in Figure 14, 90% of all writes go to the DFD or to the profile management solution. Therefore, the remaining 10% incurs the metadata overhead when writing to the OS disk.

9 VMware View 5.1 View Storage Accelerator

9.1 Overview

VMware View Storage Accelerator is a host-based cache that is available in VMware vSphere 5 and is enabled with VMware View 5.1 and later.

The goal of View Storage Accelerator is to reduce the number of input/output operations per second (IOPS) being read from local or shared storage and to improve performance by reducing read latency during activities such as boot storms, login storms, and steady-state operations. VMware View Storage Accelerator has been shown to reduce IOPS by up to 80% and to reduce throughput by up to 65%.

View Storage Accelerator supports any desktops managed by vCenter Server, such as manual desktops, automated full-clone desktops, and automated linked-clone desktops. This includes desktops that are created with NetApp Virtual Storage Console (VSC).

Note: Clones created by View Composer API Integration (VCAI) are not supported by VSA, as documented in the VMware View Administrator Guide.

View Storage Accelerator helps to reduce response time on the operating system disk, and also on any other disk that is provisioned to the virtual machine. View Storage Accelerator works best when multiple disks contain identical content, such as with the linked clone replica disk or any full or FlexClone VMs that were created from the same or a similar template.

View Storage Accelerator is a host-based memory cache that uses ESXi host memory. The cache size can be set from 100MB to a maximum of 2048MB, and it can be set differently from host to host.

Note: View Storage Accelerator cannot be used in combination with View Composer API Integration. <http://pubs.vmware.com/view-52/topic/com.vmware.view.administration.doc/GUID-77B22AC9-EF9F-4161-9856-88DADEE095DD.html#GUID-77B22AC9-EF9F-4161-9856-88DADEE095DD>

Figure 15 shows an example of host cache settings.

Figure 15) Host cache settings.

Host Cache Settings

☒ Enable host caching for View

Default host cache size: MB

Cache must be between 100 MB and 2048 MB

Hosts

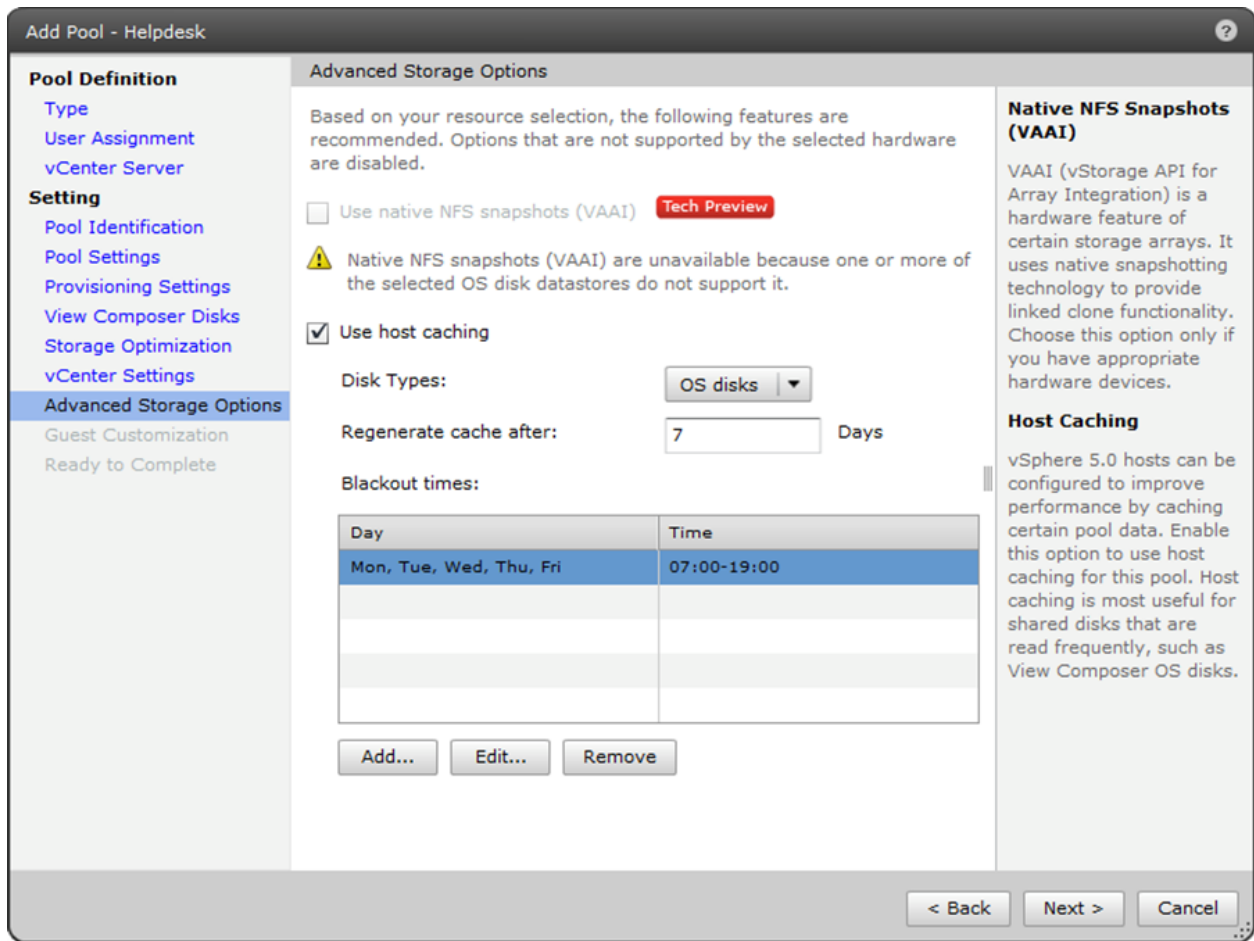
☐ Show all hosts

Host	Cache Size
/East Coast DC/host/VMware View 001/192.168.1.10	Default
/East Coast DC/host/VMware View 001/192.168.1.11	Default
/East Coast DC/host/VMware View 001/192.168.1.12	Default
/East Coast DC/host/VMware View 001/192.168.1.13	Default
/East Coast DC/host/VMware View 001/192.168.1.14	Default
/East Coast DC/host/VMware View 001/192.168.1.15	Default
/East Coast DC/host/VMware View 001/192.168.1.16	Default
/East Coast DC/host/VMware View 001/192.168.1.17	Default

OK Cancel

View Storage Accelerator is enabled on a per-pool basis from the Add Pool wizard. Through the advanced storage options, View Storage Accelerator can be enabled and configured to regenerate the digest intervals, as shown in Figure 16.

Figure 16) Enabling View Storage Accelerator.

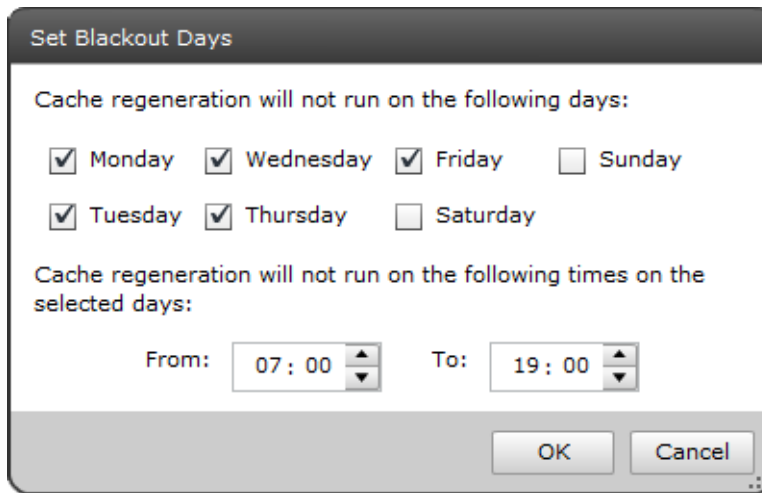


The process of regenerating the digest requires the VM to be powered off. Therefore, blackouts should be scheduled for times when digest regeneration does not occur. Figure 17 shows how to set blackout days.

Best Practice

Set blackout times for off-peak hours, because the VMs must be powered off to regenerate the digest.

Figure 17) Setting blackout days.



The dialog box is titled "Set Blackout Days". It contains two sections. The first section, "Cache regeneration will not run on the following days:", has checkboxes for Monday, Wednesday, Friday, Sunday, Tuesday, Thursday, and Saturday. Monday, Wednesday, Friday, Tuesday, and Thursday are checked. The second section, "Cache regeneration will not run on the following times on the selected days:", has "From:" and "To:" labels followed by time selection boxes. The "From:" box shows "07:00" and the "To:" box shows "19:00". At the bottom are "OK" and "Cancel" buttons.

9.2 Storage Considerations

View Storage Accelerator consumes additional storage. It creates one additional `-digest.vmdk` disk for each virtual machine disk (VMDK). This additional VMDK is a digest disk that contains the hash values for the data VMDKs. These disks consume 0.5% of the space of the original VMDK when using SHA1 hash with a 4K block size, and 1.2% when using collision detection. For example, if the VM is 50GB, the digest is 2.5GB. For more information about using collision detection, refer to [View Storage Accelerator in VMware View 5.1](#).

9.3 Additional Documentation

For more information about VMware View and View Storage Accelerator, refer to the following documents:

- [View Storage Accelerator in VMware View 5.1](#)
- [VMware View 5.1 Reference Architecture for Floating Desktops on NetApp Storage](#)

10 VMware View 5 VSC Clones

10.1 Overview

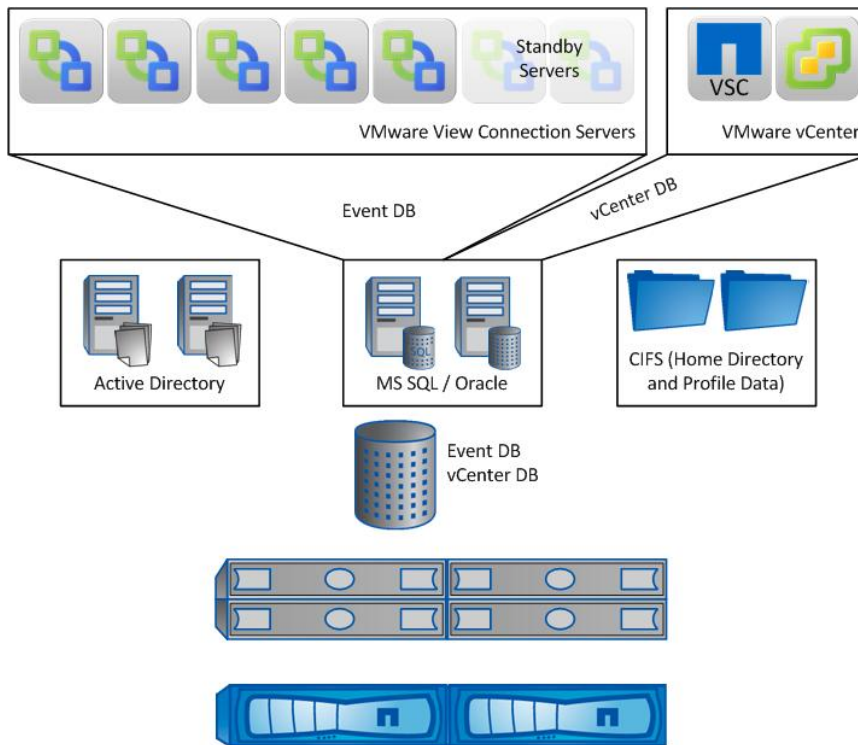
Desktops can be provisioned in many different ways for VMware View solutions that use NetApp storage. Each provisioning method has a different effect on storage consumption, performance, security, and management. Here are some of the benefits provided by clones created through the NetApp Virtual Storage Console (VSC):

- **Storage efficiency.** NetApp VSC-created clones use FlexClone technology to rapidly create prededuplicated virtual desktops that significantly reduce initial storage requirements.
- **Performance.** Performance can be significantly improved by the use of prededuplicated virtual desktops because the operating system is identical. The concept of deduplication and acceleration is called *Virtual Storage Tier (VST)*. VST can help improve boot, login, and steady-state workloads by caching data in a virtual tier (which is commonly the controller memory), in a caching PCIe card, or in a solid-state drive (SSD) based controller cache.

10.2 VMware VSC-Clone Architecture

Figure 18 shows the server components of a VSC-clone solution. The solution has seven VMware View Connection Server instances: Five are designated as active servers and two are designated as standby servers in the event of a failure. VSC clones are created by the VSC and imported into the VMware View Connection broker. The vCenter Server manages the virtual machines and provides core services, such as VMware high availability and Distributed Resource Scheduler. The VSC can be installed on a standalone server or coinstalled with the vCenter Server.

Figure 18) VSC-clone solution architecture.



10.3 Virtual Desktop Machine Architecture

A VSC-created clone contains at least one virtual disk. This virtual disk is a FlexClone copy of the VM template that contains the OS and any corporate-installed patches and applications. This VM is initially storage efficient, due to FlexClone technology. However, over time, it is necessary to continue to deduplicate the datastore in which the VSC clones are created to maintain storage efficiency. Typical deduplication savings of 50% to 85% have been seen in real-world deployments, but the actual storage savings might vary by implementation.

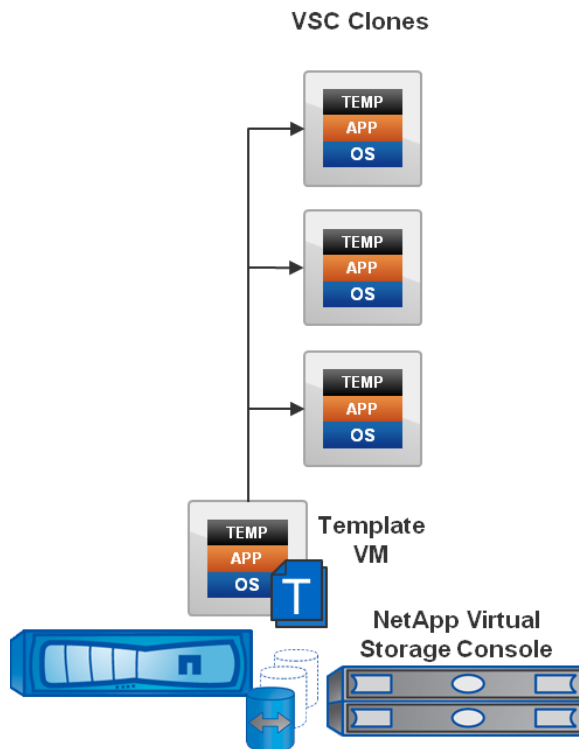
Figure 19 shows three VSC-cloned VMs. Each VM is a full virtual FlexClone copy created from the template VM. In this case, VSC clones the entire virtual desktop from the template VM, including the OS, corporate-installed patches and applications, and any temporary data that resided in the template.

Note: It is important to remove all transient data, such as installers and temporary files, from the template VM. Removing transient data reduces the size of the template and the time it takes to clone it.

Best Practices

- Remove all transient data from the template VM before deploying virtual desktops.
- When using Network File System (NFS) with VSC, perform space reclamation on the template to make the VM as small as possible.

Figure 19) VMware VSC-clone architecture.



10.4 Datastore Design for VSC Clones

The design of datastore architectures for VMware View VSC-clone environments should include at least two datastores to simplify management, improve storage efficiency, and allow different backup, recovery, and disaster recovery policies for each data type.

The first datastore should contain the template VMs. Patching and maintenance of the virtual-desktop templates can take place in this datastore. Using the NetApp VSC, the administrator can create a clone of the original template, make changes to the template, and create new virtual desktops from the newly changed template. The first datastore can then be replicated to remote sites or other environments by using the VSC datastore remote replication feature. This feature allows the administrator to update the templates in the source datastore and automatically replicate them to one or many destinations by using NetApp SnapMirror technology. Deduplication should be applied to the template datastore to achieve storage efficiency and thin replication when SnapMirror is used.

The second datastore should contain the actual virtual desktops. Table 9 summarizes the datastore settings for deduplication, backup and recovery, and replications.

Table 9) Datastore settings for deduplication, backup and recovery, and replication.

Datastore Name	Recommended or Optional	Deduplication	Backup and Recovery	Replication
Template datastore	Recommended	Yes	Yes	Yes
VM datastore	Recommended	Yes	Yes, if dedicated and persistent	Yes, if dedicated and persistent
vSwap	Optional	No	No	No

Note: Deduplication and replication of VM datastores for dedicated and persistent desktops might take a significant amount of time to complete, or they might not complete at all. Transient data (such as `pagefile` and other temporary files) that is modified daily can result in gigabytes of data changing per VM every day, thus lengthening the time that deduplication and replication take to complete.

10.5 Number of Virtual Desktops per Datastore

Two guiding principles should be observed in deciding how many virtual desktops to use per datastore: link speed and deduplication limits.

Link Speed

The faster the link, the more virtual desktops can be added to a datastore, because the current NFS clients in ESX do not distribute packets evenly across the links. A single ESX server connects only to a single datastore over a single Ethernet path, regardless of how many links are aggregated.

When using 1GbE, it is important to have multiple links, use link aggregation, and set load balancing to the Route Based on IP Hash option on the ESX server. The storage administrator can create as many IP aliases as the number of physical links participating in the link aggregation, so that each datastore can be mounted from a different IP address to better use all of the potential paths.

As a general rule, the number of datastores on a single controller should be equal to or be a multiple of the number of links available on the storage controller. The probability for equal use across each link increases with this configuration. For more information, refer to [TR-3749: NetApp Storage Best Practices for VMware vSphere](#).

Deduplication Limits

Deduplication is the process of detecting duplicate data blocks, updating the reference to the blocks, updating the reference counts of the blocks, and deleting the duplicate blocks. In Data ONTAP 8.1, the maximum number of block references increased from 255 to 32,768.

In Data ONTAP releases earlier than 8.1, if a system had 255 identical data blocks, they would deduplicate down to one block. If the system had 256 blocks, they would deduplicate to two blocks. As the range of 255 blocks was exceeded, an additional block would be referenced for each group of 255 blocks. Figure 20 shows the copy equation for Data ONTAP 8.0.x and earlier releases and an example of the application of this equation.

Figure 20) Unique copy equation and example for Data ONTAP 8.0.x and earlier.

$$\text{Round up } \left[\frac{\# \text{ Desktops}}{255} \right] = \# \text{ Unique VM Copies per Datastore}$$

$$\text{Round up } \left[\frac{1000}{255} \right] = 4$$

In Data ONTAP 8.1, if a system has 32,768 identical data blocks, they deduplicate down to one block. If the system has 32,769 blocks, they deduplicate to two blocks. As the range of 32,768 blocks is exceeded, an additional block is referenced for each group of 32,768 blocks. Figure 21 shows the copy equation for Data ONTAP 8.1 and later releases and an example of the application of this equation.

Figure 21) Unique copy equation and example for Data ONTAP 8.1 and later.

$$\text{Round up } \left[\frac{\# \text{ Desktops}}{32,768} \right] = \# \text{ Unique VM Copies per Datastore}$$

$$\text{Round up } \left[\frac{1000}{32,768} \right] = 1$$

Table 10 presents the recommended number of virtual desktops to use per datastore, depending on the version of Data ONTAP involved and on the link speed.

Table 10) Recommended number of VMs per datastore.

Data ONTAP Version	Network Type	Recommended VMs per Datastore
8.0.x or earlier	1GbE	250
8.0.x or earlier	10GbE	250
8.1 or later	1GbE	250
8.1 or later	10GbE	1,250

Note: In Data ONTAP 8.0.x and earlier, the maximum number of deduplicated block references is 255. The number 250 was chosen as a recommendation in Table 10 because it allows a round number of VMs per pool of desktops (POD) when architecting a solution.

In Data ONTAP 8.1 and later, the maximum number of deduplicated block references is 32,768. However, because of the limited bandwidth over a 1GbE link, the number selected as a recommendation for this network type was 250 VMs per datastore. These are recommendations, not requirements.

10.6 Performance Considerations

Provisioning Time

For VSC-created clones in virtual desktop deployments, cloning speed is not an issue as it might be for full clones. VSC leverages NetApp FlexClone technology, which does not copy any of the actual data but instead creates the metadata that refers to the actual data blocks. This method of cloning is much faster.

However, using VSC leverages the customization-specification method for guest customization. This method relies on Sysprep and is much slower than the VMware QuickPrep feature.

Reducing Storage Consumption

Table 11 shows the initial storage that is required to deploy 250, 5,000, and 50,000 full virtual desktops when a 25GB template VM is used. It also demonstrates how provisioning with FlexClone technology or using deduplication can reduce the storage required for large and small virtual desktop environments.

Table 11) Storage consumption based on initial deployment after deduplication.

Users at 25GB per Desktop	250	5,000	50,000
Number of template datastores	1	2	20
Number of desktop datastores	1	20	200
Total storage requirement at 25GB per desktop	6TB	115TB	1.1PB
NetApp storage needed	74GB	148GB	2.9TB
NetApp storage savings	6TB	115TB	1.1PB
NetApp storage savings percentage	99%	99%	99%

Note: The data in Table 11 is based on a 50,000-seat virtual desktop deployment on NetApp storage. For more information about this architecture, refer to [WP-7108: 50,000-Seat VMware View Deployment](#) and [TR-3949: NetApp and VMware View 5,000-Seat Performance Report](#).

Virtual Storage Tier

VST allows customers to benefit from NetApp storage efficiency and significantly increase I/O performance. VST is natively built into the Data ONTAP OS and works by leveraging block-sharing technologies (such as NetApp primary storage deduplication and file or volume FlexClone technology) to reduce the amount of required cache and eliminate duplicate disk reads. Because only one instance of any duplicate block is read into the cache, NetApp storage requires less cache than traditional storage solutions.

VMware View implementations can see up to 99% initial space savings when NetApp space-efficient cloning technologies are used. These savings translate into higher cache deduplication and high cache hit rates. VST is especially effective in addressing the simultaneous system boot (or boot storm) and login of hundreds to thousands of virtual desktop systems, which can overload a traditional legacy storage system.

I/O Profile of VSC Clone

Unlike virtual desktop clones that rely on hypervisor snapshots, VSC-created clones do not have additional overhead associated with reading and writing metadata to the OS disk (delta file). There is a one-to-one relationship between the guest I/O and the storage I/O, and no special sizing considerations need to be addressed when VSC-created clones are used.

10.7 Known Limitations

One of the main benefits of using VMware View is the availability of automated pools. Automated pools give IT the ability to automatically provision virtual desktops on demand, using a centralized management interface to both create and manage the virtual desktops. With VSC-created clones, the VSC must also be used, adding an additional management interface. However, the benefits of increased performance and storage efficiency over other cloning methods can outweigh the management overhead.

Unlike linked clones, VSC-created clones cannot refresh, recompose, or rebalance the virtual desktops. This limitation prohibits refresh-after-logoff operations and requires the administrator to manually or automatically patch each VM individually instead of patching the parent VM and recomposing it.

When using VSC-created clones in a nonpersistent manner, VSC does have a feature that is similar to recompose, called redeploy. This feature allows the administrator to patch or update the template VM and reprovision individual or multiple virtual desktops easily. Because the desktops are entirely deleted and reprovisioned, all data and applications that were installed on the guest VM are deleted.

10.8 Management Considerations

Because VSC-created clones can take advantage of NetApp FlexClone technology and deduplication, caution must be taken when using advanced hypervisor features such as Storage DRS and Storage vMotion.

Technologies that move virtual desktops from one datastore to another, such as Storage DRS, require a VM to consume the total amount of space on the destination datastore that the VM required before deduplication. If Storage vMotion is used on a VM that is deduplicated, the VM is rehydrated; that is, the VM consumes the same amount of storage on the destination that it required before deduplication.

11 VDI User Data Layout

11.1 Overview

A virtual desktop is composed of many different data types, as shown in Figure 22. Because the virtual desktop exists in the data center, administrators can separate the different data types and deliver different service levels to each one. These different service levels might include components such as backup and recovery, deduplication, disaster recovery replication, and many others. Each of the backup and recovery and disaster recovery service levels can have different recovery point objectives and recovery time objectives. This section discusses user-created data and how the different types of user data offer the user an experience similar to that of using a physical desktop or a laptop.

Figure 22) Different data types of a virtual desktop.



11.2 What Is User Data?

In this document, user data is defined as any data that is directly or indirectly created by the end user. User data can be generated by a user-executed application such as Microsoft Office (for example, a Word document), a user-installed application (such as an instant messaging client), the user login process, or system-based tasks (such as virtual machine swapping or browser cache). All of these different data types contribute to the end user's unique experience, and each can have different service-level agreements, based on the importance of the data.

11.3 Different Types of User Data

End-User Data

The first type of user data is end-user data, or home-directory data, as shown in Figure 22. This data is the intellectual property of the company and is directly generated by the end user. In a virtual desktop environment, home-directory data is located on 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. End-user data files should be deduplicated and compressed to achieve storage efficiency and to reduce the overall solution cost.

Best Practice

NetApp recommends deduplicating and compressing end-user data files stored in home directories to gain storage efficiency.

User Profile Data

The second type of user data is the user profile, also known as the persona 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` on a Microsoft Windows physical machine, and it can be redirected to a CIFS share for use in a nonpersistent virtual desktop environment. There are many profile management solutions on the market that simplify management, improve reliability, and reduce the network requirement over standard Windows folder redirection.

Using a profile management solution offers many benefits, including the following:

- Speeding the migration process from a physical desktop or laptop by first virtualizing the profile and then virtualizing the desktop
- Improving login times over folder redirection alone
- Centralizing the end-user profile for better backup, recovery, and disaster recovery of data

Best Practice

NetApp recommends using a third-party profile management solution to allow end users to customize their experience in a nonpersistent desktop environment.

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 be done by using a user-installed application solution, which redirects the data to a CIFS share or a persistent virtual disk.

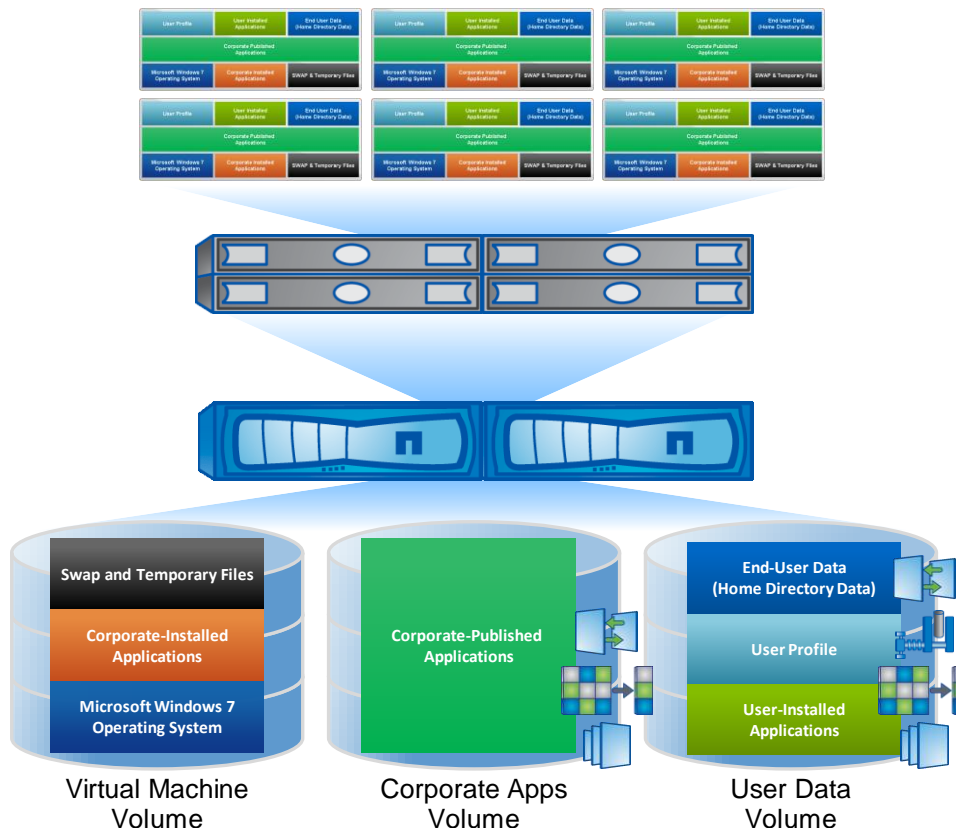
11.4 Separate User Data from the Virtual Machine

Each data type of a virtual desktop has a different purpose, and therefore each type 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 OS and corporate-installed applications in the VM, and it allows administrators to focus on the end users, their experience, and their data. This end-user-centric focus allows administrators to deliver a higher quality desktoplike experience with data center benefits, such as security, compliance, backup and recovery, and disaster recovery.

Figure 23 shows how data can be separated into the different core components. Each of these components can have different policies for DR, backup and recovery, deduplication, and so on. In this example, the administrator is using NetApp Snapshot™ technology, SnapMirror, deduplication, and compression on the user-data volume. No advanced storage capabilities are enabled on the VM volume because the data is transient. The administrator has protected the important data that cannot easily be created but has applied no data protection policies on the VM volume because that data can easily be recreated by reprovisioning the virtual desktop. In this example, all of the user data is stored in one volume, and each user has a personal directory for the home directory data, the user's profile, and the user-installed applications.

Note: The virtual machine volume shown in Figure 23, and other volumes, can be broken down more granularly, depending on the connection broker, the provisioning method, and other factors.

Figure 23) Data type organized by volume.



11.5 Profile Management Considerations

The selection of a profile management solution involves many considerations. Some of the most common questions that should be asked of the profile management vendor fall into the categories of common criteria, Windows native criteria, infrastructure criteria, and data migration criteria.

Common Criteria

- Does the solution support Windows XP, Windows 7, and Windows 8?
- Does it support multiple logins?
- Does it support physical, virtual, or different connection brokers?
- Does it support both 32- and 64-bit OSs?
- Does the solution require an agent to be installed, through SMS or other methods, for each machine that the user uses?
- Does it support location awareness for tasks such as local printing?
- Does it fully support all application virtualization platforms?

Windows Native Criteria

- How is the user profile stored on disk? Is it stored in a database or in a file share?

Note: This is an important question because solutions that use file shares to store the data simplify backup and recovery, allow the data to be stored in a native file format, and prevent vendor lock-in.

- Does the profile management solution allow the end user to easily revert to a standard profile?

Infrastructure Criteria

- Does the solution require an additional VM or physical server that requires additional licensing?
- Is the server for management, or is it required for operation?

Data Migration Criterion

- Does the solution perform the migration of the user data from the original device, PC, laptop, or VM into the profile management solution?

12 System Performance Modeler for VMware View

For VMware View desktop deployments to be successful, an assessment and sizing must be performed. The assessment captures information about what the users are doing to determine the infrastructure requirements. Assessment software, such as Liquidware Labs Stratusphere Fit, optionally can be used to collect this data. In most cases, user and application behavior is not something that can be estimated or generalized because each user interacts with data and applications differently. Therefore, assessments are necessary because they provide the data that is required to size the storage infrastructure appropriately. NetApp uses a tool called the [System Performance Modeler](#) to size virtual desktop infrastructures.

NetApp provides the Liquidware Labs Stratusphere Fit assessment software to our field and partners free of charge. It can assess up to 500 desktops for 45 days, and the license can be renewed. For more information, refer to [Liquidware Labs for VDI Profiling](#) on Field Portal.

To summarize, VMware View enables organizations to increase corporate IT control, manageability, and flexibility without increasing cost, while providing end users with a familiar desktop experience. The NetApp key value proposition of at least 50% savings in storage, power, and cooling requirements; performance acceleration; operational agility; and a best-in-class data protection and business

continuance solution makes it the solution of choice for storage and data management for VMware View. The key NetApp technologies (RAID-DP®, thin provisioning, space reclamation, FlexClone, deduplication, Snapshot, and SnapMirror) provide the foundational strengths to support these claims.

This guide has provided detailed guidance on how to architect, implement, and manage a large, scalable VMware View solution on NetApp storage. It also describes the best integration points for each of the key enabling NetApp technologies and how all of the technology concepts play a critical role and complement each other to work together as an integrated NetApp solution for VMware View at any scale.

This guide is not intended to be a definitive implementation or solutions guide. Expertise is required to solve issues with specific deployments. Contact your NetApp representative and make an appointment to speak with one of our VMware View solutions experts.

13 Feedback

Send an e-mail to xdl-vtme@netapp.com with questions or comments concerning this document.

References

NetApp Documents

- TR-3949: NetApp and VMware View 5,000-Seat Performance Report
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60932-16&m=tr-3949.pdf
- Flash Cache Technical Specifications
www.netapp.com/us/products/storage-systems/flash-cache/flash-cache-tech-specs.html
- High Availability
www.netapp.com/us/products/platform-os/active-active.html
- TR-3001: A Storage Networking Appliance
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60538-16&m=tr-3001.pdf
- TR-3066: Data Protection Strategies for NetApp Storage Systems
http://media.netapp.com/documents/wp_3066.pdf
- TR-3107: NetApp Antivirus Scanning Best Practices Guide
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61066-16&m=tr-3107.pdf
- TR-3298: RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60325-16&m=tr-3298.pdf
- TR-3347: A Thorough Introduction to FlexClone Volumes
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60767-16&m=tr-3347.pdf
- TR-3348: Block Management with Data ONTAP 7G: FlexVol®, FlexClone, and Space Guarantees
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-59763-16&m=tr-3348.pdf
- TR-3367: NetApp Systems in a Microsoft Windows Environment
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61520-16&m=tr-3367.pdf
- TR-3437: Storage Subsystem Resiliency Guide
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60064-16&m=tr-3437.pdf
- TR-3446: SnapMirror Async Overview and Best Practices Guide
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61273-16&m=tr-3446.pdf
- TR-3450: Active-Active Controller Overview and Best Practices Guidelines
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61088-16&m=tr-3450.pdf
- TR-3459: FlexShare® Design and Implementation Guide
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60718-16&m=tr-3459.pdf
- TR-3466: Open Systems SnapVault® (OSSV) Best Practices Guide
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60482-16&m=tr-3466.pdf

- TR-3505: NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60084-16&m=tr-3505.pdf
- TR-3563: NetApp Thin Provisioning
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61045-16&m=tr-3563.pdf
- TR-3747: NetApp Best Practices for File System Alignment in Virtual Environments
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61563-16&m=tr-3747.pdf
- TR-3749: NetApp and VMware vSphere Storage Best Practices
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-104655-16&m=tr-3749.pdf
- TR-3770: VMware View on NetApp Deployment Guide Using NFS
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61199-16&m=tr-3770.pdf
- TR-3801: Introduction to Predictive Cache Statistics
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60523-16&m=tr-3801.pdf
- TR-3808: VMware vSphere and ESX 3.5 Multiprotocol Performance Comparison Using FC, iSCSI, and NFS
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60368-16&m=tr-3808.pdf
- TR-4068: VMware vSphere 5 on NetApp Clustered Data ONTAP 8.1
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61602-16&m=tr-4068.pdf

VMware Documents

- Comprehensive Virtual Desktop Deployment with VMware and NetApp
www.vmware.com/files/pdf/partners/netapp-vmware-view-wp.pdf
- VMware View 5.2 Documentation Center
<http://pubs.vmware.com/view-52/index.jsp>
- VMware Horizon View Optimization Guide for Windows 7
www.vmware.com/files/pdf/VMware-View-OptimizationGuideWindows7-EN.pdf
- Virtual Storage Console 4.1 for VMware vSphere Installation and Administration Guide
https://library.netapp.com/ecm/ecm_download_file/ECMP1118859
- VMware View Administration
<http://pubs.vmware.com/view-50/topic/com.vmware.ICbase/PDF/view-50-administration.pdf>
- View Storage Accelerator in VMware View 5.1
www.vmware.com/files/pdf/techpaper/vmware-view-storage-accelerator-host-caching-content-based-read-cache.pdf
- VMware View 5.1 Reference Architecture for Floating Desktops on NetApp Storage
www.vmware.com/files/pdf/techpaper/vmware-view-reference-architecture-floating-netapp-storage.pdf

Version History

Version	Date	Document Version History
Version 1.0	June 2013	Original document
Version 1.1	November 2013	Updated recommendations for Full Cloning and Horizon View

About the Author

Chris Gebhardt has been with NetApp since 2004 and is currently a senior technical marketing engineer, leading the end-user computing solutions for the NetApp VMware Ecosystem Organization. Chris has coauthored these documents:

- [TR-3705: NetApp and VMware View Solution Guide](#)
- [TR-3770: VMware View on NetApp Deployment Guide Using NFS](#)
- [WP-7108: NetApp FAS2050HA Unified Storage: A Guide to Deploying NetApp FAS2050HA with VMware View and the 50,000-Seat VMware View Deployment Whitepaper](#)
- [TR-3949: NetApp and VMware View 5,000-Seat Performance Report](#)
- [WP-7190: 800-Seat Performance Report](#)

Chris is a contributor to the [NetApp Virtualization Blog](#). He is a VMware vExpert 2010, 2011, 2012, 2013, VMware Certified Professional 3, 4 and 5, Brocade Certified Fabric Professional, Brocade Certified SAN Designer, and NetApp Certified Implementation Engineer as well as a NetApp Accredited Virtual Desktop Infrastructure Professional. Prior to joining the Virtualization Business Unit, he was a Professional Service consultant with NetApp and was the central area practice lead for Network Storage Consolidation and Virtualization Practice, where he authored many deployment guides. Before joining NetApp, Chris was a UNIX® and NetApp administrator for seven years at a worldwide telecommunications company.

Acknowledgements

The author would like to thank Neil Glick and Abhinav Joshi for their contributions to this solutions 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®](#)