



Technical Report

NetApp and VMware View (VDI) Solution Guide

Best Practices for Design, Architecture, Deployment, and Management

Abhinav Joshi and Chris Gebhardt | NetApp
March 2009 | TR-3705 | Version 3.0

ABSTRACT

This document provides NetApp best practices on designing, architecting, deploying, and managing a scalable VMware® View (VDI) environment on NetApp® storage.

TABLE OF CONTENTS

1	EXECUTIVE SUMMARY	4
2	SCOPE	4
3	INTRODUCTION TO VMWARE VIEW	4
4	VMWARE VIEW DEPLOYMENT MODELS	5
4.1	VMWARE VIEW MANAGER DESKTOP DELIVERY MODELS	5
5	NETAPP SOLUTION AND COMPONENTS.....	7
5.1	STORAGE EFFICIENCY	7
5.2	PERFORMANCE	9
5.3	OPERATIONAL AGILITY	10
5.4	DATA PROTECTION	12
6	NETAPP AND VMWARE VIEW MANAGER DESKTOP DELIVERY MODELS.....	13
6.1	INDIVIDUAL DESKTOP	15
6.2	MANUAL DESKTOP POOL	15
6.3	AUTOMATED DESKTOP POOL	16
7	ACCELERATING VMWARE VIEW WITH INTELLIGENT CACHING	17
7.1	CONCEPTS	17
7.2	BENEFITS OF INTELLIGENT CACHING	18
7.3	DEDUPLICATION & NON-DUPLICATION TECHNOLOGIES	19
7.4	INTELLIGENT CACHING DETAILS	20
7.5	SUMMARY OF INTELLIGENT CACHING IN A VMWARE VIEW ENVIRONMENT	24
8	STORAGE SIZING BEST PRACTICES	25
8.1	GATHER ESSENTIAL SOLUTION REQUIREMENTS	25
8.2	PERFORMANCE-BASED AND CAPACITY-BASED STORAGE ESTIMATION PROCESS	27
9	STORAGE ARCHITECTURE BEST PRACTICES.....	37
9.1	STORAGE SYSTEM CONFIGURATION BEST PRACTICES.....	37
10	DEPLOYING NETAPP SPACE-EFFICIENT VM CLONES ON NFS	40
10.1	OVERVIEW OF DEPLOYING NETAPP SPACE-EFFICIENT CLONES ON NFS	40
10.2	DETAILS OF DEPLOYING NETAPP SPACE-EFFICIENT CLONES ON NFS	41
11	DEPLOYING NETAPP SPACE-EFFICIENT VM CLONES ON VMFS	53
11.1	OVERVIEW OF DEPLOYING NETAPP SPACE-EFFICIENT CLONES ON VMFS	53
11.2	DETAILS OF DEPLOYING NETAPP SPACE-EFFICIENT CLONES ON VMFS	54
12	VMWARE VIEW OPERATIONS BEST PRACTICES	62
12.1	DATA DEDUPLICATION	62
12.2	ANTIVIRUS OPERATIONS	63
12.3	MONITORING NETAPP AND VMWARE VIEW INFRASTRUCTURE	64

12.4	DATA PROTECTION SOLUTION.....	64
13	SUMMARY	65
14	ACKNOWLEDGEMENTS.....	66
15	FEEDBACK.....	66
16	REFERENCES	66
17	VERSION HISTORY	67

1 EXECUTIVE SUMMARY

NetApp technology enables companies to optimize their virtual infrastructures by providing advanced storage and data management capabilities. NetApp provides industry-leading storage solutions that simplify virtual machine (VM) provisioning, reduce virtual desktop storage costs, enable mass VM cloning, provide individual VM backup and restores, and deliver simple and flexible business continuance for virtual infrastructures.

This solution best practices guide provides guidelines and best practices for architecting, deploying, and managing VMware View (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 solution best practices guide describes them in detail.

Note: The best practices are only recommendations, not requirements. Following these best practices will not determine whether your implementation is supported by NetApp and VMware. Not all recommendations apply to every scenario. NetApp believes that its customers will benefit from thinking through these recommendations before making any implementation decisions. In addition to this document, professional services are available through NetApp, VMware, and our joint partners. These services can be an attractive means to assure optimal virtual storage architecture for your virtual infrastructure.

The target audience for this paper is familiar with concepts pertaining to VMware ESX Server 3.5, VMware vCenter Server 2.5, and NetApp Data ONTAP® 7.X. For additional information and an overview of the unique benefits that are available when creating a virtual infrastructure on NetApp storage, see [NetApp TR-3515: NetApp and VMware ESX Server 3.0: Building a Virtual Infrastructure from Server to Storage](#).

2 SCOPE

The scope of this document is to provide architectural, deployment, and management guidelines for customers planning or have already decided to implement VMware View on NetApp virtualized storage. It provides a brief overview of the VMware View technology concepts; key solution architecture considerations for implementing VMware View on NetApp; storage estimation and data layout recommendations; and solution, deployment, and management guidelines.

3 INTRODUCTION TO VMWARE VIEW

Corporate IT departments are facing a new class of desktop management issues as they strive to provide end users with the flexibility of accessing corporate IT resources using any device from any network. IT is also being asked to provide access to corporate resources for an increasingly dispersed and growing audience that includes workers in off-site facilities, contractors, partners, and outsourcing providers as well as employees traveling or working from home. All of these groups demand access to sensitive corporate resources, but IT must make sure of strict adherence to corporate security requirements and new regulatory standards.

VDI enables organizations to increase corporate IT control, manageability, and flexibility of desktop resources while providing end users with a familiar desktop experience. VMware View is an enterprise-class solution to deliver corporate PC services to end users. VMware View is not a product, but a technology concept consisting of distinct components that make up the solution:

- Virtualization hypervisor (ESX Server)
- Virtualized desktop images (for example, Windows® XP, Vista, and so on)
- Tool for centralized management, automation, provisioning and optimization (VMware vCenter, VMware View Composer)
- Connection broker (VMware View Manager)
- Thin client/PC (for example, Wyse, DevonIT)

VMware View, based on the proven VMware virtualization platform, delivers unique desktop control and manageability, while providing end users with a familiar desktop experience without any modifications to the desktop environment or applications.

4 VMWARE VIEW DEPLOYMENT MODELS

For any customer environment, the types of desktop can be classified as either persistent or non-persistent.

PERSISTENT DESKTOPS

Persistent desktop can be defined as desktops that are permanently assigned to a user, are customizable, and no other user is entitled to use this desktop. The user logs in to the same desktop every day, and the changes made to the system image, for example, the C drive, are saved across login sessions and reboots. This is exactly like a laptop or workstation with all the personalization and personal data stored locally on the C: drive. This is a common type of VDI deployment model that is used today for knowledge workers, mobile workers, and power users and is a major driver for increased shared storage requirement.

NONPERSISTENT DESKTOPS

Non-persistent desktops can be defined as desktops that are not assigned to a specific user. The user could get assigned to a different desktop at every login. This deployment model could be used for task workers (for example, call centers) and some knowledge workers who require little control of their desktops

One may choose to implement either of these models or a mix based on the business requirements, types of users, and proportion of users represented by different job functions.

4.1 VMWARE VIEW MANAGER DESKTOP DELIVERY MODELS

[VMware View Manager](#) is the VMware virtual desktop management solution that improves control, manageability, and provides a familiar desktop experience.




Figure 1) VMware View Features (Source: VMware)

At a high level, there are multiple types of desktop delivery models available in VMware View Manager:

- **Individual Desktops:** This can be classified as a pre-assigned desktop that has been pre-created, populated in VMware vCenter Server, imported into VMware View Manager, and assigned to one or more users. The users log in to the same desktop every time, and the user settings and data on the OS disk are preserved across login sessions and reboot. However, it allows only one active user on the desktop at a time. Individual desktops are considered as “non-persistent” because even though a user logs into same desktop every time, other users are also entitled to use that desktop (although not at the same time).
- **Manual Desktop Pools:** This is a pool of desktops and provides mapping of multiple users to multiple desktops, with only one active user on a desktop at a time. These types of desktops also have to be **manually pre-created, populated in VMware vCenter Server, and imported** into VMware View Manager. The desktop assignment happens at the user login.

The manual desktop pool supports two types of access modes:

- **Persistent:** Users are assigned a desktop that retains all of their documents, applications, and settings between sessions. The desktop is statically assigned the first time the user connects and is then used for all subsequent sessions.
 - **Non-persistent:** Users are not assigned to particular desktops and could get connected to a different desktop from the pool each time they connect. Also, there is no persistence of environmental or user data between sessions.
- **Automated Desktop Pools:** This is also a pool of desktops and provides mapping of multiple users to multiple desktops, with only one active user on a desktop at a time. The creation and customization of these types of desktops are driven by VMware View Manager and optionally VMware View Composer by using either of these two options:
 - **Full Clone:** Leveraging VMware vCenter Server virtual machine template to create VMware full clones
 - **Linked Clone:** Leveraging VMware View Composer feature in VMware vCenter Server to create VMware linked clones

Both these options in automated desktop pools support the two types of access modes:

- **Persistent:** Users are assigned a desktop that retains all of their documents, applications, and settings between sessions. The desktop is statically assigned the first time the user connects and is then used for all subsequent sessions.
 - **Non-persistent:** Users are not assigned to particular desktops and could get connected to a different desktop from the pool each time they connect. Also, there is no persistence of environmental or user data between sessions.
- **Terminal Server Pool:** This is a pool of terminal server desktop sources served by one or more terminal servers. Discussion on the storage best practices for this type of desktop delivery model is outside the scope of this document.

Table 1) Desktop delivery models.

Desktop Delivery Models for virtual desktops	Virtual Desktop Provisioning Method	Virtual Desktop Persistence Mode
Individual Desktop	VMs pre-created and available in VMware vCenter Server	Non-persistent
Manual Desktop Pool	VMs pre-created and available in VMware vCenter Server	Persistent
		Non-persistent

Desktop Delivery Models for virtual desktops	Virtual Desktop Provisioning Method	Virtual Desktop Persistence Mode
Automated Desktop Pool	Invoke VMware full clone to create VMs	Persistent
		Non-persistent
	Invoke VMware linked clone to create VMs	Persistent
		Non-persistent

For more details on VMware View Desktop Delivery Models and access modes, refer to the [VMware View Manager Administrator Guide](#).

5 NETAPP SOLUTION AND COMPONENTS

NetApp provides a true single, scalable, unified storage and data management solution for VMware View. The unique benefits of the NetApp solution are:

- **Storage efficiency** – Significant cost savings with multiple levels of storage efficiency.
- **Performance** – Enhanced user experience with intelligent caching that strongly complements NetApp's storage efficiency capabilities.
- **Operational agility** – Enhanced VMware View solution management with tight partner integration.
- **Data protection** – Ensure protection of both the virtual desktop OS data and the user data, with very low overhead for both cost and operations.

SINGLE SCALABLE UNIFIED ARCHITECTURE

The NetApp unified storage architecture provides customers with the most agile and scalable storage and data management platform in the industry. NetApp's innovative solutions deliver new alternatives to expand possibilities for the customer. All NetApp storage systems are configured with the Data ONTAP operating system to satisfy customer needs: SAN (FC, iSCSI), NAS (CIFS, NFS), primary storage, and secondary storage within a single platform. A single process for activities such as installation, provisioning, mirroring, backup, and upgrading is used throughout the entire product line. Having a single set of software and processes brings great simplicity to even very complex enterprise data management challenges. Unifying storage and data management software and processes reduces the complexity of data ownership, enables companies to adapt to changing business conditions without interruption, and results in a dramatic reduction in total cost of ownership.

For large, scalable VMware View environments, the NetApp solution provides the following unique benefits:

- At least 50% savings in storage, power, and cooling requirements
- Most agile and operationally efficient storage solutions
- Best-in-class data protection and business continuance solutions to address any level of data availability demands

5.1 STORAGE EFFICIENCY

One of the critical barriers to VMware View adoption is increased cost, specifically associated with increased shared storage requirement. Virtual machine deployment creates a high level of data redundancy especially for the VM OS data. Using traditional storage, this means you need storage equal to the sum of the storage required by each VM. For example, if each VM is 20GB in size and there are supposed to be 1,000 VMs in the solution, it would require at least 20TB usable data on the shared storage.

Thin provisioning, data deduplication, and FlexClone® are the critical components of the NetApp solution and offer multiple levels of storage efficiency. This helps customers save on average 50% to 90% on the cost associated with shared storage (based on existing customer deployments and NetApp solutions lab

validation). NetApp is the only storage vendor that offers data deduplication on the primary storage for several years, without any negative tradeoffs.

THIN PROVISIONING

Thin provisioning is a way of presenting more storage space to the hosts connecting to the storage system than has been physically allocated. With thin provisioning, the storage administrator is able to utilize a pool of physical disks (known as aggregate) and create logical volumes for different applications to use, but does not actually allocate any space upfront to those volumes. The space gets allocated only when the host needs it. The unused aggregate space is available for the existing thinly provisioned volumes to expand or for use in creation of new volumes. For more details on thin provisioning, review [NetApp TR-3563: NetApp Thin Provisioning](#).





Figure 2) Traditional vs. thin provisioning.



Source: Oliver Wyman Study: "Making Green IT a Reality." November 2007.
 *Thin Provisioning, clones, & multiprotocol all contribute to savings.

Figure 3) Increased disk utilization with NetApp thin provisioning.

NETAPP DEDUPLICATION

NetApp deduplication saves space on primary storage by removing redundant copies of blocks within a volume hosting hundreds of virtual desktops. This process is transparent to the application and can be enabled and disabled on the fly. In a VMware View environment, this solution provides great value when we consider that each VM has its own identical copy of the operating system (OS) and applications and that the same patches will be applied to each VM. For more information on NetApp deduplication, refer to [NetApp TR-3505: NetApp Deduplication for FAS, Deployment and Implementation Guide](#).




Figure 4) NetApp deduplication.

FLEXCLONE

NetApp FlexClone technology allows rapid creation of space-efficient, writable, point-in-time images of individual files, LUNs, or flexible volume. The use of the FlexClone technology in VMware View deployments with thousands of VMs offers high levels of scalability and significant saving in dollars, space, and energy. Both the file-level and volume-level cloning is tightly integrated with VMware vCenter Server using the NetApp Rapid Cloning Utility vCenter plugin. This gives you the flexibility to provision thousands of VMs rapidly with hundreds of VMs in each datastore. Note that from a scalability and manageability perspective, this allows you to only have very few datastores to provision and manage as compared to other solutions, which might require one datastore per VM. FlexClone volumes add a new level of agility and efficiency to storage operations. They take only a few seconds to create and are created without interrupting access to the parent FlexVol® volume or files. FlexClone copies share the same physical data space with the source and occupy negligible extra space in the storage system. FlexClone file-level or volume-level clones use space very efficiently, leveraging the Data ONTAP architecture to store only data that changes between the source and clone. In addition to all these benefits, file-level or volume-level FlexClone has the same high performance as other FlexVol volumes or files hosted on the volumes. Also, FlexClone technology provides enormous benefits for disaster recovery (DR) testing. DR testing with FlexClone is safe, risk free, and can be done during operational hours at any time. For more information on FlexClone technology concepts, see [NetApp TR-3347: FlexClone Volumes: A Thorough Introduction](#).

5.2 PERFORMANCE

Another critical barrier to VMware View adoption is performance concerns 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. Traditionally this was not a problem because each desktop had its own disks and I/O was contained within a single desktop. With VMware View using a shared storage infrastructure, significant performance issues might arise during these critical operations. This essentially means requiring lots of disks on the shared storage to meet the performance requirements, resulting in increased overall solution cost.

To solve this problem the NetApp solution provides intelligent caching capabilities. Intelligent caching is a core component of Data ONTAP and is extended with the Performance Acceleration Module (PAM). These solution components save customers money by:

- Requiring far less disks and cache
- Providing better throughput, system utilization
- Faster response times and a better overall end user experience

DATA ONTAP INTELLIGENT CACHING

Data ONTAP intelligent caching allows customers to benefit from NetApp's storage efficiency capabilities and at the same time significantly increasing I/O performance. It is natively built into the Data ONTAP

operating system and works by leveraging block sharing technologies such as NetApp primary storage deduplication and file/volume FlexClone to reduce the amount of cache required and eliminate duplicate disk reads. Only one instance of any duplicate block is read into cache, thus requiring less cache than traditional storage solutions. Since VMware View implementations can see as great as 99% initial space savings (validated in the NetApp solutions lab) using NetApp space-efficient technologies, this translates into higher cache deduplication and very high cache hit rates. Intelligent caching is especially effective in addressing the simultaneous system boot or “boot storm” of hundreds to thousands of virtual desktop systems that can overload a traditional storage system.

PERFORMANCE ACCELERATION MODULE (PAM)

The NetApp PAM is additional intelligent read cache that extends the intelligent caching abilities of Data ONTAP for virtual desktop performance, including boot times for thousands of VMs in VMware View environments. More details of the PAM module will be discussed later in this document. For more details on NetApp PAM technology, visit www.netapp.com/us/products/storage-systems/performance-acceleration-module/performance-acceleration-module.html.

FLEXIBLE VOLUMES AND AGGREGATES

Flexible volumes (also known as FlexVol volumes) and aggregates provide pools of storage in a way that all the IOPS and storage capacity is available in the pool for thousands of virtual desktops in VMware View to utilize on demand. This can be compared to the virtualization on servers where all the memory and CPU resources are virtualized and available on demand. NetApp virtualizes storage resources like server virtualization hypervisors do to the server resources.

5.3 OPERATIONAL AGILITY

Implementation and management complexities associated with deploying VMware View solution are another critical barrier to VDI adoption. The NetApp storage solution is operationally agile and provides tight integration with VMware for rapidly provisioning thousands of VMs directly from the VMware vCenter Server.

RAPID CLONING UTILITY (RCU)

NetApp Rapid Cloning Utility provides customers with the ability to rapidly provision thousands of VMs in minutes, leveraging the file and volume FlexClone technology discussed earlier. It is tightly integrated with VMware vCenter Server as a plug-in and provides capability to import the VMs directly into VMware View Manager 3.0 or higher. RCU provides integration of the NetApp storage efficiency and rapid cloning capabilities with the VMware APIs to provide a highly cost-effective VM provisioning solution.

INTEGRATION WITH VMWARE VCENTER SERVER AND VMWARE VIEW MANAGER

RCU provides a way to provision VMs directly from VMware vCenter Server and is integrated with VMware View Manager to allow import the VMs into VMware View Manager. The imported VMs can be managed as either individual desktops or manual desktop pools. Therefore, customers are able to achieve the desired storage efficiency and rapid provisioning associated with both persistent and non-persistent desktops and still be able to achieve performance acceleration with NetApp intelligent caching.




Figure 5) NetApp RCU plug-in in VMware vCenter Server.

SANSCREEN VM INSIGHT

Also available as a VMware vCenter Server plug-in, the NetApp SANSscreen® VM Insight also provides cross-domain visibility from the VM to the shared storage, allowing both storage and server administration teams to more easily manage their VMware View storage and server infrastructure. VM Insight provides service-level information for virtual servers, physical servers, and storage devices, as well as VM volume, allocated capacity, and datastore information from VMware vCenter Server. Based on this visibility, VM Insight shows the actual service paths and server performance information to allow end-to-end monitoring. The enterprise-class data warehouse provides IT the ability to access, query, and analyze VM data; and when deployed with SANSscreen Capacity Manager, it enables capacity planning and chargeback for VM environments. For further details on SANSscreen VM Insight, visit the following solutions page: www.netapp.com/us/products/management-software/sansscreen/sansscreen-vm-insight.html.




Figure 6) NetApp SANSscreen VM Insight plug-in in VMware vCenter Server.

OPERATIONS MANAGER

NetApp Operations Manager provides a comprehensive monitoring and management solution for the VMware View storage infrastructure. It provides comprehensive reports of utilization and trends for capacity planning and space usage. It also monitors system performance, storage capacity, and health to resolve potential problems. For further details on Operations Manager, visit the following solutions page:

www.netapp.com/us/products/management-software/operations-manager.html.




Figure 7) NetApp Operations Manager.

SNAPMANAGER FOR VIRTUAL INFRASTRUCTURE

NetApp's SnapManager® for Virtual Infrastructure (SMVI) is a unique, scalable, data protection solution for VMware View. SMVI integrates VMware snapshot functionality with NetApp Snapshot™ and is a complete solution for protecting VMware View environments. For further details, read the following section on data protection.

5.4 DATA PROTECTION

The availability of thousands of virtual desktops is dependent on the availability of shared storage. Thus, the RAID level of the storage is very critical. Also, being able to protect the virtual desktop images and/or user data is very important. RAID-DP® and SMVI are critical components of the NetApp solution that help address storage availability.

RAID-DP

With any VMware View deployment, data protection is critical, because any RAID failure could result in hundreds to thousands of end users being disconnected from their desktops. NetApp FAS systems provide a level of data protection that is unique and unmatched in the storage industry. RAID-DP provides performance that is comparable to that of RAID 10, with much higher resiliency. It provides protection against double disk failure as compared to RAID 5, which can only protect against one disk failure. For more information on RAID-DP, see [NetApp TR-3298: RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection](#).

SNAPMANAGER FOR VIRTUAL INFRASTRUCTURE

NetApp SnapManager for Virtual Infrastructure (SMVI) is a unique, scalable, integrated data protection solution and is an excellent choice for protecting VMware View environments. SMVI integrates VMware snapshot functionality with the NetApp array-based block-level Snapshot copies to provide consistent

backups for the virtual desktops. SMVI is NetApp deduplication aware and also integrates with the NetApp SnapMirror® replication technology that preserves the deduplicated storage savings from the source and destination storage arrays. Deduplication is not required to be rerun on the destination storage array. SMVI also provides a user-friendly GUI that can be used to manage the data protection schemes. Additionally, when a VMware View environment is replicated with SnapMirror, the replicated data can be quickly brought online to provide production access in the event of a site or data center outage. Also, the SnapMirror technology is fully integrated with VMware Site Recovery Manager (SRM) and NetApp FlexClone technology to instantly create writable zero cost copies of the replicated virtual desktops at the remote site that can be used for DR testing or for test and development work. For more information on SnapMirror, see [NetApp TR-3446: SnapMirror Best Practices Guide](#). For more information on VMware SRM Integration, see [NetApp TR-3671: VMware Site Recovery Manager in a NetApp Environment](#). For more detailed information on SMVI, see [NetApp TR-3737: SMVI Best Practices](#).




Figure 8) NetApp SMVI

6 NETAPP AND VMWARE VIEW MANAGER DESKTOP DELIVERY MODELS

At a high level, there are three virtual desktop delivery models available in VMware View Manager:

- Individual Desktops
- Manual Desktop Pool
- Automated Desktop Pool

Both the **manual desktop pool** and **automated desktop pool** models offer two user access modes (persistent and non-persistent). Automated desktop pool offers two options in which virtual desktops can be provisioned: using VMware full clones and VMware linked clones.

Considering this flexibility in the architecture, there are a total of seven ways in which virtual desktops can be provisioned and managed in VMware View Manager:

1. Individual desktops in non-persistent access mode
2. Manual desktop pool in persistent access mode
3. Manual desktop pool in non-persistent access mode
4. Automated desktop pool, leveraging VMware full clones, in persistent access mode
5. Automated desktop pool, leveraging VMware full clones, in non-persistent access mode
6. Automated desktop pool, leveraging VMware linked clones, in persistent access mode
7. Automated desktop pool, leveraging VMware linked clones, in non-persistent access mode

Most businesses have a mix of user profiles (for example, finance, HR, engineering, help desk, call centers, teleworkers, data entry operators, and so on), and each use case might require one or more of these seven ways in which the virtual desktops can be provisioned and managed. Based on the requirements and mix of user profiles, companies might choose to implement one of these models or a specific mix. For example, a company requirement might dictate a need the following mix of desktop types:

- 10% individual desktops (for example, kiosks, part-time workers)
- 40% manual desktop pool in persistent access mode (for example, management staff, human resources, analysts, R&D)
- 20% automated desktop pool, utilizing VMware full clones in persistent access mode (for example, software developers)
- 30% automated desktop pool, utilizing VMware linked clones in non-persistent access mode (for example, help desk representatives)

Considering this to be a real-world scenario, the desktop mix might vary over time as more and more desktops are virtualized. However, the ultimate objective is to make sure that the complete end-to-end solution that is implemented can be easily managed, be cost effective, and provide the best-in-class end user experience.

The NetApp value proposition of storage efficiency, performance, operational agility, and data protection strongly complements all these seven ways in which virtual desktops can be deployed and managed, as shown in the table below.

Table 2) NetApp value propositions.

VMware View Manager			NetApp and VMware Solution				
Desktop Delivery Model	Virtual Desktop Provisioning Method	Access Mode	Storage Efficiency		Performance	Operational Agility	Data Protection
			OS Data	User Data			
Individual Desktop	NetApp RCU	Non-persistent	✓	✓	✓	✓	✓
Manual Desktop Pool	NetApp RCU	Persistent	✓	✓	✓	✓	✓
		Non-persistent	✓	✓	✓	✓	✓

VMware View Manager			NetApp and VMware Solution				
Desktop Delivery Model	Virtual Desktop Provisioning Method	Access Mode	Storage Efficiency		Performance	Operational Agility	Data Protection
			OS Data	User Data			
Automated Desktop Pool	VMware full clone	Persistent	✓	✓	✓	✓	✓
		Non-persistent	✓	✓	✓	✓	✓
	VMware linked clone	Persistent	✓	✓	✓	✓	✓
		Non-persistent	✓	✓	✓	✓	✓

Let us go into the details of each of the VMware View Manager desktop delivery models and access modes and validate how the NetApp solution strongly complements each one of them to achieve your ultimate goals and objectives.

6.1 INDIVIDUAL DESKTOP

If you are planning to implement individual desktops as part of the VMware View solution, the desired storage savings for both the VM operating system (OS) data (50% to 90%) and user data (up to 50%) can be achieved using the NetApp thin provisioning, deduplication, and FlexClone components of the solution. Note that the NetApp best practice is to store the user data and profile in a CIFS share on a NetApp NAS volume and served using [Microsoft® roaming profiles](#) and/or [folder redirection](#). This allows better management and protection of the important user data. NetApp intelligent caching and PAM enhance the end user experience. Space-efficient VMs can be rapidly created and imported into VMware View Manager using RCU, the NetApp plug-in in VMware vCenter Server. Effective solution monitoring, management, and data protection for both the OS data and user data can be achieved using VM Insight, Operations Manager, SMVI, SnapVault®, SnapMirror, and RAID-DP.

Table 3) Individual Desktops.

VMware View Manager			NetApp and VMware Solution				
Desktop Delivery Model	Virtual Desktop Provisioning Method	Access Mode	Storage Efficiency		Performance	Operational Agility	Data Protection
			OS Data	User Data			
Individual Desktop	NetApp RCU	Non-persistent	Thin provisioning FlexClone Deduplication	CIFS home directories Thin provisioning Deduplication	Intelligent caching PAM	RCU VM Insight Operations Manager SMVI	RAID-DP SMVI SnapVault SnapMirror

6.2 MANUAL DESKTOP POOL

If you are planning to implement manual desktop pool as part of the solution, you might choose to implement either the **persistent** or **non-persistent** access modes. For both of these types of access modes, the desired storage efficiency for the VM OS data (50% to 90%) and user data (up to 50%) can be achieved using the NetApp thin provisioning, deduplication, and FlexClone components of the solution. Note that the

NetApp best practice is to store the user data and profile in a CIFS share on a NetApp NAS volume and served using [Microsoft roaming profiles](#) and/or [folder redirection](#). This allows better management and protection of the important user data. NetApp intelligent caching and PAM help achieve the desired end user experience. Space-efficient VMs can be rapidly created and imported into VMware View Manager using RCU, the NetApp plug-in in VMware vCenter Server. Effective solution monitoring, management, and data protection can be achieved using VM Insight, Operations Manager, SMVI, SnapVault, SnapMirror, and RAID-DP.

Table 4) Manual Desktop Pool.

VMware View Manager			NetApp and VMware Solution				
Desktop Delivery Model	Virtual Desktop Provisioning Method	Access Mode	Storage Efficiency		Performance	Operational Agility	Data Protection
			OS Data	User Data			
Manual Desktop Pool	NetApp RCU	Persistent /Non-persistent	Thin provisioning FlexClone Deduplication	CIFS home directories Thin provisioning Deduplication	Intelligent caching PAM	RCU VM Insight Operations Manager SMVI	RAID-DP SMVI SnapVault SnapMirror

6.3 AUTOMATED DESKTOP POOL

If you are planning to implement automated desktop pools as part of the solution, you might choose to implement either a VMware full clone solution, a VMware linked clone solution, or a mix of the two.

VMWARE FULL CLONE

This solution supports use of both **persistent** and **non-persistent** access modes. With both of these types of access modes, the desired storage efficiency for the VM OS data (50% to 90%) and user data (up to 50%) can be achieved using NetApp thin provisioning and deduplication components of the solution. NetApp best practice is to store the user data and profile in a CIFS share on a NetApp NAS volume and served using [Microsoft roaming profiles](#) and/or [folder redirection](#). This allows better management and protection of the important user data. NetApp intelligent caching and PAM helps achieve the desired end user experience. Effective solution monitoring, management, and data protection can be achieved using VM Insight, Operations Manager, SMVI, SnapVault, SnapMirror, and RAID-DP.

Table 5) VMware full clone.

VMware View Manager			NetApp and VMware Solution				
Desktop Delivery Model	Virtual Desktop Provisioning Method	Access Mode	Storage Efficiency		Performance	Operational Agility	Data Protection
			OS Data	User Data			
Automated Desktop Pool	VMware full clone	Persistent/ Non-persistent	Thin provisioning Deduplication	CIFS home directories Thin provisioning Deduplication	Intelligent caching PAM	VM Insight Operations Manager SMVI	RAID-DP SMVI SnapVault SnapMirror

VMWARE LINKED CLONE

VMware linked clone solution also supports both the persistent and non-persistent access modes.

- For **persistent** access mode, the desired storage efficiency (50% to 90%) for the VM **“OS data disk”** can be achieved using VMware linked clones and NetApp thin provisioning components of the solution. The storage efficiency for the **“user data disk”** (up to 50%) can be achieved using NetApp deduplication and thin provisioning. Alternatively, depending on your data management and protection requirements for the important user data, you might choose to follow the same process of storing the user data and profile in a CIFS share on a NetApp NAS volume and served using [Microsoft roaming profiles](#) and/or [folder redirection](#).
- For **non-persistent** access mode, the desired storage efficiency (50% to 90%) for the VM **“OS data disk”** can be achieved using VMware linked clones and NetApp thin provisioning components of the solution. The storage efficiency (up to 50%) for the user data can be achieved using NetApp deduplication and thin provisioning. Note that the NetApp best practice is to store the user data and profile in a CIFS share on a NetApp NAS volume and served using [Microsoft roaming profiles](#) and/or [folder redirection](#). This allows better management and protection of the important user data.

For both persistent and non-persistent access modes, NetApp intelligent caching and PAM help enhance the end user experience. Effective solution monitoring, management and data protection can be achieved using VM Insight, Operations Manager, SMVI, SnapVault, SnapMirror, and RAID-DP.

Table 6) VMware linked clones.

VMware View Manager			NetApp and VMware Solution				
Desktop Delivery Model	Virtual Desktop Provisioning Method	Access Mode	Storage Efficiency		Performance	Operational Agility	Data Protection
			OS Data	User Data			
Automated Desktop Pool	VMware linked clone	Persistent /Non-persistent	VMware linked clones Thin provisioning Deduplication	CIFS home directories (optional) Thin provisioning Deduplication	Intelligent caching PAM	VM Insight Operations Manager SMVI	RAID-DP SMVI SnapVault SnapMirror

To summarize, NetApp solution strongly complements all the desktop delivery models and user access modes in VMware View Manager to provide a highly cost-effective, high-performing, operationally agile, and integrated VMware View solution.

7 ACCELERATING VMWARE VIEW WITH INTELLIGENT CACHING

7.1 CONCEPTS

VMware View by nature is a read-intensive technology. Some critical VMware View operational challenges experienced by most scalable deployments today are:

- Boot storms
- Login storms
- Virus scan storms

Boot storm is a somewhat rare but real condition when some or all virtual desktop machines start simultaneously. This can happen as a result of rolling out mass OS updates, new application installation, maintenance window, server failures, or any number of practical issues or interruptions. Login storms and

virus scan storms are also critical situations as they create spikes of activity for the entire solution. The traditional methods for addressing these critical challenges are:

- Increase cache for both ESX servers and storage devices
- Increase the spindle count
- Add more storage devices

The NetApp and VMware View solution addresses these challenges in a unique way, with NO negative tradeoffs to the customer environment. The NetApp intelligent caching in Data ONTAP, Performance Acceleration Module, FlexClone and deduplication are key components of the NetApp and VMware View solution that create an *intelligent deduplication-aware read cache* that helps customers:

- Dramatically reduce their storage footprint by more than 80% (validated in the NetApp solutions lab)
- Increase performance by up to 70% in the form of better network utilization (validated in the NetApp solutions lab)
- Decrease overall storage solution costs

NetApp intelligent caching eliminates the requirement for a large number of spindles to serve read-intensive VMware View workloads, while NetApp FlexClone and deduplication can further reduce the number of spindles required to store data.

7.2 BENEFITS OF INTELLIGENT CACHING

The following are some of the key benefits of using intelligent caching.

- **Increased performance:** With intelligent caching, in combination with FlexClone and deduplication, latencies decrease significantly by a factor of 10X versus serving data from the fastest spinning disks available, giving sub-millisecond data access (validated in the NetApp solutions lab). Decreasing the latency results in higher throughput and lower disk utilization, which directly translates into fewer disk reads.
- **Lowering TCO:** Requiring fewer disks and getting better performance allow customers to increase the number of virtual machines on a given storage platform, resulting in a lower total cost of ownership.
- **Green benefits:** Power and cooling costs are reduced as the overall energy needed to run and cool the PAM is significantly less than even a single shelf of Fibre Channel disks. A standard disk shelf of 300GB 15K RPM disks can consume as much as 340 Watts/hr and generate heat up to 1394BTU/hr. In contrast, the PAM card consumes only a mere 18 Watts/hr and only generates 90BTU/hr. By not deploying a single shelf, the power savings alone can be as much as 3000kWh/year per shelf (validated in the NetApp solutions lab). In addition to the environmental benefits of heating and cooling, each shelf not used saves 3U of rack space. For real-world deployment, a NetApp solution (with PAM as a key component) would typically replace several such storage shelves; therefore, the savings could be considerably higher than comparing to one disk shelf.





Figure 9) Power and heat savings for PAM vs. FC 15K disk shelf.

7.3 DEDUPLICATION & NON-DUPLICATION TECHNOLOGIES

Using NetApp Deduplication and File FlexClone not only can reduce the overall storage footprint of VMware View desktops but also can improve performance by leveraging intelligent caching. Data that is deduplicated or non-duplicated, in the case of File FlexClone data, on disk will only exist in intelligent cache once per volume. All subsequent reads from any of the VM disks (VMDKs) of a block that is already in cache will be read from cache and not from disk, therefore improving performance by 10x. Any non-deduplicated data that is not in cache must be read from disk. Data that is deduplicated but does not have as many block references as a heavily deduped VMDK will appear in cache only once but based on the frequency of access might be evicted earlier than data that has many references or is heavily used.



Deduplication within VMware Environments

VMware View deployments can reduce storage footprint by up to 99%.

This diagram demonstrates the initial deployment where all blocks are duplicate blocks.

Figure 10) NetApp deduplication in VMware environments.

Deduplication Guidelines

- Deduplication is configured and operates on the flexible volumes only.
- Data can be deduplicated up to 255:1 without consuming additional space.
- Each storage platform has different deduplication limits.
- Each volume has a dense and non-dense size limit.
- Deduplication is configured using the command line.
- Requires Data ONTAP 7.2.5.1, 7.3P1, or later.
- Both a_sis and NearStore® must be licensed for deduplication to work.
- Run deduplication prior to creating Snapshot copies or running SnapMirror or SnapVault updates.

For more detailed information on deduplication, refer to [NetApp TR-3505: NetApp Deduplication for FAS, Deployment and Implementation Guide](#).

7.4 INTELLIGENT CACHING DETAILS

Intelligent caching is performed natively within Data ONTAP and can be extended with the use of a Performance Acceleration Module. PAM has both a hardware component (the PAM module) and software component (FlexScale). This section will describe these components and the NetApp best practices to use them in a VMware View environment.

7.4.1 Intelligent Caching in Data ONTAP

Data ONTAP will only store a single block in cache for up to 255 physical blocks. Data ONTAP intelligent caching is available in all versions of Data ONTAP 7.3.1 or higher. This means that intelligent caching can be used in every FAS, V-Series, & IBM N Series that support Data ONTAP 7.3.1 and block sharing technologies.

How Data ONTAP Intelligent Caching Functions

When a data block is requested, Data ONTAP reads the block into main memory (also known as the WAFL® buffer cache). If that data block is a deduplicated block, in that it has multiple files referencing the same physical block, each subsequent read of that same physical block will come from cache as long as it has not been evicted from cache. Heavily referenced blocks that are frequently read will reside in cache longer than blocks that have fewer references or less frequent access. The effect this has is that since main memory can be accessed much faster than disk, latency is decreased, disk utilization is decreased, and network throughput is increased thus improving overall performance and end user experience without any negative tradeoffs.




Figure 11) Intelligent caching function with dedupe.

How Data ONTAP Intelligent Caching Functions with PAM

NetApp has the ability to extend Data ONTAP's native Intelligent Caching with the use of the Performance Acceleration Module. As long as that block is not evicted from both caches, all subsequent reads will be performed from main memory or PAM, thus improving performance by not having to go to disk. Again the more heavily the data is deduplicated and more frequently accessed the longer it will stay in cache. Intelligent caching combined with NetApp disk deduplication provides cost savings on many levels without any negative tradeoffs.




Figure 12) Intelligent caching with PAM and deduplication.

The decision on whether to use PAM in addition to Data ONTAP intelligent caching is based on the amount of deduplicated data and the percentage of reads within the environment. As users of the VMware View environment create more data, the amount of deduplicated data will change, thus affecting the cache hit rate. Thus, more cache might be needed if the data becomes more unique (even after running regular deduplication operation on the new data).

NetApp recommends when possible to always use Data ONTAP 7.3.1 or later for VMware View environments. For environments with greater than 500 virtual desktops per NetApp storage controller, NetApp recommends the use of both Data ONTAP caching and at least one PAM module per storage controller.

7.4.2 The PAM Module

The PAM is a PCI Express card that can be installed in FAS3000, V3000, FAS6000, and V6000 series platforms running Data ONTAP 7.3 or higher. Each module contains 16GB of DRAM. The VMware View solution can be composed of one (16GB) to 10 (160GB) PAM cards per FAS or V-Series storage cluster depending on the customer requirements. Details on the number of modules per platform can be found at www.netapp.com/us/products/storage-systems/performance-acceleration-module/performance-acceleration-tech-specs.html.

How the PAM Functions Without Deduplication

The PAM works by receiving data blocks that have been evicted from main memory. After evicted from main memory, should the same block be required for a second time and that block has not been evicted from PAM, that block is read from PAM and placed into main memory. Every block, whether or not it contains the same data as another block, is read first from disk. This is how legacy storage arrays operate in that the first of all reads must come from disk, and subsequent reads depend on the size of the cache. This is the reason legacy vendors require large amounts of cache.




Figure 13) NetApp PAM without deduplication.

How the PAM Functions with Deduplication

PAM receives data blocks that have been evicted from main memory. After eviction from main memory, should a block be required for a second time, that block is read from PAM, a cache hit, and placed into main memory. If the block being requested is a duplicate block that has been deduplicated (also known as a shared block), the block is read from PAM to main memory. As long as that block is not evicted from cache, all subsequent reads will be performed from PAM, thus improving performance by not having to go to disk. Intelligent deduplicated caching combined with NetApp disk deduplication provides cost savings on many levels without any negative tradeoffs.




Figure 14) NetApp PAM with deduplication.

7.4.3 FlexScale

FlexScale is the tunable software component to the PAM. It is a licensed feature of Data ONTAP 7.3 or greater. FlexScale allows different caching modes to be used based on the type of workload. The different modes of caching are metadata only, normal user data, and low-priority blocks. Extensive scalable VMware View testing within the NetApp solution labs has shown that significant performance improvements can be gained by turning on all three caching modes in FlexScale. Below are the instructions for licensing and enabling FlexScale.

Table 7) Installing the FlexScale license.

Step	Action
1	Connect to the controller system's console, using either SSH, telnet, or serial console.
2	Check to see if the FlexScale license has already been installed by typing license and finding line that says flex_scale: <i>license</i>
3	If FlexScale is not licensed, you can license it by issuing the following command. If you do not have your license available, you can locate it within the NOW™ (NetApp on the Web) site: http://now.netapp.com . <i>license add <License_Key></i>

Table 8) Changing the FlexScale caching modes for use with VMware View workloads.

Step	Action
1	Connect to the controller system's console, using either SSH, telnet, or serial console.

2	Change the following options with the following commands. This turns on metadata, normal user data, and low-priority blocks caching. <pre>options flexscale.enable on options flexscale.normal_data_blocks on options flexscale.lopri_blocks on</pre>
3	You can verify these settings have been changed: <pre>options flexscale</pre>

7.4.4 Predictive Cache Statistics (PCS)

PCS is a method of simulating cache to determine if one or more Performance Acceleration Modules would be necessary for the given VMware View workload. If you have not already purchased a PAM card, you can engage a NetApp or partner systems engineer to assist with collecting predictive cache statistics. The PCS counters are available in Data ONTAP 7.3 or higher.

7.5 SUMMARY OF INTELLIGENT CACHING IN A VMWARE VIEW ENVIRONMENT

NetApp has seen significant performance improvement with intelligent caching during VMware View boot storm benchmarking at customer sites and internal NetApp solutions lab. In the NetApp solutions lab, 4,096 Windows XP VMs were simultaneously booted from a single FAS3070 active-active midtier storage controller. During these tests detailed controller statistics and ESX statistics were collected. The tests were performed with and without PAM and with and without deduplication. The test results demonstrated that with the PAM in the VMware View solution, we can provide:

- Up to 70% increase the network throughput
- Up to 67% decrease in the number of reads from disk
- Up to 40% improvement in the boot time

To summarize, in our labs we were able to boot 4,096 VMs using PAM and deduplication in under 30 minutes. Intelligent caching with Data ONTAP or PAM and deduplication provide significant value in a VMware View environment with no negative tradeoffs.

The Intelligent Caching Value

- **Cache efficiency:** Deduplication occurs not only on disk but also in cache. Working sets of data are deduplicated so larger caches are not needed as in traditional storage solutions.
- **Performance acceleration:** Blocks read from cache are served 10 times faster, as latency is reduced by a factor of 10.
- **Storage efficiency:** Allows you to reduce the spindle count even further as large percentage of the I/O requests are served up directly from intelligent cache.
- **Lower TCO:** NetApp intelligent caching and deduplication reduce rack space, power, and cooling.

NetApp Recommendations

Since read IO can be greatly reduced by using intelligent caching NetApp recommends the use of Data ONTAP 7.3.1 or later. This version supports intelligent caching across our entire truly Unified Storage product line. For Mid & High Tier systems the use of Intelligent Caching with both Data ONTAP and extension of cache using the PAM card provides extended capabilities for larger environments.

The net result of intelligent caching is two fold, customers will buy less storage to serve IO's from intelligent cache and they will go faster by utilizing deduplication and NetApp's Intelligent Caching in VMware View environments. The end result will be enhanced user experience

8 STORAGE SIZING BEST PRACTICES

Storage estimation for deploying VMware View solutions on NetApp includes the following steps:

- Gather essential solution requirements
- Perform performance-based and capacity-based storage estimation
- Getting recommendations on storage system physical and logical configuration

8.1 GATHER ESSENTIAL SOLUTION REQUIREMENTS

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

- Total number of VMs for which the system has to be designed (for example, 2000 Windows XP VMs).
- The types and percentage of different types of desktops being deployed. For example if VMware View Manager is used, different desktop delivery models might require special storage considerations.
- Size per VM (for example, 20GB C: drive, 2GB data disk).
- VM OS (for example, Windows XP, Vista, and so on).
- Worker workload profile (type of applications on the VM, IOPS requirement, if known).
- Number of years for which the storage growth has to be considered.
- Disaster recovery/business continuance requirements.
- Backup requirements (frequency and retention policies) for C: drives.
- Size of NAS (CIFS) home directories.
 - NetApp strongly recommends storing user data on NAS (CIFS) home drives. Using NAS home drives, companies can more efficiently manage and protect the user data and eliminate the need to back up the virtual desktops.
- For most of the VMware View deployments, companies might also plan to implement roaming profiles and/or folder redirection. For detailed information on implementing these technologies, consult the following documentation:
 - Microsoft [Configuring Roaming User Profiles](#)
 - NetApp [TR-3367: NetApp Systems in a Microsoft Windows Environment](#)
 - Microsoft [Configuring Folder Redirection](#)
- **VMware View Manager considerations:** When implementing VMware View Manager, decide on the following:
 - Types of desktops that will be deployed for different user profiles
 - Data protection requirements for different data components (OS disk, user data disk, CIFS home directories) for each desktop type being implemented
 - For automated desktop pools utilizing linked clone in persistent access mode, the user data and profile can be hosted on a separate “user data disk.” Being a vmdk file, it is important to decide on the user data disk size upfront. NetApp thin provisioning, deduplication, and SMVI data protection solution components can be leveraged to achieve the desired storage efficiency and data protection for the “user data disk”.

8.1.1 Performance Requirements

ESTIMATING ENVIRONMENT WORKLOAD

For proper storage sizing, it is critical to determine the IOPS requirement per virtual desktop. This involves analyzing how busy the virtual desktops will be and the percentage of users who are heavy workers (knowledge workers) versus light workers (for example, data entry workers). Important factors to be considered are:

- Hourly, daily, monthly, and quarterly user workload (best case and worst case scenarios)
- Percent reads versus writes
- Commonality of data and how well the data will deduplicate as this is directly related to cache efficiency
- Impact of antivirus operation (scanning and virus definition updates) requirements (frequency, schedules, and so on). For antivirus operations, one could either take a performance hit during scheduled AV operations and impact the end user experience or appropriately design the VMware View solution to make the AV operations seamless. The first option is definitely not desirable. The second option is highly recommended. Long-term solution would be to implement a security technology such as [VMware VMsafe™](#) as and when it becomes available to help offload the virus scanning processing from the guest.
- Any recommendations specific to VMware on storage performance and IOPS requirements for best-case and worst-case situations for the customer environment. Also, VMware has provided some guidelines on IOPS per heavy and light user in the [“VMware VDI Storage Considerations” white paper](#).

PERFORMANCE DATA COLLECTION METHODS

This performance data can be collected in many ways. If the VMware View environment is not new, one of the following methods could be used:

- **NetApp data collection tool:** NetApp storage data collector and analyzer tool collect storage-specific performance counters from a range of Windows clients and help analyze the collected data so that they can be effectively used with the NetApp storage sizing tools. For details on how to obtain and run the tool in your environment, contact your NetApp account team.
- **VMware Capacity Planner data collector, Windows Logman tool, PlateSpin, and TekTools**

In addition to the NetApp data collection tool, VMware Capacity Planner, Windows Logman, Perfmon, PlateSpin, or TekTools can be used to analyze the existing physical desktops within the environment to understand the I/O requirements.

Any of these methods can produce data that will assist in sizing the storage platform and the spindle count required to service the workload.

Example formula:

Total storage IOPS requirement = (sum of all max IOPS/number of desktops tested) x number of virtual desktops

Example:

During the performance data collection over 30 days, the max number of IOPS for 10 clients totaled 1,327. The customer is looking to deploy 100 seats. From this sample one can calculate the estimated IOPS requirement. From this number as it is a maximum, one would have to decide the probability of all the clients reaching the max I/O requirement at the same time and adjusting an acceptable maximum, as architecting a solution for this concurrency might not be necessary. The average of all IOPS on all clients can be used in to get a better understanding of the daily load:

13,270 total max IOPS = (1,327 sum of max IOPS/10 desktops tested) x 100 future virtual desktops

8.2 PERFORMANCE-BASED AND CAPACITY-BASED STORAGE ESTIMATION PROCESS

There are two important considerations for sizing storage for VMware View. The storage system should be able to meet both the performance and capacity requirements of the project and be scalable to account for future growth. The steps for calculating these storage requirements are:

1. Determine storage sizing building block
2. Perform detailed performance estimation
3. Perform detailed capacity estimation
4. Obtain recommendations on the storage system physical and logical configuration

8.2.1 Determine Storage Sizing Building Block

The purpose of this step involves determining the logical storage building block or POD size. This means deciding on the following parameters:

- Storage building block scope: NetApp recommends the VMware View storage sizing building block to be based on the number of datastores required per ESX cluster because it provides benefits of planning and scaling the storage linearly with the number of ESX clusters required for the solution.
- Determine the usable storage required per storage building block (per ESX cluster). For VMFS datastores, there can be multiple LUNs per flexible volume where each LUN is a datastore. For NFS datastores, each volume can represent a datastore with more number of VMs as compared to VMFS datastores.
- Flexible volume layout across storage controllers: All the flexible volumes belonging to an ESX cluster should be evenly split across the two controllers of the hosting NetApp storage cluster. This is recommended for better performance as the VMware View deployment scales out from one HA cluster to multiple, hundreds to thousands to tens of thousands of virtual desktops.

CONSIDER VMWARE ESX CONFIGURATION MAXIMUMS

Carefully review the VMware documentation on configuration maximums associated with the various storage-related parameters critical to the system design. For ESX 3.5, review the document at www.vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_config_max.pdf.

The important configuration parameters critical to the design are:

- Number of virtual CPUs per server: This information is important to understand the maximum limit on the number of VMs that can be hosted on the physical server, irrespective of the number of cores per server.
- Number of virtual CPUs per core for VMware View workloads: This information will determine the upper limit on the numbers of VMs that can be supported per physical ESX host, but it cannot be more than the limit on number of virtual CPUs that can be hosted per server. Consult your VMware SE for their recommendation on number of VMs that can be supported per ESX Server host.
- Number of VMs managed per vCenter Server instance: This information will help determine the maximum number of ESX hosts that can be managed by a single vCenter Server instance.
- Number of NAS datastores per cluster: This information is critical for sizing scalable virtual desktops on NFS datastores.
- Number of VMFS datastores configured per server (for FC/iSCSI): This information is critical for sizing scalable VMware View solutions on VMFS datastores.
- Number of VMs per VMFS datastore: This information is critical for sizing scalable VMware View solutions on VMFS datastores. For NFS, there are no VMware recommendations on the maximum number of VMs per datastore.
- Number of hosts per HA/DRS cluster

The above information should help determine the following design parameters:

- Proposed number of VMs per ESX host

- Proposed number of ESX host per ESX cluster
- Proposed number of datastores per ESX cluster
- Proposed number of VMs per ESX cluster
- Number of ESX clusters managed by a vCenter instance
- Proposed number of VMs per datastore
- Total number of datastores required for the project

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

DECIDE ON STORAGE PROTOCOL

The two shared storage options available for VMware View are:

- VMFS-based datastores over FC or iSCSI
- NFS-based datastores

NetApp is a true unified multiprotocol storage system that has the capability to serve storage for both the options listed above from a single storage cluster without the use of additional SAN or NAS gateway devices.

Both of these are viable and scalable options for VMware View. Consider reading [NetApp TR-3697: Performance Report: Multiprotocol Performance Test of VMware ESX 3.5 on NetApp Storage Systems](#) for results from a technical performance study conducted jointly by NetApp and VMware on different storage protocols. Also, perform a cost benefit analysis for your environment and decide on the storage protocol to be used. The key NetApp value proposition for VMware View holds true across all the protocols.

8.2.2 Perform Detailed Performance Estimation

This step involves estimating the total number of disk IOPS and PAM cards required for the VMware View solution based on the requirements. Intelligent caching capabilities available in NetApp solution help significantly reduce the amount of data disks required. The calculations are performed based on the IOPS requirement per VM, version of Data ONTAP, intelligent caching, and the customer environment workload characteristics. For getting detailed estimation on savings achieved for your environment, contact your NetApp account team. The output of this step includes:

- Total number of PAM cards required
- Total IOPS required in order to meet the performance needs
- Percentage of the total IOPS that require data disks considering the NetApp intelligent caching and PAM capabilities.

VMware View Manager Considerations: This step is applicable to all the seven virtual desktop types available in VMware View Manager 3.0 and can help reduce the total number of spindles required.

8.2.3 Perform Detailed Capacity Estimation

The flowchart below describes the various steps involved in the capacity-based storage estimation process. Each of these steps is discussed in detail below.




Figure 15) Overview of capacity estimation process.

DETERMINE STORAGE REQUIRED FOR VM FILES

There are several files associated with each VM, and these files require shared storage in addition to the actual VM.

Table 9) VMware file listing.

Files	Purpose	Storage Required
.vmdk	Two VMDK files that make up a VM. The –flat.vmdk file is the actual data disk, and the .vmdk file is the descriptor file that describes the disk geometry of the –flat file, (<2K). If considering using VMware Snapshot copies, the –delta.vmdk files must be considered as a part of the storage requirements.	Size of the VM (for example, 20GB)
.vswp	Each VM will have its own VM-specific .vswp file based on the amount of memory assigned to that VM. For example if you are sizing for 100 VMs, each with 512MB of RAM, then you need to plan for 50GB storage for .vswp. <i>NetApp recommendation is to move the vswp file to a different datastore if Snapshot copies are being considered. See NetApp TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices for more details on this recommendation.</i>	Amount of memory assigned to each VM (for example, 512MB/VM)
.vmsd/.vmsn	A VMSN file stores the exact state of the VM when the Snapshot copy was created. A VMSD file stores information and metadata about Snapshot copies. It is important that space for these files be accounted, if temporary Snapshot copies have to be created. Consider the number of Snapshot copies, how long they are likely to be kept, and how much they would grow. If the Snapshot copies also include memory, then the space requirements can grow very quickly.	The total space required for the non-VMDK files might be 5% if the vswp is moved to a separate datastore. It will be higher (15%) if the vswp is located on the same datastore. This number will be
.vmx	A vmx file is the primary configuration file for a VM and stores the important configuration parameters.	

Files	Purpose	Storage Required
.vmxf	This is a supplemental configuration file in text format for VMs that are in a team. Note that the .vmxf file remains if a VM is removed from the team.	even more if there are specific requirements for creating and retaining VMware Snapshot copies.
.vmss	This is the suspended state file, which stores the state of a suspended VM.	
.nvram	This is the file that stores the state of the VM BIOS.	
.log	The log files contain information on the VM activity and are useful in troubleshooting VM problems.	
Total storage for VM files		Size of VM + 5–15% space for non-VMDK files

VMware View Manager Considerations:

- If you are planning to implement VMware View Manager, from calculation purpose, it is important to understand the total space required per VM for individual desktops, manual desktop pool, and automated desktop pool leveraging full clones. The actual storage required (output of the storage estimation process) will be far less considering the NetApp solution space efficiency components.
- For “automated desktop pool” leveraging linked clones, estimate the space required by all the files that make up the parent VM, replica VMs, OS data disks, user data disks, and vswap, considering the policies for linked clone desktop refresh, recompose, and rebalance. For further details, refer to the [VMware View Manager System Administrator Guide](#).

DETERMINE PROJECTED STORAGE ESTIMATION LIFETIME

Determine the total number of years for which the storage growth has to be factored in. This is important because when using the NetApp FlexClone and deduplication solution components, initially the VMs hosted in the FlexClone volumes will not consume any space. But the new writes will require storage to commit the writes.

PROVISIONING WITH FLEXCLONE

FlexClone provides simple, rapid, space-efficient, zero upfront cost provisioning of VMware View datastores using NetApp RCU. FlexClone technology uses the power of NetApp Snapshot technology to rapidly provision VMFS or NFS datastores. The process of provisioning with FlexClone involves creating a gold datastore and making multiple FlexClone of the datastore depending on the total number of VMs required. This gold datastore is the base of the FlexClone volumes. For example, if you have an NFS volume with 200 VMs contained within the volume/datastore, leverage FlexClone to provision additional volumes of 200 VMs within seconds. Using the power of FlexClone you can eliminate the need to clone each VM sequentially, which can be a very time-consuming process, especially when provisioning hundreds or thousands of VMs. For more information on FlexClone technology concepts, see [NetApp TR-3347: FlexClone Volumes: A Thorough Introduction](#).

VMware View Manager Considerations: If you are planning to implement VMware View Manager, the use of NetApp FlexClone technology is applicable to the “individual desktops” and “manual desktop pool” desktop delivery models. For “automated desktop pool” leveraging full clones, estimate the space required by each full clone. Automated desktop pools can also benefit from the use of VMware linked clones.

ESTIMATE GOLD VOLUMES/DATASTORES

The next step is to size the gold datastore/volumes appropriately. The gold datastore is a datastore that will be used as a base that will be cloned quickly and efficiently using FlexClone. Using gold datastores will provide the ability to utilize the provisioning benefits associated with FlexClone. From an overall management and scalability perspective, consider splitting the datastores for each ESX cluster equally across the two controllers of the NetApp storage system. This would require two gold datastores (one per

NetApp controller) per ESX cluster. It also helps multiple ESX clusters have no dependencies on one another for sharing storage. When using NetApp RCU with Data ONTAP 7.3.1 and higher, file-level FlexClone can be leveraged to rapidly make space-efficient clones inside a gold volume.

This example demonstrates the sizing of a gold volume for an NFS environment in which we are creating one datastore with 200 VMs. We have budgeted 10% overhead within the gold volume.

Example formula 1 (Data ONTAP 7.3.1 or higher):

Size of gold volume = [(space required per VM) x (number of datastores per flexible volume)] + 10% buffer space

Example:

28GB gold volume (approx.) = [(25GB per VM) x (1 datastore per volume)] + 10%

Example formula 2 (Data ONTAP 7.3 or earlier):

Size of gold volume = [(space required per VM) x (number of VMs per datastore) x (number of datastores per flexible volume)] + 10% buffer space

Example:

5.5TB gold volume = [(25GB per VM) x (200 VMs per datastore) x (1 datastore per volume)] + 10%

Note: The buffer space in formula 2 is only a temporary requirement so that the gold volume doesn't get 100% full when all the VMs are added in the datastore. Set the Snapshot reserve to zero and disable the default Snapshot schedule.

CONSIDER GOLD VOLUME DEDUPLICATION SAVINGS

Once the storage required per gold volume is determined, estimate the deduplication savings. For flexible volumes with multiple identical VMs, the deduplication savings will be very high as the OS part of all the VMs in the volume is almost identical.

Note: For deployments leveraging Data ONTAP 7.3.1 or higher, the gold datastore is already deduplicated as a result of FlexClone file-level cloning inside the gold datastore. Therefore, factoring in deduplication savings is valid only for deployments with Data ONTAP 7.3 or lower.

For VMware View on NetApp environments, we estimate the initial deduplication savings to be at least 80% to 95% for the gold datastores as all virtual machines within the datastore are identical, given that there are more than eight VMs per gold datastore.

For more general information on deduplication, refer to [NetApp TR-3505: NetApp Deduplication for FAS, Deployment and Implementation Guide](#).

The output of this step will give the storage required per gold flexible volume after the deduplication savings have been factored in (for Data ONTAP 7.3 or lower). It might be also necessary to factor in additional space for Snapshot reserve. The Snapshot reserve can be less for environments where the gold volume will not change much and few Snapshot copies will be retained. Consult a NetApp SE for recommendations for your data protection requirements.

Example formula:

Deduplicated gold volume size = [(size of gold FlexVol volume) x (% deduplication)] x % Snapshot reserve space

Example:

121GB gold volume size = [(5.5TB FlexVol volume) x (98% savings)] x 10%

VMware View Manager Considerations: If you are planning to implement VMware View Manager, the use of NetApp deduplication for gold datastore is applicable to the "individual desktops" and "manual desktop pool" desktop delivery models. Deduplication scenarios for automated desktop pools leveraging linked clones are discussed later in the document.

ESTIMATE ADDITIONAL STORAGE CONSUMED PER VM

All the VMs hosted on the flexible volumes provisioned with FlexClone will initially access the same data blocks as the deduplicated gold volume. However, the new writes for every VM will consume new storage blocks. Over the years for which the storage growth has to be considered, every FlexClone volume will have consumed some new storage space.

Calculate the projected storage consumed by each VM over this time period, across all the VM storage elements (OS/system files, data files on C drive) and non-VMDK files in the datastore:

1. Projected new storage consumed per year (GB) for OS part of the VM.
2. Projected new storage consumed per year (GB) for data part of the VM.
3. Projected new storage consumed per VM for non-VMDK files in the datastore.
4. Multiply each of these elements above with the total number of years for which the storage estimation is being done.

Example formula:

Total additional storage/VM for OS part of the C drive = (yearly OS delta) x number years

Total additional storage/VM for data part of C drive = (yearly data delta) x number years

Total additional storage/VM for non-VMDK files in the datastore = (yearly non-VMDK files delta) x number years

Example:

6GB total additional storage/VM for OS = (2GB OS) x 3 years

15GB total additional storage/VM for data = (5GB data) x 3 years

0.9GB total additional storage/VM for data = (300MB data) x 3 years

VMware View Manager Considerations: If you are planning to implement VMware View Manager, considerations should be given to each desktop delivery model.

- The method described above is valid for VMs being used in the “individual desktops” and “manual desktop pool” desktop delivery models.
- For automated desktop pools, leveraging linked clone, considerations should be given to the linked clone desktop disk usage.
 - For the VMs provisioned using linked clones in persistent access mode, based on your user profiles, determine how much will be growth of the OS data disks between the refresh, recompose, and/or rebalance operations.
 - Also, for the VMs provisioned using linked clones in persistent access mode, determine the projected growth of the “user data disk” over the lifespan of the user accessing this disk. It is important to decide on the size of the vmdk file representing the user data disk upfront.
 - For VMs provisioned using linked clones in non-persistent access mode, based on your user profile, determine how much will be the OS data disk growth.
- For further details, refer to the [VMware View Manager System Administrator Guide](#).

ESTIMATE STORAGE REQUIRED FOR FLEXCLONE VOLUMES (CONSIDERING DEDUPE SAVINGS)

The storage required per FlexClone volume can be obtained by multiplying the results of the previous step with the total number of VMs planned per FlexClone volume and discounting the deduplication savings. This will also be a factor of number of VMs hosted per FlexClone volume. The deduplication savings for the OS image part of the C: drive will be higher because all the patches and upgrades applied to each VM will essentially be the same. A conservative number for the OS part would be between 50% and 70% (based on existing customer deployments and NetApp solutions lab validation). However, the deduplication savings for the data part of the C: drive might not be as dense as the OS images as it is unique to each individual VM. A

conservative number for the data part would be between 20% and 50%, as is seen in CIFS home directory deduplication savings for customer deployments.

Example formula:

Per FlexVol volume storage consumption = [(Total additional storage per VM for OS) x (number of VMs per FlexClone volume) x deduplication savings] + [(total additional storage per VM for data) x (number of VMs per FlexClone volume) x deduplication savings] + [(total additional storage/VM for non-VMDK files in the datastore) x (number of VMs per FlexClone volume) x deduplication savings]

Example:

1950GB per FlexVol volume storage consumption = (6GB x 200 VMs) x 70% savings + (15GB x 200 VMs) x 50% savings + (0.9GB x 200 VMs) x 50% savings

For environments in which Snapshot and/or mirroring will not be used for the VM C drives, NetApp recommends setting the Snapshot reserve to 0% and disabling the Snapshot schedule. If Snapshot copies or SnapMirror will be used, the snap reserve should be set to a value that will allow for the planned number of Snapshot copies.

Table 10) Adjusting Snapshot reserve and Snapshot schedule.

Step	Action
1	Connect to the controller system's console, using either SSH, telnet, or serial console.
2	Set the volume Snapshot reserve: <i>snap reserve <vol-name> ##</i>
3	Set the volume Snapshot schedule: <i>snap sched <vol-name> 0 0 0</i>

The output of this step is the total usable storage required by each FlexClone volume. This number can be extrapolated to calculate the total usable storage required for the gold volumes and associated FlexClone volumes per ESX cluster and ultimately for the entire environment depending on the total number of ESX clusters required. Again, NetApp recommends splitting the datastores for each ESX cluster across the two controllers of the NetApp storage system.

View Manager Considerations:

- The considerations in this step are valid for VMs provisioned using NetApp RCU (leveraging FlexClone technology) configured either as individual desktops and part of manual desktop pool.
- For VMs provisioned as part of automated desktop pools using VMware full clones, 50% to 90% storage savings can be achieved using NetApp deduplication. Consider these savings in the storage estimation process.
- For VMs provisioned as part of automated desktop pools using VMware linked clones, significant storage savings can be achieved for the "user data disk" using NetApp deduplication. A conservative number to consider would be between 20% and 50%, as seen for home directories.

FACTOR IN SCRATCH SPACE STORAGE REQUIREMENTS

If required, factor in additional storage (scratch space) for test and dev operations or any other reason. This step is not mandatory, but highly recommended (for future or last-minute design changes).

SUMMARY OF CAPACITY-BASED STORAGE ESTIMATION PROCESS

The capacity calculations have provided guidance to the following essential storage architecture elements:

- Total number of datastores per gold volume
- Total number of FlexClone volumes per gold volume
- Total storage required per gold volume

- Total storage required per FlexClone volume
- Total storage required for each gold and FlexClone volume combination
- Total number of gold and FlexClone volume combinations
- Total storage required for all the gold and FlexClone volume combinations
- All the storage considerations for different desktop delivery models available in VMware View Manager

SUMMARY OF STORAGE CONSIDERATIONS FOR DIFFERENT DESKTOP DELIVERY MODELS IN VMWARE VIEW MANAGER

The table below summarizes the storage sizing considerations for different desktop deployment models in VMware View Manager, specifically with linked clones.

Table 11) Summary of storage considerations for desktop delivery models.

Desktop Delivery Model	Virtual Desktop Provisioning Method	Generic Recommendations	Special Storage Considerations
Individual Desktop	NetApp RCU (FlexClone)	• Determine the types of desktops that will be deployed for different user profiles • Determine data protection requirements for different data components (OS disk, user data disk, CIFS home directories) for each desktop type	Primary objective of the discussion in this section of the document
Manual Desktop Pool	NetApp RCU (FlexClone)	• Consider reading the VMware View Manager System Administrator Guide	
Automated Desktop Pool	VMware full clone		• Estimate the space required by each full clone • Leverage NetApp deduplication to achieve 50% to 90% storage efficiency • Consider the storage savings in the storage estimation process
	VMware linked clone		• Estimate the space required by all the files that make up the parent VM, replica, OS data disks, user data disks, and vs wap considering the policies for linked clone desktop refresh, recompose, and rebalance.

Desktop Delivery Model	Virtual Desktop Provisioning Method	Generic Recommendations	Special Storage Considerations
			• Consider space in each datastore for different replica(s). The number of replica(s) in a datastore will depend on the total number of parent VMs and Snapshot copies to which the active linked clone VMs in each datastore are anchored. • Decide on your storage overcommit policies to determine the number of linked clone VMs that can be hosted in a datastore. For detailed information, refer to the VMware View Manager System Administration Guide. • Give consideration to the linked clone desktop disk usage as in some instances, the OS data disk can grow to the size of the parent VM (for details, see the VMware View Manager System Administration Guide). • For the persistent access mode, based on your user profiles, determine how much will be the growth of the OS data disks between the refresh, recompose, and/or rebalance operations. • For the persistent access mode, determine the projected growth of the user data disk over the lifespan of the user accessing this disk. This is important to decide in the design phase of the project. • For the persistent access mode, significant storage savings can be achieved for the user data disk using NetApp thin provisioning and data deduplication. A conservative number to consider would be between 20% and 50%, as seen for home directories. • For the persistent access mode, NetApp SMVI data protection solution can be leveraged to achieve the desired storage efficiency and data protection for the "user data disk". • For non-persistent access mode, based on your user profile, determine the growth rate of the OS data disk. Consider policies such as delete VM after first use to conserve storage space.




Figure 16) FlexClone scalability.

8.3 Getting Recommendations on Storage System Physical and Logical Configuration

Provide the total capacity and performance requirements to the NetApp SE and obtain appropriate storage system configuration. If required, NetApp can help you in each phase of the process discussed above. NetApp has detailed VMware View specific sizing tools that can help architect VMware View deployments of any scale. The tools are designed to factor in all the NetApp storage efficiency and performance acceleration components discussed above.

This step also involves planning the logical architecture (the total number of gold and the associated FlexClone volumes that should be provisioned per aggregate). The recommendation is to provision fewer large aggregates over more, smaller aggregates. The advantages to larger aggregates are that the I/O has more disks to write across, therefore increasing the performance of all volumes contained within the aggregate. Based on the estimated volume size from capacity calculations section above, determine the number of gold and associated FlexClone volumes that can be hosted in the largest possible aggregate. It is also a good idea to leave some room to grow the aggregates to handle situations when unexpected growth occurs. Also, disable scheduled aggregate Snapshot copies and set the aggregate snap reserve to 0. Make sure the data disk in the aggregate satisfies the performance requirements for the proposed number of VMs for volumes to be hosted in the aggregate.

9 STORAGE ARCHITECTURE BEST PRACTICES

In a VMware View environment, the availability and performance of the storage infrastructure are very critical because thousands of users will be impacted by storage outages or performance issues. Thus the storage architecture must provide the level of availability and performance typical for business-critical applications. NetApp has all the software and hardware solutions that address the availability and performance for large, scalable VMware View environments.

9.1 STORAGE SYSTEM CONFIGURATION BEST PRACTICES

This section of the solution guide provides a high-level overview of the components and features that should be considered when deploying a VMware View infrastructure on NetApp. For detailed information on storage resiliency, refer to the following TRs:

- [NetApp TR-3437: Storage Best Practices and Resiliency Guide](#)
- [NetApp TR-3450: Active-Active Controller Overview and Best Practices Guidelines](#)

9.1.1 Building a Resilient Storage Architecture

ACTIVE-ACTIVE NETAPP CONTROLLERS

The controller in a storage system can be a single point of failure if not designed correctly. Active-active controllers provide controller redundancy and simple automatic transparent failover in the event of a controller failure to deliver enterprise-class availability. Providing transparent recovery from component failure is critical as all desktops rely on the shared storage. For more details, visit www.netapp.com/us/products/platform-os/active-active.html.

MULTIPATH HIGH AVAILABILITY (HA)

Multipath HA storage configuration further enhances the resiliency and performance of active-active controller configurations. Multipath HA-configured storage enhances storage resiliency by reducing unnecessary takeover by a partner node due to a storage fault, improving overall system availability and promoting higher performance consistency. Multipath HA provides added protection against various storage faults, including HBA or port failure, controller-to-shelf cable failure, shelf module failure, dual intershell cable failure, and secondary path failure. Multipath HA helps provide consistent performance in active-active configurations by providing larger aggregate storage loop bandwidth. For more information, visit <http://media.netapp.com/documents/tr-3437.pdf>.

RAID DATA PROTECTION

Data protection against disk drive failure using RAID is a standard feature of most shared storage devices, but with the capacity and subsequent rebuild times of current hard drives where exposure to another drive failure can be catastrophic, protection against double disk failure is now essential. NetApp RAID-DP is an advanced RAID technology that is provided as the default RAID level on all FAS systems. RAID-DP provides performance that is comparable to that of RAID 10, with much higher resiliency. It provides protection against double disk failure as compared to RAID 5, which can only protect against one disk failure. NetApp strongly recommends using RAID-DP on all RAID groups that store VMware View data. For more information on RAID-DP, refer to NetApp white paper 3298 at www.netapp.com/us/library/white-papers/wp_3298.html.

REMOTE LAN MANAGEMENT (RLM) CARD

The RLM card improves storage system monitoring by providing secure out-of-band access to the storage controllers, which can be used regardless of the state of the controllers. The RLM offers a number of remote management capabilities for NetApp controllers, including remote access, monitoring, troubleshooting, logging, and alerting features. The RLM also extends AutoSupport capabilities of the NetApp controllers by sending alerts or "down storage system" notification with an AutoSupport message when the controller goes down, regardless of whether the controller can send AutoSupport messages. These AutoSupport messages also provide proactive alerts to NetApp to help provide faster service. For more details on RLM, visit http://now.netapp.com/NOW/download/tools/rlm_fw/info.shtml.

NETWORKING INFRASTRUCTURE DESIGN (FIBRE CHANNEL OR IP)

A network infrastructure (FC or IP) should have no single point of failure. A highly available solution includes having two or more FC or IP network switches, two or more HBAs or NICs per host, and two or more target ports or NICs per storage controller. In addition, if using Fibre Channel, two independent fabrics are required to have a truly redundant architecture.

For additional information on designing and deploying an FC solution, refer to the FC/iSCSI configuration guide at

http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/fc_iscsi_config_guide.pdf.

For additional information on designing and deploying an IP solution, refer to [TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices](#).

9.1.2 Top Resiliency Best Practices

- Use RAID-DP, the NetApp high-performance implementation of RAID 6, for better data protection.
- Use multipath HA with active-active storage configurations to improve overall system availability as well as promote higher performance consistency.
- Use the default RAID group size (16) when creating aggregates.
- Allow Data ONTAP to select disks automatically when creating aggregates or volumes.
- Use the latest Data ONTAP general deployment release available on the NOW site.
- Use the latest storage controller, shelf, and disk firmware available on the NOW site.
- Disk drive differences are FC, SAS, SATA disk drive types, disk size, and rotational speed (RPM).
- Maintain two hot spares for each type of disk drive in the storage system to take advantage of Maintenance Center.
- Do not put user data into the root volume.
- Replicate data with SnapMirror or SnapVault for disaster recovery (DR) protection.
- Replicate to remote locations to increase data protection levels.
- Use an active-active storage controller configuration (clustered failover) to eliminate single points of failure (SPOFs).
- Deploy SyncMirror® and RAID-DP for the highest level of storage resiliency.

For more details, refer to [NetApp TR-3437: Storage Best Practices and Resiliency Guide](#).

9.1.3 Building a High-Performance Storage Architecture

A VMware View workload can be very I/O intensive, especially during the simultaneous bootup, login, and virus scan within the virtual desktops. These first two workloads are commonly known as a “boot storm” and “login storms.” A boot storm, depending on how many ESX Servers and guests are attached to the storage, can create a significant performance impact if the storage is not sized properly. A boot storm can affect both the speed in which the VMware desktops are available to the customer and overall customer experience. A “virus scan storm” is similar to a boot storm in I/O but might last longer and can significantly affect customer experience. A virus scan storm is when a virus scan within the guest is initiated on all the clients at once. Due to these factors, it is important to make sure that the storage is architected in such a way to eliminate or decrease the impact of these events.

AGGREGATE SIZING

An aggregate is NetApp’s virtualization layer, which abstracts physical disks from logical data sets that are referred to as flexible volumes. Aggregates are the means by which the total IOPS available to all of the physical disks are pooled as a resource. This design is well suited to meet the needs of an unpredictable and mixed workload. NetApp recommends that whenever possible a small aggregate should be used as the

root aggregate. This aggregate stores the files required for running and providing GUI management tools for the storage system. The remaining storage should be placed into a small number of large aggregates. The overall disk I/O from VMware environments is traditionally random by nature, so this storage design gives optimal performance because a large number of physical spindles are available to service I/O requests. On smaller storage systems, it might not be practical to have more than a single aggregate, due to the restricted number of disk drives on the system. In these cases, it is acceptable to have only a single aggregate.

HARD DISK DRIVE CHARACTERISTICS

When architecting a VMware View storage environment, it is best to use FC disk drives. These drives have the ability to serve more I/Os per second than a 7200 RPM SATA. The speed of the disks might not improve performance on a per desktop basis, but during a boot or virus storm, these higher speed FC disk drives will decrease boot time and allow users to log on sooner.

PAM

The PAM is an intelligent read cache that improves read performance and in turn increases throughput and decreases latency. It provides greater system scalability by removing IOPS limitations due to disk bottlenecks and lowers cost by providing the equivalent performance with fewer disks. Leveraging the PAM card in a dense (deduplicated) volume allows all the shared blocks to be accessed directly from the intelligent, faster PAM card versus disk. The PAM card provides great benefits in a VMware View environment, especially during a boot storm, login storm, or virus storm, as only one copy of deduplicated data will need to be read from the disk (per volume). Each subsequent access of a shared block will be read from PAM cache and not from disk, increasing performance and decreasing latency and overall disk utilization.

DISK CONFIGURATION SUMMARY

When sizing your disk solution, consider the number of desktops being served by the storage controller/disk system and the number of IOPS per desktop. This way one can make a calculation to arrive at the number and size of the disks needed to serve the given workload. Remember, keep the aggregates large, the spindle count high, and rotational speed fast. When one factor needs to be adjusted, the PAM card can help eliminate any potential bottlenecks to the disk.

FLEXIBLE VOLUMES

Flexible volumes contain either LUNs or virtual disk files that are accessed by VMware ESX Servers. NetApp recommends a one-to-one alignment of VMware datastores to flexible volumes. This design offers an easy means to understand the VMware data layout when viewing the storage configuration from the storage system. This mapping model also makes it easy to implement Snapshot backups and SnapMirror replication policies at the datastore level, because NetApp implements these storage side features at the flexible volume level.

LUNS

LUNs are units of storage provisioned from a NetApp storage controller directly to the ESX Servers. The LUNs presented to the ESX Server are formatted with the VMware File System (VMFS). This shared file system is capable of storing multiple virtual desktops and is shared among all ESX Servers within the HA/DRS cluster. This method of using LUNs with VMFS is referred to as a VMFS datastore. For more information, see the VMware Storage/SAN Compatibility Guide for ESX Server 3.5 and ESX Server 3i.

10 DEPLOYING NETAPP SPACE-EFFICIENT VM CLONES ON NFS

This chapter demonstrates the steps involved in deploying NetApp space-efficient clones using the NetApp [Rapid Cloning Utility version 2.0](#). The NetApp Rapid Cloning Utility, or RCU, allow administrators to leverage the power of NetApp FlexClone (both file and volume), thin provisioning, and deduplication capabilities directly from the VMware vCenter Server GUI. Integrating NetApp capabilities into VMware vCenter Server allows VMware administrators to provision one to thousands of new virtual machines in minutes without requiring the administrator to log in to the storage. RCU utilizes both the NetApp and VMware APIs to create a robust, fully supported solution. No end user customization or scripting is required. In this section, we will demonstrate how to properly configure the virtual machine template and deploy one or thousands of virtual machines right from the VMware vCenter Server interface.

NOTE: Windows XP has been chosen as the reference desktop OS because it makes up the largest share of virtual desktop environments today. RCU is not limited to deploying Windows XP desktops but supports Windows 2003 as well for deploying virtual servers.

10.1 OVERVIEW OF DEPLOYING NETAPP SPACE-EFFICIENT CLONES ON NFS




Figure 17) NetApp RCU deployment overview.

10.1.1 Provision NetApp Storage

This step involves preparing the NetApp storage for provisioning VMs. The detailed steps involved are as follows:

1. Create aggregate.
2. Create gold volume.
3. Configure the NFS export.
4. Set NFS tunables.
5. Mount the gold datastore on ESX Server(s).

10.1.2 Build Template Virtual Machine

This step involves creating and customizing the template VM that will be used to deploy the VMs within the environment. The detailed steps involved are as follows:

1. Create virtual machine for use as a template.
2. For Windows XP virtual machines, perform guest partition alignment for the empty vmdk as per instructions in [TR-3747](#).
3. Install Windows XP on the template VM.
4. Disable NTFS last access.
5. Change disk timeout value.
6. Install all necessary applications and modify any additional system settings.
7. Power off VM and mark as template.

10.1.3 Deploy Space-Efficient Clones with RCU

This step involves using the NetApp RCU to deploy virtual machines from the template VM. This step assumes that the NetApp RCU has already been installed and configured on a server. The detailed steps involved are as follows:

1. Create a customization specification within VMware vCenter Server.
2. Use the NetApp RCU to create the NetApp volumes, attach the datastores, provision the virtual machines, and import into VMware View Manager 3.0.

10.2 DETAILS OF DEPLOYING NETAPP SPACE-EFFICIENT CLONES ON NFS

10.2.1 Provision NetApp Storage

The first step in the VMware View mass deployment process is to provision NetApp storage. Figure 18 shows the complete process.




Figure 18) Provision NetApp storage.

CREATE AGGREGATE

Create aggregates on the storage system as per the design. The details on storage design considerations are discussed in section 8. Basic recommendations include use RAID-DP and configure fewer, larger aggregates but make sure to follow the storage design.

CREATE GOLD VOLUME

As per the storage design, create one gold volume on one of the storage controllers. This datastore will be used to house the gold template VM or VMs. Follow the configuration recommendations mentioned in [NetApp TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices](#), specifically for thin provisioning, automatic Snapshot copies, volume auto growth, Snapshot auto delete, and other generic volume settings (for example no_atime_update), qtree security style, and so on. If a gold/template volume already exists within the environment, verify the volume has the proper settings.

CONFIGURE THE NFS EXPORTS

Export the gold volume to at least one ESX Server in the cluster to the VM kernel IP address. Make sure the export is persistent across reboot and also provide root and read/write access to the ESX Server. Persistence can be achieved by using either the FilerView® interface, using the export -p command from the command line, or manually editing the /etc/exports file and using exportfs -r to reread the exports file. Follow the NFS datastore provisioning recommendations mentioned in [NetApp TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices](#).

SET NFS TUNABLES

NetApp recommends setting the following ESX values if using NFS to optimize scalability and resiliency.

Table 12) Setting ESX NFS values.

Step	Action
1	Connect to the controller system's console, using either SSH, telnet, or serial console.
2	<code>esxcfg-advcfg -s 32 /NFS/MaxVolumes</code>
3	<code>esxcfg-advcfg -s 12 /NFS/HeartbeatFrequency</code>
4	<code>esxcfg-advcfg -s 5 /NFS/HeartbeatTimeout</code>
5	<code>esxcfg-advcfg -s 10 /NFS/HeartbeatMaxFailures</code>
6	<code>esxcfg-advcfg -s 30 /Net/TcpIpHeapSize</code>
7	<code>esxcfg-advcfg -s 120 /Net/TcpIpHeapMax</code>
8	Reboot the ESX Servers for the changes to take effect.

NetApp recommends setting the following Data ONTAP value for `nfs.tcp.recvwindowsize` to make more efficient use of the network by reducing the frequency of acknowledgements between client and server.

Table 13) Setting Data ONTAP NFS TCP receive window size.

Step	Action
1	Connect to the controller system's console, using either SSH, telnet, or serial console.
2	On each NetApp controller enter: <code>options nfs.tcp.recvwindowsize 64240</code>

MOUNT THE GOLD DATASTORE ON AN ESX SERVER

Follow the procedure mentioned in [NetApp TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices](#) to mount the gold datastore to at least one ESX Server using either ESX service console (not available in ESX 3i) or VMware vCenter Server.

10.2.2 Build Template Virtual Machine

In this phase, a VM is created that will be the template image for the virtual desktops that you plan to provision for the end users. This process should be familiar, because it is the same process that VMware administrators follow when creating virtual server templates. The process includes building a VM, performing the guest partition alignment, installing the OS, applying patches, installing applications and general system settings, and powering off the VM. This process starts by making sure that the guest OS, required business, and infrastructure applications are installed in the gold VM and it is in a condition to be duplicated. Figure 19 shows the process of building a template VM.




Figure 19) Build template virtual machine.

CREATE NEW VIRTUAL MACHINE IN VCENTER

The first step in the building of a virtual machine is to create a VM within the VMware vCenter Server. Use your company's hardware configuration for creating new virtual machines. Make sure to attach all virtual hardware, including networking that will be needed during production.

PERFORM GUEST PARTITION ALIGNMENT

Perform the guest partition alignment for the gold VM's VMDK prior to the installation of the guest operating system as described in [TR-3747: NetApp Best Practices for File System Alignment in Virtual Environments](#).

INSTALL OPERATING SYSTEM

Install Windows XP on the gold VM. For Windows XP desktops follow VMware recommendations in the Windows XP Deployment Guide at www.vmware.com/files/pdf/XP_guide_vdi.pdf. Install all the baseline applications and service packs.

DISABLING NTFS LAST ACCESS

It is important to change the value of last access on the gold VM as this will reduce the number of unnecessary I/Os to disk.

Table 14) Disabling NTFS last access.

Step	Action
1	Log in to the gold VM.
2	Open a CMD window by going to Start > Run and enter cmd and press Enter.
3	Check the current value for disable last access: <i>fsutil behavior query disablelastaccess</i>
4	If the value is zero enter the following to disable last access: <i>fsutil behavior set disablelastaccess 1</i>

For more information, visit <http://technet2.microsoft.com/windowsserver/en/library/9fcf44c8-68f4-4204-b403-0282273bc7b31033.mspx?mfr=true>.

CHANGE DISK TIMEOUT VALUE

Increasing this value increases the amount of time before a SCSI timeout within the guest OS.

Table 15) Changing disk timeout value.

Step	Action
1	Log in to the gold VM.
2	Open a regedit by going to Start > Run and enter regedit and press Enter.
3	Find the TimeOutValue by following the path [HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk].
4	Change the key "TimeOutValue"=dword:00000190.
5	Reboot the VM now or at the end of the installation of applications and general system settings.

INSTALL NECESSARY APPLICATIONS AND GENERAL SYSTEM SETTINGS

Install all the necessary infrastructure and business applications in the gold VM. A few examples include Microsoft Office, antivirus software, Adobe Acrobat Reader, a connection broker agent (if planning to use a connection broker); allow specific users or groups RDP access to the VMs.

POWER OFF GOLDEN VM AND MARK AS TEMPLATE

After performing all the template customizations and software installations, power off the virtual machine as they will need to be powered off to deploy. Then change the VM to be a template. This reduces the risk of accidental powering on of the VM.

10.2.3 Deploy Space-Efficient Clones with RCU

This section describes in detail the use of the Rapid Cloning Utilities to clone one to thousands of virtual machines. The NetApp Rapid Cloning Utility is a VMware vCenter Server plug-in that allows VMware administrators the ability to provision virtual machines directly from vCenter without having to leave the vCenter client.

This section assumes that the NetApp RCU has been installed and registered according to the RCU Installation and Administration Guide located [here](#).

The following are the detailed steps involved in the process.




Figure 20) Deploy space-efficient clones with RCU.

CREATE CUSTOMIZATION SPECIFICATION


Create a customization specification for use with the deployment of the VMs. The customization specification creates the information necessary for sysprep to successfully customize a guest OS from the VMware vCenter Server. It includes information on hostname, network configuration, license information, domain membership, and other information necessary to customize a guest OS. This procedure can be found in the VMware Basic System Administration Guide on page 224 at www.vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_admin_guide.pdf. This customization specification can be used by RCU to personalize each VM.


DEPLOY SPACE-EFFICIENT CLONES USING RCU 2.0



In this example 2,000 virtual machines will be deployed using RCU. RCU has already been installed on the VMware vCenter Server. RCU will be used to create eight datastores that have 250 virtual machines each. RCU will perform the following steps:



- 1. Create the clones with file FlexClone.
- 2. Clone the datastores with volume FlexClone.
- 3. Mount the NFS datastores to the vCenter Server.
- 4. Create the virtual machines from the cloned vmrk.
- 5. Customize the virtual machines using the customization specification.
- 6. Power on the virtual machines.
- 7. Import virtual machines into VMware View Manager.



Table 16) Deploy space-efficient clones using RCU 2.0.



Step	Action
1	Log in to the VMware vCenter Server and click the NetApp icon and enter the RCU GUI. 



Step	Action
2	Select the Storage Controllers tab and select "Add Storage Controller...." 
3	Enter the IP address of the storage controller as well as the username and password. If SSL has been enabled on the controller, select "Use SSL" If you are unsure if SSL has been enabled, check the box and try to connect. If enabled you will connect, if not it will reject the connection then uncheck the box and try again. 
4	Click Ok. 

Step	Action
5	Once storage controllers have been added, select the inventory button to get back to the servers and VMs. Right-click the VM to be cloned and select “Create NetApp Rapid Clones.” 
6	Choose the storage controller with the drop-down arrow and click Next.  Additionally, if the VMware VI client is not running, select Advanced Options and enter the password for the vCenter Server.

Step	Action
7	Select the data center, cluster, or server to provision the VMs to and select Next. 
8	Select where to place the new clones, in an existing datastore or new datastore, and click Next. 

Step	Action
9	Enter in the number of clones per datastore, the total number of datastores, the clone name prefix, and the starting clone number. Then if guest customization is required, select the check box and the customization specification that will be applied. Then choose if the virtual machine will be powered on after the clones are created. 
10	Then if using VMware View Manager one can select to automatically create an import file. Add the View Manager Server name or IP address, the domain name, username, and password. Then select if these VMs will be imported as individual desktops or manual desktop pool and specify the pool name. Then select the validate credentials to test the connection to the View Manager Server. 

Step	Action
11	After the credentials have been validated, click Ok. 
12	Select Next. 

Step	Action
13	Using the sizing and design information, provide the size of the datastore, the datastore prefix, and the hosting aggregate. Then select Thin Provisioning if you don't want to pre-allocate the storage from the aggregate for each volume. Additionally you can select to manually set the datastore names. 
14	Review the configuration and click Apply. 

IMPORT RCU DEPLOYED VMS INTO VMWARE VIEW MANAGER


In this step the virtual machines will be imported into VMware View Manager 3.0.

Table 17) Import RCU deployed VMs into VMware View Manager.

[illegible]

4

Click the Desktop Sources tab to verify the VMs that were created by RCU are now imported in View Manager and are part of the pool.



11 DEPLOYING NETAPP SPACE-EFFICIENT VM CLONES ON VMFS

The VMware View mass deployment process requires few basic steps, including the creation of a gold VM, building a gold datastore populated with several VMs and then creating production FlexClone datastores from the deduplicated gold datastore, and mass deployment. The mass deployment process utilizes the NetApp Rapid Cloning Utilities (RCU) and is available for download on the NetApp NOW Toolchest at <http://now.netapp.com/NOW/download/tools/rcu>. These tools are provided as is and are supported by PVR. To obtain official support from NetApp, contact your sales representative or sales engineer.

Windows XP has been chosen as the reference desktop OS because it makes up the largest share of virtual desktop environments today.

11.1 OVERVIEW OF DEPLOYING NETAPP SPACE-EFFICIENT CLONES ON VMFS

Figure 21 shows a high-level overview of the VMware View mass deployment process.




Figure 21) NetApp RCU deployment overview for VMFS.

11.1.1 Provision NetApp Storage

This step involves preparing the NetApp storage for provisioning VMs. The detailed steps involved are as follows:

1. Create aggregate.
2. Create gold volume.

3. Configure VMFS LUN.
4. Mount the gold datastore on an ESX host.

11.1.2 Build Template Virtual Machine

This step involves preparing the template VM that will be used as the basis creating the gold datastore. The detailed steps involved are as follows:

1. Create virtual machine for use as a template.
2. For Windows XP virtual machines, perform guest partition alignment as per instructions in [TR-3747: NetApp Best Practices for File System Alignment in Virtual Environments](#).
3. Install OS on template VM.
4. Disable NTFS last access.
5. Change disk timeout value.
6. Install all necessary applications and general system settings.
7. Sysprep the template VM and shut down.
8. Unregister gold VM from VMware vCenter Server.

11.1.3 Build Gold Datastore

This step involves populating a gold datastore with multiple copies of the vmdk. The detailed steps involved are as follows:

1. Populate gold datastore using the NetApp Rapid Cloning Utilities RFWM tool.
2. Deduplicate gold datastore.
3. Adjust volume size to realize deduplicated savings.

11.1.4 Deploy Space-Efficient Clones with RCU 1.0

This step involves using FlexClone to duplicate the gold datastore, attaching and configuring the guest OSs. The detailed steps involved are as follows:

1. Create template VM.
2. Create a customization specification within VMware vCenter Server.
3. Use the NetApp Rapid Cloning Utilities na_vi_clone.exe tool to:
 - a. Create FlexClone volumes
 - b. Attach FlexClone volumes to ESX Servers in the cluster
 - c. Clone the VMs
 - d. Perform guest customization
4. Power on VMs for use.

11.2 DETAILS OF DEPLOYING NETAPP SPACE-EFFICIENT CLONES ON VMFS

11.2.1 Provision NetApp Storage

The first step in the VMware View mass deployment process is to provision NetApp storage. Figure 22 shows the complete process.




Figure 22) Provision NetApp storage.

CREATE AGGREGATE

Create aggregates on the storage system as per the design. The details on storage design considerations are discussed in Section 8 earlier. Basic recommendations include, use RAID-DP, configure fewer, larger aggregates but make sure to follow the storage design.

CREATE GOLD VOLUME

As per the storage design, create the gold volume on one storage controller. The gold volume will be hosting the gold datastores that will contain several identical VMs. This datastore will be used for mass deployment. Follow the configuration recommendations mentioned in [NetApp TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices](#), specifically for thin provisioning, automatic Snapshot copies, generic volume settings, NFS-specific volume settings, and VMFS-specific volume settings.

CONFIGURE VMFS LUN

For VMFS (FC/iSCSI) datastores, create one or more LUNs inside the gold volume and map it to the ESX Server that will be used to build the gold VM. Follow the VMFS LUN provisioning best practices mentioned in [NetApp TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices](#).

MOUNT THE GOLD DATASTORE AND SET ESX AND NETAPP VALUES

Follow the procedure mentioned in [NetApp TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices](#) (for VMFS datastores) to mount the gold datastore on an ESX host using either ESX host console or VMware vCenter Server.

SETTING ESX TCP VALUES

NetApp recommends setting the following ESX values if using iSCSI to optimize scalability and resiliency.

Table 18) Setting ESX TCP values.



Step	Action
1	Connect to the controller system's console, using either SSH, telnet, or serial console.
2	<code>esxcfg-advcfg -s 30 /Net/TcpIpHeapSize</code>
3	<code>esxcfg-advcfg -s 120 /Net/TcpIpHeapMax</code>
4	Reboot the ESX Servers for the changes to take effect.

SETTING ESX VALUES WHEN USING VMFS FILE SYSTEMS

NOTE: This step should only be performed when connecting to FlexClone VMFS volumes. After connecting to a FlexClone volume, change the value back to the default of zero.

This action of connecting to FlexClone volumes may be performed when performing a restore operation or provisioning a new datastore from a FlexClone volume. Caution should be used when leaving this value set as unintended datastores might appear when rescanning the storage.

Table 19) Setting ESX values for VMFS file systems.



Step	Action
1	Enable VMFS resignaturing. This will allow volumes that have been created with FlexClone to be added to the ESX server. This can be done at the command line using the following command or by using the VI client. This step must be performed on each host within the ESX HA/DRS cluster.
2	On the command line of the service console of the ESX Servers run the following: <i>esxcfg-advcfg -s 1 /LVM/EnableResignature</i>
3	From VI client click Advanced Settings from the Hardware menu. 
4	Select LVM from the Advanced Settings window and change the LVM.EnableResignature from 0 to 1. 

CHECKING QUEUE DEPTH FOR HBA AND LUNS

By default, the ESX Servers should have appropriate queue depth values for the HBAs and LUNs. Should queuing occur on the ESX host, refer to the VMware Infrastructure 3 Fibre Channel SAN Configuration Guide at www.vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_san_cfg.pdf. Should the queue depth need to be changed also verify the Disk.SchedNumReqOutstanding value is set the same as the queue depth.

Table 20) Checking ESX Disk.SchedNumReqOutstanding values.

Step	Action
------	--------

Step	Action
1	Log in to the ESX service console.
2	On the command line of the service console of the ESX Servers run the following: <i>esxcfg-advcfg -g /Disk/SchedNumReqOutstanding</i>
3	Or from VI client click Advanced Settings from the Hardware menu. 
4	Select Disk from the Advanced Settings window and verify the Disk.SchedNumReqOutstanding value. Should the value need to be changed it can be done here. As you will see this example demonstrates the changing from the default of 32 to 64. 

11.2.2 Build Template Virtual Machine

In this phase, a VM is created that will be the template image for the virtual desktops that you plan to provision for the end users. This process should be familiar, because it is the same process that VMware administrators follow when creating virtual server templates. The process includes building a VM, performing the guest partition alignment, installing the OS, applying patches, installing applications and general system settings, and powering off the VM. This process starts by making sure that the guest OS, required business, and infrastructure applications are installed in the gold VM and it is in a condition to be cloned. Figure 23 shows the process of building a template VM.




Figure 23) Build template virtual machine.

CREATE NEW VIRTUAL MACHINE IN VCENTER

The first step in the building of a virtual machine is to create a VM within the VMware vCenter Server. Use your company's hardware configuration for creating new virtual machines. Make sure to attach all virtual hardware including networking that will be needed during production.

PERFORM GUEST PARTITION ALIGNMENT

Perform the guest partition alignment for the template VM's VMDK prior to the installation of the guest operating system as described in TR-3747: NetApp Best Practices for File System Alignment in Virtual Environments.

INSTALL OPERATING SYSTEM

Install Windows XP on the template VM. For Windows XP desktops follow VMware recommendations in the Windows XP Deployment Guide at www.vmware.com/files/pdf/XP_guide_vdi.pdf. Install all the baseline applications and service packs.

DISABLING NTFS LAST ACCESS

It is important to change the value of last access on the template VM as this will reduce the number of unnecessary I/Os to disk.

Table 21) Disabling NTFS last access.

Step	Action
1	Log in to the gold VM.
2	Open a CMD window by going to Start > Run and enter cmd and press Enter.
3	Check the current value for disable last access: <i>fsutil behavior query disablelastaccess</i>
4	If the value is zero enter the following to disable last access: <i>fsutil behavior set disablelastaccess 1</i>

For more information, visit <http://technet2.microsoft.com/windowsserver/en/library/9fcf44c8-68f4-4204-b403-0282273bc7b31033.mspx?mfr=true>.

CHANGE DISK TIMEOUT VALUE

Increasing this value increases the amount of time before a SCSI timeout within the guest OS.

Table 22) Changing disk timeout value.

Step	Action
1	Log in to the gold VM.
2	Open a regedit by going to Start > Run and enter regedit and press Enter.
3	Find the TimeOutValue by following the path [HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk].
4	Change the key "TimeOutValue"=dword:00000190.
5	Reboot the VM now or at the end of the installation of applications and general system settings.

INSTALL NECESSARY APPLICATIONS AND GENERAL SYSTEM SETTINGS

Install all the necessary infrastructure and business applications in the template VM. A few examples include Microsoft Office, antivirus software, Adobe Acrobat Reader, a connection broker agent (if planning to use a connection broker); allow specific users or groups RDP access to the VMs.

UNREGISTER TEMPLATE VM FROM VMWARE VCENTER SERVER

This step can either be done from the VMware vCenter Server by right-clicking the VM and clicking "Remove from Inventory" or from the ESX console by using the following command. This is done to make sure that the gold VM is not accidentally powered on and changed.

Table 23) Unregistering template VM from the ESX service console.

Step	Action
1	Log in to the ESX host console.
2	Run the following command: <code>vmware-cmd -s unregister /vmfs/volumes/<gold-datastore> /<gold-vm1>/<gold-vm1>.vmx</code>

11.2.3 Build Gold Datastore

The process of building a gold datastore involves using NetApp deduplication and FlexVol resizing. This step involves populating a gold datastore with multiple copies of the vmdk. The detailed steps involved are shown in Figure 24.




Figure 24) Build gold datastore.

POPULATE GOLD DATASTORE USING THE NETAPP RAPID CLONING UTILITIES RFWM

This step includes populating the gold datastore using the NetApp Rapid Cloning Utilities RFWM application. The advantage of using the NetApp Rapid Cloning Utilities RFWM application is that it is faster than performing an ndmccopy or copy operation and places less load on the NetApp storage system. This tool has the ability to create as many clones of the machine as necessary. Detailed instructions on using the NetApp Rapid Cloning Utilities RFWM can be found on the NetApp NOW Toolchest: <http://now.netapp.com/NOW/download/tools/rcu/>.

PERFORM DEDUPLICATION OF THE VOLUME HOSTING THE GOLD DATASTORE

After you have populated the gold datastore, you can use deduplication to reduce the storage requirement of the datastore down to that of an individual VM. Because the datastore was populated with identical data, the storage savings can be considerable. As show in the diagram, this step is optional and not required.

Table 24) Enable and initialize deduplication on the flexible volume.

Step	Action
1	Connect to the controller system's console, using either SSH, telnet, or serial console.
2	Execute the following command to enable deduplication for the gold volume: <i>sis on <gold volume path></i>
3	Execute the following command to start processing existing data: <i>sis start -s <gold volume path></i>
4	Execute the following command to monitor the status of the deduplication operation: <i>sis status</i>
5	After the deduplication has finished you can use the following command to see the savings: <i>df -s</i>

RESIZE GOLD VOLUME

When the deduplication process is complete, there might be a large amount of free space in the gold volume. The size of the gold volume can be reduced so that the additional space can be used for other volumes. It is not recommended to reduce the volume to exactly the size of the data in the gold volume. It is recommended to leave a couple of percent free space within the volume. Also factor in some space for Snapshot reserve, as this volume will be used as the basis for creating FlexClone volumes, which require a base Snapshot copy.

DEDUPLICATION CONSIDERATIONS WITH VMFS LUNS

Enabling deduplication when provisioning LUNs also produces storage savings. However, the default behavior of a LUN is to reserve an amount of storage equal to the provisioned LUN. This design means that although the storage array reduces the amount of storage capacity consumed, any gains made with deduplication are, for the most part, unrecognizable because the space reserved for LUNs is not reduced.

To recognize the storage savings of deduplication with LUNs, you must enable LUN thin provisioning. See [NetApp TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices](#) for details on thin provisioning. In addition, although deduplication reduces the amount of consumed storage, the VMware administrative team does not see this benefit directly in the VMware vCenter Server or ESX console because their view of the storage is at a LUN layer, and LUNs always represent their provisioned capacity, whether they are traditionally or thin provisioned at the storage end.

If deploying scalable VMware View solution (thousands of VMs), it is best practice to split the number of VMs per ESX cluster across both controllers of the NetApp storage cluster. The most efficient way to copy gold volume from one storage controller to another is by using the NetApp SnapMirror technology. SnapMirror preserves the deduplication savings. This step is not mandatory; it is more of a best practice for scaling a V solution.

To summarize, in this step, you build a deduplicated datastore populated with VMs to be used as a template to mass deploy VMs. There are other methods that will work to fill the gold datastore with copies of the gold VM directory. Feel free to use the tool of your preference.

11.2.4 Deploy Space-Efficient Clones with RCU 1.0

This step involves cloning the gold datastore using FlexClone, attaching to ESX Servers, and configuring the guest OSs. The detailed steps involved follow.




Figure 25) Deploy space-efficient clones with RCU 1.0.

CREATE TEMPLATE VIRTUAL MACHINE

The purpose of this VM is to be used to clone and attach the cloned vmdks. This template will not have a VMDK.

Table 25) Create template VM.

Step	Action
1	Log in to VMware vCenter Server.
2	Create a new VM as documented in the VMware Basic System Administration Guide on page 143 at www.vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_admin_guide.pdf .
3	Edit newly created VM and remove virtual disk.
4	Right-click the newly created VM and "Convert to Template."

CREATE CUSTOMIZATION SPECIFICATION

Create a customization specification for use with the deployment of the VMs. The customization specification creates the information necessary for sysprep to successfully customize a guest OS from the VMware vCenter Server. It includes information on hostname, network configuration, license information, domain membership, and other information necessary to customize a guest OS. This procedure can be found in the VMware Basic System Administration Guide on page 224 at www.vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_admin_guide.pdf.

RUN NA_VI_CLONE TOOL

NetApp Rapid Cloning Utilities na_vi_clone.exe is a tool used to deploy hundreds or thousands of virtual machines with little to no administrative effort. The application handles the following tasks, which if left to do manually would take days of administrative work to accomplish:

- Creation of the FlexClone volumes of the golden datastore. FlexClone gives NetApp the unique ability to create a near instantaneous, space-efficient clone of a datastore that potentially could contain hundreds of VMs with no performance penalty. NetApp has the ability to create hundreds of FlexClone volumes from a single datastore and do so incredibly efficiently.
- After the FlexClone volumes are created, the tool attaches the volumes created with FlexClone to the ESX Servers in the data center.
- Once the storage has been provisioned, NetApp Rapid Cloning Utilities na_vi_clone.exe deploy the VMs from the template that was created in the first step. Since the disk was removed from the

configuration in the first step, the cloning process of a blank VM takes only a second. The vmdk files are then moved into the VM's directory and a reconfiguration is done to add the virtual disk to the VM.

- Running the guest customization on each new VM. This involves applying the customization specification created in the second step to each new clone. This step configures the individual VMs so that they will be licensed, connected to the network, be a part of Active Directory, and ready for use. This step is optional and can be performed later based on the requirements.
- The final step is to power on all the VMs, and add them to the connection broker of your choice. This step is optional and can be performed later based on the requirements.

12 VMWARE VIEW OPERATIONS BEST PRACTICES

12.1 DATA DEDUPLICATION

Production VMware View environments can benefit from the cost savings associated with NetApp deduplication as discussed earlier. Each VM will consume storage as new writes happen. Scheduling and monitoring deduplication operations for the NetApp volumes hosting VMs are very important.

DEDUPLICATION FOR VMS PROVISIONED AS VMWARE FULL CLONES

Using NetApp deduplication, VMs provisioned using VMware full clones can also achieve similar storage savings as seen with the use case of provisioning VMs with NetApp RCU. Follow these steps to configure deduplication on the datastores hosting these VMs:

Table 26) Deduplication for VMs provisioned as VMware full clones.

Step	Action
1	Connect to the controller system's console, using either SSH, telnet, or serial console.
2	Execute the following command to enable deduplication for the volume(s) hosting the datastore(s): <i>sis on <volume path></i>
3	Execute the following command to start processing existing data in the datastores: <i>sis start -s <volume path></i>
4	Execute the following command to monitor the status of the deduplication operation: <i>sis status</i>
5	After the deduplication has finished you can use the following command to see the savings: <i>df -s</i>

CONFIGURING DEDUPLICATION SCHEDULES

It is important to schedule the deduplication operations to run during off peak hours so that there is no impact on the end user experience. Also, it is important to understand the number of simultaneous dedupe operations that can be performed on the storage controller. Planning for dedupe operations will ultimately depend on your environment. There are multiple scheduling options available:

- Specific days of the week and hours, for example, run every day from Sunday to Friday at 11 p.m.
- Automatic: This means that deduplication is triggered by the amount of new data written to the flexible volume, specifically when there are 20% new fingerprints in the change log.
- Specific hour ranges on specific day(s)

Table 27) Configuring deduplication schedules.

Step	Action
1	Connect to the controller system's console, using either SSH, telnet, or serial console.

Step	Action
2	Execute the following command to set the deduplication schedule for the volume(s) hosting the datastore(s): <pre>sis config [-s sched]<volume path></pre> Where "sched" can be in four ways: <pre>[day_list][@hour_list]</pre> <pre>[hour_list][@day_list]</pre> - auto The day_list specifies which days of the week deduplication should run. It is a comma-separated list of the first three letters of the day: sun, mon, tue, wed, thu, fri, sat. The names are not case sensitive. Day ranges such as mon-fri can also be used. The default day_list is sun-sat. The hour_list specifies which hours of the day deduplication should run on each scheduled day. The hour_list is a comma-separated list of the integers from 0 to 23. Hour ranges such as 8-17 are allowed. Step values can be used in conjunction with ranges. For example, 0-23/2 means "every 2 hours." The default hour_list is 0; that is, midnight on the morning of each scheduled day. If "-" is specified, there is no scheduled deduplication operation on the flexible volume. The auto schedule causes deduplication to run on that flexible volume whenever there are 20% new fingerprints in the change log. This check is done in a background process and occurs every hour.

MONITORING DEDUPLICATION OPERATIONS

Deduplication operations should be monitored carefully as it might require readjusting the schedules for multiple reasons as the environment scales. For example, the deduplication schedule for a new volume (storage controller on the East Coast) hosting a datastore representing a set of users on the West Coast starts too early to be running during production hours.

Pooling virtual desktops with similar characteristics on the same datastore(s) will make it easier to manage dedupe schedules.

- The status of dedupe operations can be monitored using the `sis status` command on the storage controller console.
- The storage savings can be seen using the `df -s` command on the storage controller console.

For further details on NetApp deduplication, refer to [NetApp TR-3505: NetApp Deduplication for FAS, Deployment and Implementation Guide](#).

12.2 ANTIVIRUS OPERATIONS

For antivirus operations, one could either take a performance hit during scheduled AV operations and impact the end user experience or appropriately design the VMware View solution to make the AV operations seamless. The first option is definitely not desirable. The second option can be approached in two different ways:

- Optimize the AV operation policies for VMware View. Since VMware View involves moving from a completely distributed CPU (on the end user desktops) to centralizing much of the processing (in the VMs), the overall AV model should be thought about in a different way. Optimizing the traditional AV policies means better planning the scheduled AV scan and virus definition update so that not all the virtual desktops run AV scan or virus definition updates at the same time, creating CPU contention within the environment. By staggering the scheduled AV operations and distributing the load at different

points in time, you can avoid a large percentage of this contention. In addition to modifying the schedules, it is important to verify that these schedules do not interfere with other scheduled events such as backup or replication. In addition, we suggest that AV scanning of CIFS home directories should be done on the storage side, where the storage arrays and AV servers can dedicate processing to this activity. This will take some load off the virtual desktops. For more details, read TR-3107: NetApp Antivirus Scanning Best Practices Guide.

- Select intelligent, optimized, low-cost components for the VMware View solution that can effectively deal with the bursty performance requirement (like AV operations) without increasing the overall cost of the solution.

Optimizing the AV operations for thousands of virtual desktops is not straightforward and requires even more intelligence, especially from the back-end shared storage. NetApp intelligent caching and deduplication add a lot of value here. This not only significantly reduce the storage requirements for the otherwise redundant VMware View data, but also provides the capability to effectively deal with the bursty AV operations, without increasing the overall costs or inducing any negative tradeoffs.

12.3 MONITORING NETAPP AND VMWARE VIEW INFRASTRUCTURE

NETAPP OPERATIONS MANAGER

As discussed earlier, NetApp Operations Manager is a comprehensive monitoring and management solution for the VMware View storage infrastructure. It provides comprehensive reports of system utilization and trends for capacity planning, space usage and so on. It also monitors system performance and health to resolve potential problems. For further details on Operations Manager, visit the following solutions page: www.netapp.com/us/products/management-software/operations-manager.html.

NETAPP ESUPPORT

NetApp's proactive eSupport suite provides an early warning system that can actually reduce the number and severity of technical support cases. Automation tools identify issues early, before they have an impact, and can initiate a fix without customer burden, before people even know there's a potential problem. Support Automation works 24x7 to benchmark system status, collect diagnostic information behind the scenes, and issue proactive alerts. You can view the full scope of your NetApp environment on demand, at the company or device level.

NetApp's eSupport suite of Support Automation tools includes:

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

For more information on NetApp eSupport, visit the following: www.netapp.com/us/support/esupport.html.

SANSCREEN VM INSIGHT

As discussed earlier, consider implementing the NetApp SANscreen VM Insight. It provides cross-domain visibility from the VM to the shared storage, allowing both storage and server administration teams to more easily manage their VMware View storage and server architectures. For further details on SANscreen VM Insight, visit the following solutions page: www.netapp.com/us/products/management-software/sanscreen/sanscreen-vm-insight.html.

12.4 DATA PROTECTION SOLUTION

SNAPMANAGER FOR VIRTUAL INFRASTRUCTURE

As discussed earlier, NetApp SMVI is a unique, scalable, integrated data protection solution and is an excellent choice for protecting VMware View environments. It integrates VMware snapshots with the NetApp array-based block-level Snapshot copies to provide consistent backups for the virtual desktops. It is NetApp primary storage data deduplication aware and also integrates with the NetApp SnapMirror replication technology, which preserves the storage savings across the source and destination storage arrays. You do

not need to rerun dedupe on the destination storage array. SMVI also provides a user-friendly GUI that can be used to manage the data protection schemes. The following are some of the important benefits of SMVI:

- SMVI Snapshot backups are based on the number of 4KB blocks changed since the last backup, as compared to the number of files changed in a traditional backup solutions (which for virtual desktops can be several GB in size). This means that significantly fewer resources are required and the backups can be completed well within the window. Also, since SMVI is a storage block-based data protection solution, daily full backups of vmdk files are not required, resulting in a lower TCO.
- Backup failure rate has always been a concern with traditional, server-based backup solutions because of various moving parts in the solution: for example, backup server farms, CPU and memory limitation per server, network bandwidth, backup agents, and so on. With a NetApp solution, the number of moving parts is significantly less and requires very few policies to manage the backups for thousands of virtual desktops as multiple datastores can be part of the same protection policy, also resulting in higher success rate and not introducing new operational complexities.
- There will be significantly less net-new investment for backup infrastructure as SMVI leverages the inherent capabilities in the NetApp storage array to perform backups and does not require new backup server farms.
- SMVI allows flexible Snapshot retention scheduling that can be hourly, daily, and weekly to allow you to meet your level of RTO and RPO objectives.
- With SMVI, there are no concerns about the growth or existence of the vmdk delta files while the hot VM backups are being performed because SMVI does not require the existence of the vmdk delta files for the entire duration of backups. They have to exist only as part of the preprocessing step before SMVI invokes the NetApp Snapshot copy on the storage array, which can be few seconds to few minutes. Also, you have an option to configure only the NetApp Snapshot copies and ignore performing VMware snapshots.
- Also, since the backups are storage array based and only deduplicated data gets transferred from the primary storage array to the secondary array, there are significantly less storage requirements on the secondary storage array. Also the resource utilization on the various solution components (servers, storage, network, and so on) will be significantly less.

Scheduling considerations: Running of data protection policies should be properly planned to make sure that they don't interfere with the deduplication operations. Backup jobs should be scheduled to run after the deduplication operations so that minimum possible data has to be replicated.

VMware View Manager Considerations: As discussed earlier, NetApp data protection solution can be leveraged for all the seven virtual desktop types. An excellent use case is the data protection for "user data disk," where the user data is encapsulated in a vmdk file. Protecting this data using traditional methods requires performing full backups of vmdk files every day or every time when the policy is scheduled to run. Another important consideration for user data change rate is the delta associated with the "rebalance" operations performed on user data disk.

13 SUMMARY

To summarize, VMware View enables organizations to increase corporate IT control, manageability, and flexibility without increasing cost and providing end users with a familiar desktop experience. NetApp's key value proposition of at least 50% savings in storage, power, and cooling requirements; performance acceleration, operational agility, and best-in-class data protection and business continuance solution makes it a perfect choice as a solution for storage and data management for VMware View. The key NetApp technologies (RAID-DP, thin provisioning, FlexClone, deduplication, Snapshot copies, 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 provides details on the best integration points for each of the key enabling NetApp technologies and how each of the technology concepts plays a critical role and complements each other to work together as an integrated NetApp solution for VMware View of any scale.

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

14 ACKNOWLEDGEMENTS

The authors of this best practice guide would like to thank George Costea, Eric Forgette, Brian Gracely, Peter Learmonth, Jack McLeod, Mike Slisinger, Vaughn Stewart, Larry Touchette, and Paul Updike for their contributions to this document.

15 FEEDBACK

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

16 REFERENCES

[Total Cost Comparison: IT Decision-Maker Perspectives on EMC and NetApp Storage Solutions in Enterprise Database Environments](#)

[Wikipedia RAID Definitions and Explanations](#)

[VMware Introduction to Virtual Infrastructure](#)

[ESX Server 3i Configuration Guide](#)

[Storage/SAN Compatibility Guide for ESX Server 3.5 and ESX Server 3i](#)

[VMware View Manager System Administrator Guide](#)

[VMware View Composer Deployment Guide](#)

[VMware VMworld Conference Sessions Overview](#)

[VMware Recommendations for Aligning VMFS Partitions](#)

[VMware Basic System Administration Guide](#)

[NetApp VM Insight with SANscreen](#)

[NetApp TR-3612: NetApp and VMware Virtual Desktop Infrastructure](#)

[NetApp TR-3001: A Storage Networking Appliance](#)

[NetApp TR-3466: Open Systems SnapVault \(OSSV\) Best Practices Guide](#)

[NetApp TR-3348: Block Management with Data ONTAP 7G: FlexVol, FlexClone, and Space Guarantees](#)

[NetApp TR-3298: RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection](#)

[NetApp TR-3563: NetApp Thin Provisioning](#)

[NetApp TR-3347: FlexClone Volumes: A Thorough Introduction](#)

[NetApp TR-3505: NetApp Deduplication for FAS, Deployment and Implementation Guide](#)

[NetApp TR-3459: FlexShare Design and Implementation Guide](#)

[NetApp TR-3066: Data Protection Strategies for NetApp Storage Systems.](#)

[NetApp TR-3671: VMware Site Recovery Manager in a NetApp Environment](#)

[NetApp TR-3446: SnapMirror Best Practices Guide](#)

[NetApp TR-3737: SMVI Best Practices](#)

[NetApp Rapid Cloning Utilities](#)

17 VERSION HISTORY

Table 28) Version history.

Version	Date	Document Version History
Version 1.0	September 2008	Original document
Version 2.0	November 2008	Updates to intelligent caching with NetApp PAM and deduplication
Version 3.0	March 2009	Update for VMware View Manager, intelligent caching, sizing, operational best practices, and RCU 2.0



www.netapp.com

© 2009 NetApp. All rights reserved. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FilerView, FlexClone, FlexShare, FlexVol, NearStore, NOW, RAID-DP, SANscreen, SnapManager, SnapMirror, Snapshot, SnapVault, SyncMirror, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft and Windows are registered trademarks of Microsoft Corporation. VMware is a registered trademark and VMsafe is a trademark of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. **TR-3705**