

EMC Infrastructure for Virtual Desktops

Enabled by EMC Celerra Unified Storage (FC),
VMware vSphere 4.1, VMware View 4.5, and
VMware View Composer 2.5

Proven Solution Guide

EMC Unified Storage Solutions



vmware®

Copyright © 2010 EMC Corporation. All rights reserved.

Published August, 2010

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, this workload should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.

All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly.

EMC Corporation does not warrant or represent that a user can or will achieve similar performance expressed in transactions per minute.

No warranty of system performance or price/performance is expressed or implied in this document. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Part number: h8042

Table of Contents

Chapter 1: About this document	4
Overview	4
Audience and purpose	5
Scope	5
Reference architecture	6
Hardware and software resources	8
Prerequisites and supporting documentation.....	10
Terminology.....	11
Chapter 2: VMware View Infrastructure	12
Overview	12
VMware View 4.5	13
vSphere 4.1 infrastructure.....	15
Windows infrastructure.....	16
Chapter 3: Storage Design.....	17
Overview	17
CLARiiON storage architecture.....	18
Celerra storage architecture.....	20
Chapter 4: Network Design	22
Overview	22
Considerations	23
Celerra network configuration	24
ESX network configuration.....	25
Cisco 6509 configuration.....	26
Fibre Channel network configuration	27
Chapter 5: Installation and Configuration.....	28
Overview	28
VMware components	29
Storage components	32
Chapter 6: Testing and Validation.....	36
Overview	36
Testing overview	37
Boot storm results	40
View refresh results.....	46
View recompose results	52
Antivirus results.....	58
Patch install results	65
LoginVSI results	71

Chapter 1: About this document

Overview

Introduction

EMC's commitment to consistently maintain and improve quality is led by the Total Customer Experience (TCE) program, which is driven by Six Sigma methodologies. As a result, EMC has built Customer Integration Labs in its Global Solutions Centers to reflect real-world deployments in which TCE use cases are developed and executed. These use cases provide EMC with an insight into the challenges that are currently facing its customers.

This Proven Solution Guide summarizes a series of best practices that were discovered, validated, or otherwise encountered during the validation of the EMC Infrastructure for Virtual Desktops Enabled by EMC Celerra Unified Storage (FC), VMware vSphere 4.1, VMware View 4.5, and VMware View Composer 2.5 solution by using the following products:

- EMC® Celerra® unified storage
- VMware View Manager 4.5
- VMware View Composer 2.5
- VMware vSphere 4.1
- EMC PowerPath® Virtual Edition

Use case definition

The following six use cases are examined in this solution:

- Boot storm
- View refresh operation
- View recompose operation
- Antivirus scan
- Security patch install
- User workload simulated with LoginVSI tool

[Chapter 6: Testing and Validation](#) contains the test definitions and results for each use case.

Contents

This chapter contains the following topics:

Topic	See Page
Audience and purpose	5
Scope	5
Reference architecture	6
Hardware and software resources	8
Prerequisites and supporting documentation	10
Terminology	11

Audience and purpose

Audience The intended audience for this Proven Solution Guide is:

- Internal EMC personnel
 - EMC partners
 - Customers
-

Purpose The purpose of this use case is to provide a consolidated and virtualized out-of-the-box solution for virtual desktops that are powered by VMware View 4.5, View Composer 2.5, VMware vSphere 4.1, EMC Celerra unified storage (FC), Fully Automated Storage Tiering (FAST), EMC FAST Cache, and storage pools.

This solution includes all the attributes required to run this environment, such as hardware and software, including Active Directory, and the required View configuration.

Information in this document can be used as the basis for a solution build, white paper, best practices document, or training. It can also be used by other EMC organizations (for example, the technical services or sales organization) as the basis for producing documentation for technical services or sales kit.

Scope

Scope This Proven Solution Guide contains the results of testing the EMC Infrastructure for Virtual Desktops Enabled by EMC Celerra Unified Storage (FC), VMware vSphere 4.1, VMware View 4.5, and VMware View Composer 2.5 solution. The objectives of this testing are to establish:

- A reference architecture of validated hardware and software that permits easy and repeatable deployment of the solution.
 - The storage best practices to configure the solution in a manner that provides optimal performance, scalability, and protection in the context of the midtier enterprise market.
-

Not in scope Implementation instructions are beyond the scope of this document. Information on how to install and configure VMware View 4.5 components, vSphere 4.1, and the required EMC products is out of the scope of this document. Links to supporting documentation for these products are supplied where applicable.

Reference architecture

Corresponding reference architecture


This Proven Solution Guide has a corresponding Reference Architecture document that is available on EMC Powerlink® and EMC.com. Refer to the *EMC Infrastructure for Virtual Desktops Enabled by EMC Celerra Unified Storage (FC), VMware vSphere 4.1, VMware View 4.5, and VMware View Composer 2.5 — Reference Architecture* for details.

If you do not have access to the document, contact your EMC representative.

The Reference Architecture and the results in this Proven Solution Guide are valid for 500 Windows 7 virtual desktops conforming to the workload described in [“Observed user workload”](#) on page 37.

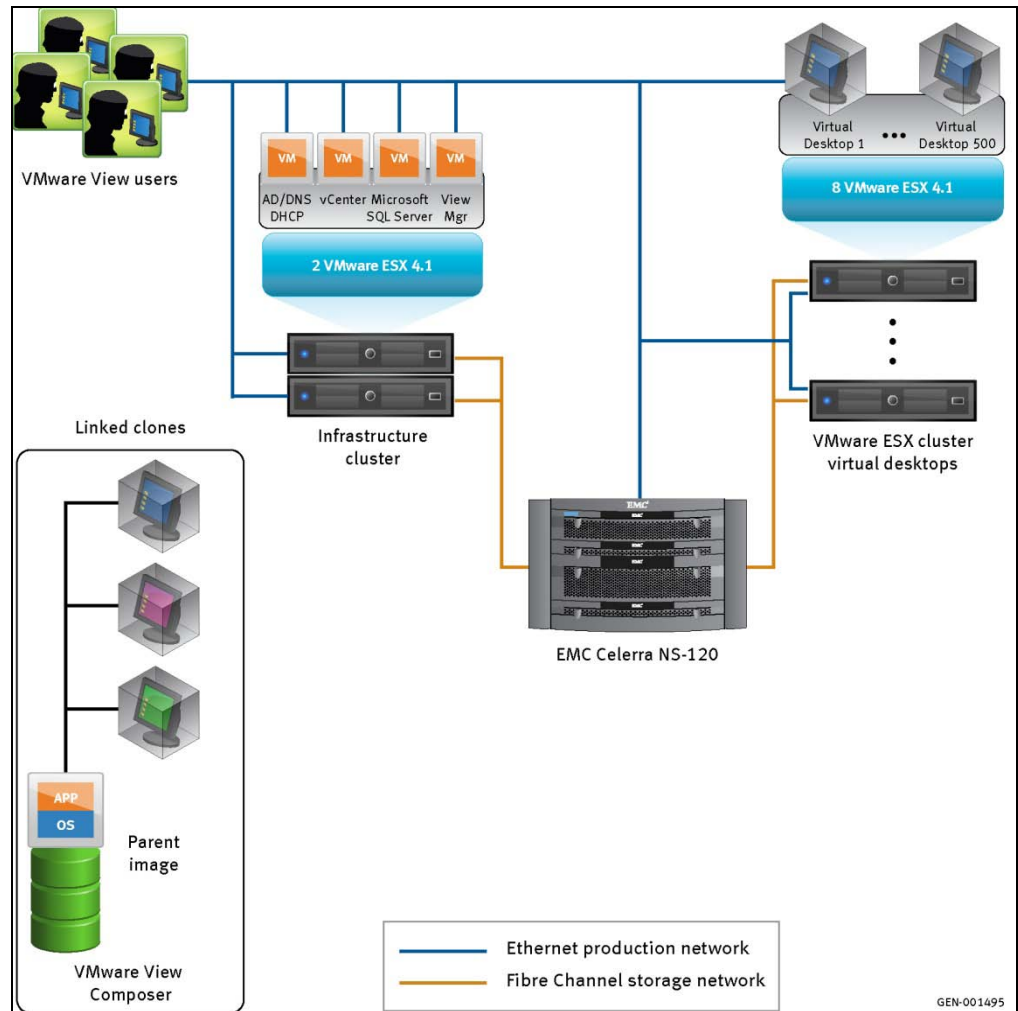
Reference architecture physical diagram

The following diagram depicts the overall physical architecture of this solution.



Reference architecture (logical)

The following diagram depicts the logical architecture of the midsize solution.



Hardware and software resources

Test results

[Chapter 6: Testing and Validation](#) provides more information on the performance results.

Hardware resources

The following table lists the hardware used to validate the solution.

Hardware	Quantity	Configuration	Notes
EMC Celerra NS-120	1	Three DAEs configured with: <ul style="list-style-type: none">• Five 300 GB 15k rpm FC disks• Fifteen 450 GB 15k rpm disks• Fifteen 1 TB 7.2k rpm SATA disks• Five 100 GB EFDs	Celerra shared storage for file systems and snaps
Dell PowerEdge R710	8	<ul style="list-style-type: none">• Memory: 64 GB RAM• CPU: Dual Xeon X5550 @ 2.67 GHz• NIC: Quad-port Broadcom BCM5709 1000Base-T	Virtual desktop ESX cluster
Dell PowerEdge 2950	2	<ul style="list-style-type: none">• Memory: 16 GB RAM• CPU: Dual Xeon 5160 @ 3 GHz• NIC: Gigabit quad-port Intel VT	Infrastructure virtual machines (vCenter Server, DNS, DHCP, Active Directory, and RRAS)
Cisco 9509	1	<ul style="list-style-type: none">• WS-6509-E switch• WS-x6748 1 Gb line cards• WS-SUP720-3B supervisor	Host connections distributed over two line cards
Brocade DS5100	2	Twenty four 8 Gb ports	Redundant SAN A/B configuration
QLogic HBA	1	<ul style="list-style-type: none">• Dual-port QLE2462• Port 0 – SAN A• Port 1 – SAN B	One dual-port HBA per server connected to both fabrics
Desktop/virtual machines	Each	<ul style="list-style-type: none">• Windows 7 Enterprise 32-bit• Memory: 768 MB• CPU: 1 vCPU• NIC: e1000 (connectivity)	Peak active memory measured at 688 MB

Software resources

The following table lists the software used to validate the solution.

Software	Configuration
Celerra NS-120 (shared storage, file systems, and snaps)	
NAS/DART	Release 6.0
CLARiiON® FLARE®	Release 30
ESX servers	
ESX	ESX 4.1
vCenter Server	
OS	Windows 2008 R2
VMware vCenter Server	4.1
VMware View Manager	4.5
VMware View Composer	2.5
PowerPath Virtual Edition	5.4 SP2
Desktops/virtual machines	
Note: This software is used to generate the test load.	
OS	MS Windows 7 Enterprise (32-bit)
VMware tools	8.3.1
Microsoft Office	Office 2007 SP2
Internet Explorer	8.0.7600.16385
Adobe Reader	9.1.0
McAfee Virus Scan	8.7.0i Enterprise

Prerequisites and supporting documentation

Technology

It is assumed that the reader has a general knowledge of the following products:

- VMware vSphere 4.1
 - View Composer 2.5
 - VMware View 4.5
 - EMC unified storage platforms
-

Supporting documents

The following documents, located on Powerlink, provide additional, relevant information. Access to these documents is based on your login credentials. If you do not have access to the following content, contact your EMC representative.

- *EMC Infrastructure for Virtual Desktops Enabled by EMC Celerra Unified Storage (FC), VMware vSphere 4.1, VMware View 4.5, and VMware View Composer 2.5 — Reference Architecture*
 - *Deploying Microsoft Windows 7 Virtual Desktops with VMware View — Applied Best Practices Guide*
 - *EMC Performance Optimization for Microsoft Windows XP for the Virtual Desktop Infrastructure — Applied Best Practices*
 - *EMC Infrastructure for Deploying VMware View in the Enterprise EMC Celerra Unified Storage Platforms — Solutions Guide*
-

Third-party documents

The following documents are available on the VMware website:

- *Introduction to VMware View Manager*
 - *VMware View Manager Administrator Guide*
 - *VMware View Architecture Planning Guide*
 - *VMware View Installation Guide*
 - *VMware View Integration Guide*
 - *VMware View Reference Architecture*
 - *Storage Deployment Guide for VMware View*
 - *VMware View Windows XP Deployment Guide*
 - *VMware View Guide to Profile Virtualization*
-

Terminology

Introduction

This section defines the terms used in this document.

Term	Definition
FAST	Fully Automated Storage Tiering (FAST) is a pool-based feature of FLARE release 30 that supports scheduled migration of data to different storage tiers based on the performance requirements of individual 1 GB slices in a storage pool.
FAST Cache	EMC CLARiiON FAST Cache is a new feature introduced in FLARE release 30. This feature enables to use EFD as an expanded cache layer for the array.
Linked clone	Virtual desktop created by VMware View Composer from a writeable snapshot paired with a read-only replica of a master image.
Login VSI	A third-party benchmarking tool developed by Login Consultants that simulates real-world VDI workload by using an AutoIT script and determines the maximum system capacity based on the users' response time.
Replica	Read-only copy of a master image used to deploy linked clones.
VMware View Composer	Integrates effectively with VMware View Manager to provide advanced image management and storage optimization.

Chapter 2: VMware View Infrastructure

Overview

Introduction The general design and layout instructions described in this chapter apply to the specific components used during the development of this solution.

Contents This chapter contains the following topics:

Topic	See Page
VMware View 4.5	13
vSphere 4.1 infrastructure	15
Windows infrastructure	16

VMware View 4.5

Introduction

VMware View delivers rich and personalized virtual desktops as a managed service from a virtualization platform built to deliver the entire desktop, including the operating system, applications, and user data. With VMware View 4.5, desktop administrators virtualize the operating system, applications, and user data and deliver modern desktops to end users. With VMware View 4.5, you can get centralized automated management of these components for increased control and cost savings. VMware View 4.5 improves business agility while providing a flexible high-performance desktop experience for end users across a variety of network conditions.

Deploying VMware View components

This solution is deployed using a single View Manager server instance that is capable of scaling up to 2,000 virtual desktops. Deployments of up to 10,000 virtual desktops are possible by using multiple View Manager servers.

The core elements of a VMware View 4.5 implementation are:

- View Manager Connection Server
- View Composer 2.5
- vSphere 4.1

Additionally, the following components are required to provide the infrastructure for a View 4.5 deployment:

- Microsoft Active Directory
 - Microsoft SQL Server
 - DNS Server
 - DHCP Server
-

View Manager Connection Server

The View Manager Connection Server is the central management location for virtual desktops and has the following key roles:

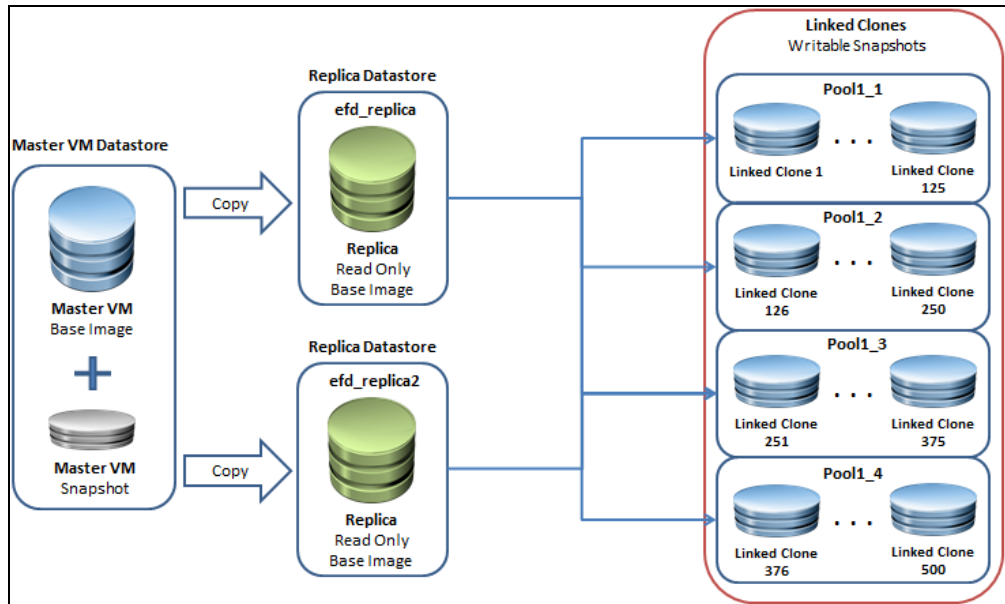
- Broker connections between users and virtual desktops
 - Control the creation and retirement of the virtual desktop images
 - Assign users to desktops
 - Control the state of the virtual desktops
 - Control access to the virtual desktops
-

View Composer 2.5


View Composer 2.5 works directly with vCenter Server to deploy, customize, and maintain the state of virtual desktops when using linked clones. The tiered storage capabilities of View Composer 2.5 allow the read-only replica and the linked clone disk images to be on the dedicated storage. This allows superior scaling in large configurations.

View Composer linked clones

VMware View with View Composer 2.5 uses the concept of linked clones to quickly provision virtual desktops. This solution uses the new tiered storage feature of View Composer to build linked clones and place their replica images on separate datastores as shown in the following diagram:



In this configuration, the operating system reads from the common read-only replica image and writes to the linked clone. Any unique data created by the virtual desktop is also stored in the linked clone. A logical representation of this relationship is shown in the following diagram:



Automated pool configuration

- All 500 desktops are deployed in two automated desktop pools by using a common Windows 7 master image. Dedicated datastores are used for the replica images and linked clone storage. The linked clones are distributed across four datastores:
- efd_replica and efd_replica2 contain the read-only copies of the master Windows 7 image. These datastores are backed by two 100 GB EFDs. Each replica image supports 256 linked clones.
 - Pool1_1 through Pool1_4 are used to store the linked clones. Each desktop pool is configured to use all four Pool1_x datastores so that all the virtual desktops are evenly distributed across the datastores.

vSphere 4.1 infrastructure


vSphere overview

VMware vSphere 4.1 is the market-leading virtualization hypervisor used across thousands of IT environments around the world. VMware vSphere 4.1 can transform or virtualize computer hardware resources, including CPU, RAM, hard disk, and network controller to create a fully functional virtual machine that runs its own operating system and applications just like a physical computer.

The high-availability features in VMware vSphere 4.1 along with Distributed Resource Scheduler (DRS) and Storage vMotion enable seamless migration of virtual desktops from one ESX server to another with minimal or no impact to customers usage.

vCenter Server cluster

The following figure shows the cluster configuration from vCenter Server.



The Infrastructure cluster holds the following virtual machines:

- Windows 2008 R2 domain controller — provides DNS, Active Directory, and DHCP services
- Windows 2008 R2 SQL 2008 — provides databases for vCenter Server, View Composer, and other services in the environment
- Windows 2003 R2 View 4.5
- Windows 2008 R2 vCenter Server — provides management services for the VMware clusters and View Composer
- Windows 7 Key Management Service (KMS) — provides a method to activate Windows 7 desktops

The View 4.5 cluster (called View45 in the diagram) consists of 500 virtual desktops.

Windows infrastructure

Introduction Microsoft Windows provides the infrastructure used to support the virtual desktops and includes the following elements:

- Microsoft Active Directory
 - Microsoft SQL Server
 - DNS Server
 - DHCP Server
-

Microsoft Active Directory The Windows domain controller runs the Active Directory service that provides the framework to manage and support the virtual desktop environment. The Active Directory has several functions such as the following:

- Manage the identities of users and their information
 - Apply group policy objects
 - Deploy software and updates
-

Microsoft SQL Server Microsoft SQL Server is a relational database management system (RDBMS). SQL Server 2008 is used to provide the required databases to vCenter Server and View Composer.

DNS Server DNS is the backbone of Active Directory and provides the primary name resolution mechanism of Windows servers and clients. In this solution, the DNS role is enabled on the domain controller.

DHCP Server The DHCP Server provides the IP address, DNS Server name, gateway address, and other information to the virtual desktops. In this solution, the DHCP role is enabled on the domain controller.

Chapter 3: Storage Design

Overview

Introduction The storage design described in this chapter applies to the specific components of this solution.

Contents This chapter contains the following topics:

Topic	See Page
CLARiiON storage architecture	18
Celerra storage configuration	20

CLARiiON storage architecture

Introduction

This solution uses both block-based and file-based storage to leverage the benefits that each of the following provides:

- Block-based storage over Fibre Channel (FC) is used to store the VMDK files for all virtual desktops. This has the following benefits:
 - Block storage leverages the VAAI APIs (introduced in vSphere 4.1) that includes a hardware-accelerated copy to improve the performance and for granular locking of the VMFS to increase scaling.
 - PowerPath Virtual Edition (PP/VE) allows better performance and scalability over the native multipathing options.
- File-based storage is provided by a CIFS export from Celerra NS-120. This has the following benefits:
 - Redirection of user data and roaming profiles to a central location for easy backup and administration
 - Single instancing and compression of unstructured user data to provide the highest storage utilization and efficiency.

This section covers the configuration of the block-based storage to be provided over FC to the ESX cluster to store the VMDK images.

Storage layout

The following diagram shows the storage layout of the disks.

Storage Pool Layout - View 4.5 - Release 30															
Slot	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Bus0 Enc0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 5 SYSTEM 146 GB FC RG - 0	Free	Free	Free	Free	Free	RAID 1/0 Replica 100 GB RG - 3	RAID 1/0 Replica 100 GB RG - 3	RAID 1 FAST C 100 GB EFD	RAID 1 FAST C 100 GB EFD	HOT SP EFD RG - 201
	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	HOT SP FC RG - 202
Bus0 Enc2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Link Cln 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	HOT SP SATA Pool1 RG - 203

RAID protection

The observed workload from Windows 7 was characterized by small random I/Os and showed a read to write ratio of 45:55. Workloads with a write ratio of over 30 percent are considered write heavy. For these types of workloads, EMC recommends using RAID 1/0.

Storage layout overview

The following storage configuration is used in the solution. Celerra NS-120 has a single back-end bus and all the drives are on bus 0. Therefore, the disk numbers are given in ENCLOSURE_DISK format.

- FC disks (0_0 to 0_4) are system LUNs for both CLARiiON and Celerra. During the installation of a Celerra system, the free space on these drives is allocated to a storage pool.
- Disks 0_14, 1_14, and 2_14 are hot spares. These disks are denoted in yellow in the storage layout diagram.
- EFDs (0_10 and 0_11) on the RAID 1/0 group are used to store the linked clone replicas. The EFDs are denoted in purple.

- EFDs (0_12 and 0_13) are used for EMC FAST Cache. There are no user-configurable LUNs on these drives. These EFDs are denoted in red.
- FC disks (1_0 to 1_13) with 450 GB and 15k rpm and SATA disks (2_8 – 2_13) with 1 TB and 7.2k rpm on the RAID 1/0 pool are used to store linked clones. The storage pool uses FAST with FC and SATA disks to optimize both performance and capacity across the pool. FAST Cache is enabled for the entire pool. These disks are denoted in blue. Four LUNs of 1.25 TB each are carved out of the pool and presented to the ESX servers.
- SATA disks (2_0 to 2_7) with 1 TB and 7.2k rpm on the RAID 1/0 group are used to store user data and roaming profiles. These disks are denoted in green. Two file systems are created on two LUNs, one for profiles and the other for user data.

EMC FAST Cache

EMC FAST Cache is a new feature introduced in FLARE release 30. This enables EFDs to be used as an expanded cache layer for the array. Celerra NS-120 is configured with two 100 GB EFDs in a RAID 1 configuration for a 91 GB read/write capable cache. Larger arrays support FAST Cache sizes up to 2 TB.

FAST Cache is an array-wide feature available for both NAS and FC storage. FAST Cache works by examining 64 KB chunks of data in FAST Cache enabled objects on the array. Frequently accessed data is copied to the FAST Cache and subsequent accesses to that data chunk are serviced by FAST Cache. This allows immediate promotion of very active data to EFDs. This dramatically improves the response time for very active data and reduces the data hot spots that can occur within the LUN.

The FAST Cache is an extended read/write cache that can absorb read-heavy activities such as boot storms and antivirus scans, and write-heavy workloads such as operating system patches and application updates.

EMC FAST

FAST is a pool-based feature of FLARE release 30 available for CLARiiON LUNs that automatically migrates data to different storage tiers based on performance requirements of the data.

The pool1_x LUNs are built on a storage pool configured with a mix of FC and SATA drives. Initially, the linked clones are placed on the FC tier. The data created by the linked clones that is not frequently accessed is automatically migrated to the SATA storage tier. This releases space in the faster FC tier for more active data.

EMC PowerPath Virtual Edition

Each datastore that is used to store VMDK files is placed on the FC storage. PowerPath Virtual Edition (PP/VE) is enabled for all FC-based LUNs to efficiently use all the available paths to storage and to minimize the effect of micro-bursting I/O patterns.

vCenter Server storage layout

The datastore configuration in vCenter Server is as follows:

- Pool 1_1 through pool 1_4 — Each of the 1.25 TB datastores accommodates 125 users. This allows each desktop to grow to a maximum size of 10 GB. The pool of desktops created in View Manager is balanced across all these datastores.
- efd_replica and efd_replica2 — These datastores are on two 100 GB EFDs with RAID 1. The input/output to these LUNs is strictly read-only except during operations that require copying a new replica into the datastore.

Celerra storage architecture

Introduction

The EMC Celerra unified storage platform is a dedicated network server optimized for file and block access, delivering high-end features in a scalable and easy-to-use package. For high scalability, Celerra unified storage platforms leverage both the innovative EMC CLARiiON Fibre Channel RAID storage, delivering best-in-class availability and data protection, and the availability, performance, and ease of management of EMC Celerra.


Celerra unified storage systems deliver a single box block and file solution offering a centralized point of management for distributed environments. This makes it possible to dynamically grow, share, and cost-effectively manage multi-protocol file systems and provide multi-protocol block access. Administrators can take advantage of simultaneous support for NFS and CIFS protocols by allowing Windows and Linux/UNIX clients to share files using the Celerra system's sophisticated file-locking mechanisms and by leveraging iSCSI or FC for high-bandwidth or latency-sensitive applications.

EMC Celerra provides five 9s (99.999 percent) availability through advanced failover, high availability, and fault-tolerant networking options. The reliability and speed of backup and recovery in traditional backup application environments can be improved by using Celerra unified storage platforms as a backup-to-disk target.

Celerra file systems

Virtual desktops use two Celerra file systems, one for user profiles and the other to redirect user storage. Each file system is exported to the environment through a CIFS share.

The following diagram shows the dvols that are used to build the file systems for user profiles and redirected user storage.



The following table shows the mapping of dvols to the file system:

File system	dvols	Size
profiles_fs	d31	1 TB
userdata1_fs	d32	2 TB

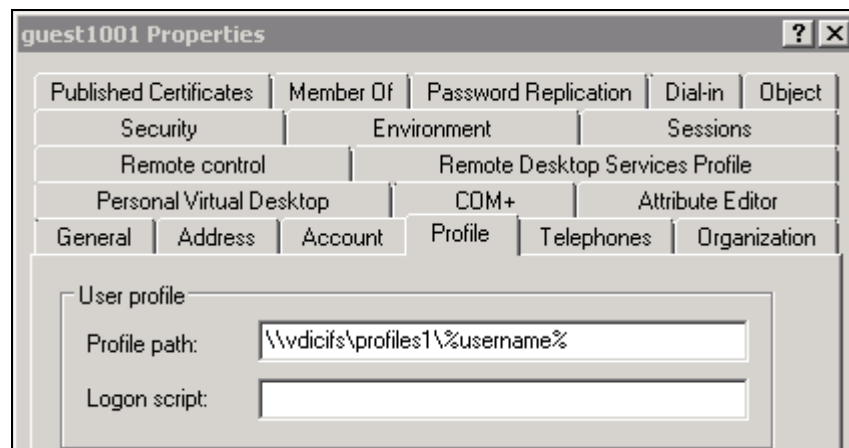
Celerra Home Directory feature

The Celerra Home Directory feature uses the userdata1_fs file system to automatically map the H: drive of each virtual desktop to the users' own dedicated subfolder on the share. This ensures that each user has a dedicated home drive share with exclusive rights to that folder. This export does not need to be created manually. The Home Directory feature automatically maps this for each user.

The Documents folder of the users is also redirected to this share. This allows users to recover the data in the Documents folder by using the Previous Versions functionality if snapshots are applied on the file system. The file system is set at an initial size of 1 TB. However, it can extend itself automatically when more space is required.

Profile export

The profiles_fs file system is used to store user roaming profiles. It is exported through CIFS. The UNC path to the export is configured in Active Directory for roaming profiles as shown in the following figure:



Capacity

The file systems leverage Virtual Provisioning™ and compression to provide flexibility and increased storage efficiency. If single instancing and compression are enabled, unstructured data such as user documents typically leads to a 50 percent reduction in consumed storage.

The Celerra file systems for user profiles and documents are configured as follows:

- profiles_fs is configured to consume 1 TB of space. Assuming 50 percent space savings, each profile can grow up to 2 GB in size. The file system can be extended if more space is needed.
- Userdata_fs is configured to consume 2 TB of space. Assuming 50 percent space savings, each user will be able to store 8 GB of data. The file system can be extended if more space is needed.

Chapter 4: Network Design

Overview

Introduction This chapter describes the network design used in this solution.

Contents This chapter contains the following topics:

Topic	See Page
Considerations	23
Celerra network configuration	24
ESX network configuration	25
Cisco 6509 configuration	26
Fibre Channel network configuration	27

Considerations

Physical design considerations

EMC recommends that switches support gigabit Ethernet (GbE) connections and Link Aggregation Control Protocol (LACP), and the ports on switches support copper-based media.

Logical design considerations

This validated solution uses virtual local area networks (VLANs) to segregate network traffic of various types to improve throughput, manageability, application separation, high availability, and security.

The IP scheme for the virtual desktop network must be designed such that there are enough IP addresses in one or more subnets for the DHCP Server to assign them to each virtual desktop.

Link aggregation

Celerra unified storage provides network high availability or redundancy by using link aggregation. This is one of the methods used to address the problem of link or switch failure.

A link aggregation is a high-availability feature that enables multiple active Ethernet connections to appear as a single link with a single MAC address and potentially multiple IP addresses.

In this solution, LACP is configured on Celerra, which combines two GbE ports into a single virtual device. If a link is lost in the Ethernet port, the link fails over to another port. All traffic is distributed across the active links.

Celerra network configuration

Data Mover ports

The Celerra NS-120 consists of two blades. These blades can be configured in an active/active or active/passive configuration. In the active/passive configuration, the passive blade serves as a failover device for the active blade. In this solution, the blades operate in the active/passive mode.

The Celerra NS-120 blade consists of four GbE controller ports. Ports cge0 and cge1 are configured by using LACP to support virtual machine traffic, home folder access, and external access for roaming profiles. Ports cge2 and cge3 are left free for further expansion.

The external_interface device is used for administrative purposes to move data in and out of the private network on VLAN 274. Both interfaces exist on the LACP1 device configured on cge0 and cge1.

The ports are configured as follows:

```
external_interface protocol=IP device=lacp1
    inet=10.6.121.55 netmask=255.255.255.0
broadcast=10.6.121.255
    UP, Ethernet, mtu=1500, vlan=521,
macaddr=0:60:16:26:19:0
lacp1_int protocol=IP device=lacp1
    inet=192.168.80.5 netmask=255.255.240.0
broadcast=192.168.95.255
    UP, Ethernet, mtu=9000, vlan=274,
macaddr=0:60:16:26:19:0
```

LACP configuration on the Data Mover

To configure the link aggregation that uses cge0 and cge1 on server_2, type the following at the command prompt:

```
$ server_sysconfig server_2 -virtual -name <Device Name> -
create trk
-option "device=cge0,cge1 protocol=lacp"
```

To verify if the ports are channeled correctly, type:

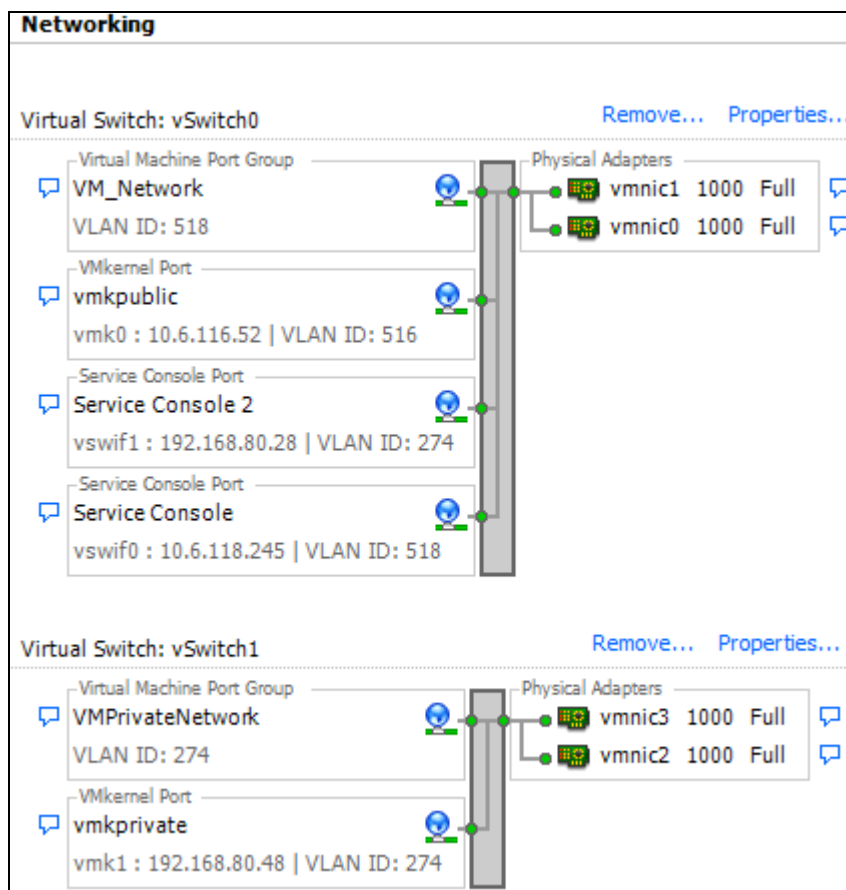
```
$ server_sysconfig server_2 -virtual -info lacp1
server_2:
*** Trunk lacp1: Link is Up ***
*** Trunk lacp1: Timeout is Short ***
*** Trunk lacp1: Statistical Load C is IP ***
Device      Local Grp   Remote Grp Link  LACP Duplex Speed
-----
cge0        10000       5888      Up    Up    Full   1000 Mbs
cge1        10000       5888      Up    Up    Full   1000 Mbs
```

The remote group number must match for both cge ports and the LACP status must be "Up." Verify if the appropriate speed and duplex are established as expected.

ESX network configuration

ESX NIC teaming

All network interfaces in this solution use 1 GbE connections. The Dell R710 servers use four on-board Broadcom GbE Controllers for all network connections. The following diagram shows the vSwitch configuration in vCenter Server.



The following table lists the configured port groups in vSwitch0 and vSwitch1.

Virtual switch	Configured port groups	Used for
vSwitch0	VM_Network	External access for administrative virtual machines
vSwitch0	Vmkpublic	Mounting NFS datastores on the public network for OS installation and patch installs
vSwitch0	Service Console 2	Private network administration traffic
vSwitch0	Service Console	Public network administration traffic
vSwitch1	VMPrivateNetwork	Network connection for virtual desktops, LAN traffic
vSwitch1	Vmkprivate	Mounting multiprotocol exports from the Celerra system on the private VLAN for administrative purposes

Cisco 6509 configuration

Overview The nine-slot Cisco Catalyst 6509-E switch provides high port densities that are ideal for many wiring closet, distribution, and core network deployments as well as data center deployments.

Cabling In this solution, the ESX server and Celerra Data Mover cabling are evenly spread across two WS-x6748 1 Gb line cards to provide redundancy and load balancing of the network traffic.

Server uplinks The server uplinks to the switch are configured in a port channel group to increase the utilization of server network resources and to provide redundancy. The vSwitches have been configured to load balance network traffic on the originating port ID.

The following is an example of the configuration for one of the server ports:

```
description 8/10 9048-43 rtpsol189-1
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 274,516-527
switchport mode trunk
no ip address
spanning-tree portfast trunk
```

Data Movers The network ports for each NS-120 Data Mover are connected to the 6509-E switch and are configured so that server_2 has interfaces cge0 and cge1 split between the two 1 Gb line cards and are configured with LACP, which provides redundancy in case of a NIC or port failure.

The following is an example of the switch configuration for one of the Data Mover ports:

```
description 7/4 9047-4 rtpsol22-dm2.0
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 274,516-527
switchport mode trunk
mtu 9216
no ip address
spanning-tree portfast trunk
channel-group 23 mode active
```

Fibre Channel network configuration

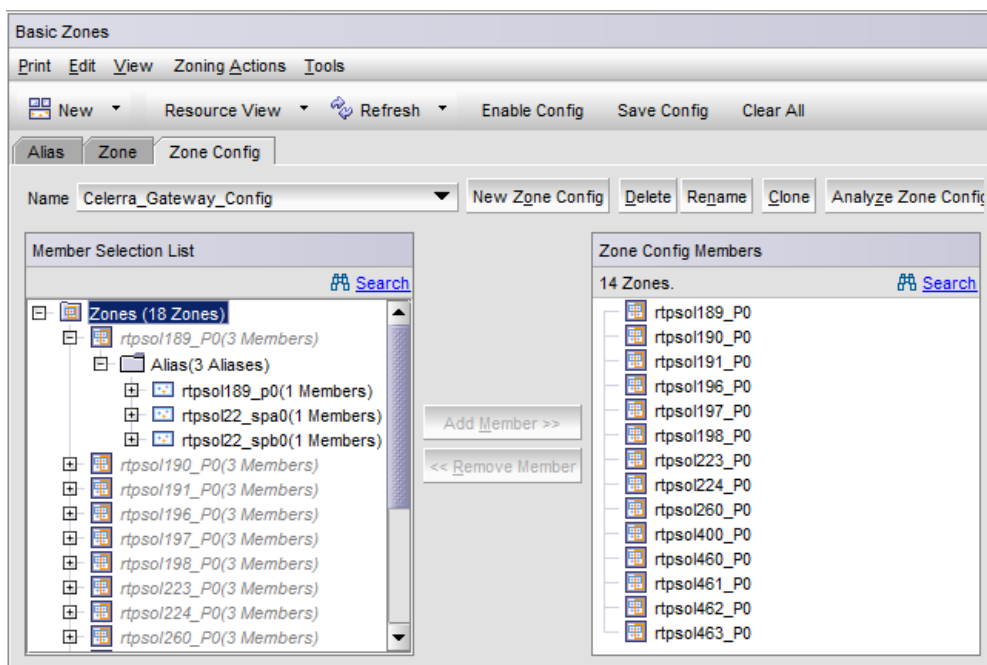
Introduction

Two Brocade DS5100 series FC switches are used to provide the storage network for this solution. The switches are configured in a SAN A/SAN B configuration to provide a fully redundant fabric.

Each server has a single connection to each fabric to provide load balancing and failover capabilities. Each storage processor has two links to the SAN fabrics for a total of four available front-end ports. The zoning is configured so that each server has four available paths to the storage array.

Zone configuration

Single initiator, multiple target zoning is used in this solution. Each server initiator is zoned to two storage targets on the array. The zone configuration for the SAN A fabric is shown in the following figure.



Chapter 5: Installation and Configuration

Overview

Introduction

The installation and configuration steps for the following components are available on the VMware website (www.vmware.com):

- VMware View Connection Server
- VMware View Composer 2.5
- VMware ESX 4.1
- VMware vSphere 4.1

This chapter provides an overview of the configuration of the following components:

- Desktop pools
- Storage pools
- FAST Cache
- Auto-tiering (FAST)
- Celerra Home Directory
- PowerPath Virtual Edition (PP/VE)

The following components are not covered:

- Microsoft System Center Configuration Manager (SCCM)
- Microsoft Active Directory, DNS, and DHCP
- Install and configuration of vSphere components
- Microsoft SQL Server 2008 R2

Contents

This chapter contains the following topic:

Topic	See Page
VMware components	29
Storage components	32

VMware components

VMware View installation overview

The *VMware View Installation Guide* available on the VMware website has detailed procedures about installing View Connection Server and View Composer 2.5. There are no special configuration instructions required for this solution.

The *ESXi Installable and vCenter Server Setup Guide* available on the VMware website has detailed procedures about installing vCenter Server and ESXi and is not covered in further detail in this paper. There are no special configuration instructions required for this solution.

View setup

Before deploying desktop pools, ensure that the following steps from the *VMware View Install Guide* have been completed:

- Prepare Active Directory
 - Install View Composer 2.5 on vCenter Server
 - Install View Connection Server
 - Add a vCenter Server instance to View Manager
-

View desktop pool configuration

VMware recommends using a maximum of 250 desktops per replica image, which requires creating a unique pool for every 250 desktops. In this solution, persistent automated desktop pools were used.

To create a persistent automated desktop pool as configured for this solution, complete the following steps:

1. Log in to the **VMware View Administration** page, which is located at <https://server/admin>, where “server” is the IP address or DNS name of the View Manager server.
2. Click the **Pools** link in the left pane.
3. Click the **Add** button under the **Pools** banner.
4. In the **Type** page, select **Automated Pool** and click **Next**.
5. In the **User Assignment** page, select the **Dedicated** radio button and ensure that the **Enable automatic assignment** checkbox is selected. Click **Next**.
6. In the **vCenter Server** page, select **View Composer linked clones** and select a vCenter Server that supports View Composer, as shown in the following figure. Click **Next**.

Add Pool

Pool Definition

- Type
- User Assignment
- vCenter Server**

Setting

- Pool Identification
- Pool Settings
- View Composer Disks
- Provisioning Settings
- vCenter Settings**
- Guest Customization
- Ready to Complete

vCenter Server

☐ Full virtual machines

☒ View Composer linked clones

vCenter Server	View Composer
viewvc.vdi.local(vd\administrator)	Yes
viewvc.vdi.local(viewmanager)	Yes

View Composer

View Composer linked clones share the same base image and use less storage space than full virtual machines.

The user profile for linked clones can be redirected to persistent disks that will be unaffected by OS updates and refreshes.

Supported Features

- ✓ Local Mode
- ✓ PCoIP
- ✓ Storage savings
- ✓ Recompose and refresh
- ✓ QuickPrep guest customization
- ✓ Sysprep guest customization

< Back Next > Cancel

- In the **Pool Identification** page, enter the required information and click **Next**.
- In the **Pool Settings** page, make any required changes and click **Next**.
- In the **View Composer Disk** page, select the **Do not redirect Windows profile** radio button and click **Next**.
- In the **Provisioning Settings** page, select a name for the desktop pool and enter the number of desktops to provision, as shown in the following figure. Click **Next**.

Add Pool

Pool Definition

- Type
- User Assignment
- vCenter Server

Setting

- Pool Identification
- Pool Settings
- View Composer Disks
- Provisioning Settings
- vCenter Settings**
- Guest Customization
- Ready to Complete

Provisioning Settings

Basic

☒ Enable provisioning

☒ Stop provisioning on error

Virtual Machine Naming

☐ Specify names manually

0 names entered Enter names...

☐ Start desktops in maintenance mode

Unassigned desktops kept powered on: 1

☒ Use a naming pattern

Naming Pattern: view45pool

Pool Sizing

Max number of desktops: 250

Number of spare (powered on) desktops: 250

☐ Provision desktops on demand

Min number of desktops: 1

☒ Provision all desktops up-front

Naming Pattern

Virtual machines will be named according to the specified naming pattern. By default, View Manager appends a unique number to the specified pattern to provide a unique name for each virtual machine.

To place this unique number elsewhere in the pattern, use '{n}'. (For example: vm-{n}-sales.). See the help for more naming pattern syntax options.

< Back Next > Cancel

- In the **vCenter Settings** page, browse to select a default image and folder for the virtual machines, the cluster hosting the virtual desktops, the resource

pool to hold the desktops, and the datastores that will be used to deploy the desktops, as shown in the following figure. Click **Next**.

Add Pool

Pool Definition

- Type
- User Assignment
- vCenter Server

Setting

- Pool Identification
- Pool Settings
- View Composer Disks
- Provisioning Settings
- vCenter Settings**
- Guest Customization
- Ready to Complete

vCenter Settings

Virtual Machine Settings

Default image: windows7vsi - finalworkingsnap Browse...

VM folder: /VDI/vm/vsi Browse...

Resource Settings

Host or cluster: /VDI/host/View45 Browse...

Resource pool: /VDI/host/View45/Resources/vsi Browse...

Datastores:

Name	Capacity (...	Free (GB)	Use For
efd_replica	91.5	80.13	Replica disk
pool1_1	1279.75	1147.87	Linked clone
pool1_2	1279.75	1147.83	Linked clone
pool1_3	1279.75	1146.7	Linked clone
pool1_4	1279.75	1146.03	Linked clone

< Back Next > Cancel

- In the **Select Datastores** page, select **Use different datastore for View Composer replica disks** and select the datastores for replica and linked clone images, as shown in the following figure.

Select Datastores

Select the datastores to use for this pool. Only datastores that can be used by the selected host or cluster can be selected.

☐ Show incompatible datastores Local datastore ☒ Shared datastore ☒

	Datastore	Capacity (GB)	Free (GB)	Type	Use For	Storage Overcommit
<input checked="" type="checkbox"/>	efd_replica	91.5	80.13	VMFS	Replica disks	Conservative
<input type="checkbox"/>	fc_replica	99.75	88.37	VMFS		
<input checked="" type="checkbox"/>	pool1_1	1279.75	1147.87	VMFS	Linked clones	Conservative
<input checked="" type="checkbox"/>	pool1_2	1279.75	1147.83	VMFS	Linked clones	Conservative
<input checked="" type="checkbox"/>	pool1_3	1279.75	1146.7	VMFS	Linked clones	Conservative

☐ Use different datastores for OS disks and View Composer persistent disks

☒ Use different datastore for View Composer replica disks

Data Type	Selected Free Space (GB)	Min Recommended (GB)	50% utilization (G)	Max Recommended (G)
Linked clones	4,588.43	375.00	1,687.50	3,187.50
Replica disks	80.13	24.00		24.00

OK Cancel

- In the **Guest Customization** page, select the **domain** and **AD container**, and then select the **Use QuickPrep** radio button. Click **Next**.
- In the **Ready to Complete** page, verify the settings for the pool, and then click the **Finish** button to start the deployment of the virtual desktops.

**PowerPath
Virtual Edition**

PowerPath/VE 5.4.1 supports ESX 4.1. Support for ESX4i is available through RPQ. An update to PowerPath/VE later this year will add support to ESX4i without any caveats.

The *EMC PowerPath/VE for VMware vSphere Installation and Administration Guide* available on Powerlink provides the procedure to install and configure PowerPath/VE. There are no special configuration instructions required for this solution.

The PowerPath/VE binaries and support documentation are available on Powerlink at the following location:
Home > Support > Software Downloads and Licensing > Downloads P-R > PowerPath for VMware

Storage components

Storage pools

Storage pools in FLARE release 30 support heterogeneous drive pools. In this solution, a 20-disk pool was configured from 14 FC disks and 6 SATA drives. Four thick LUNs, each 1.25 TB in size, were created from this storage pool, as shown in the following figure. FAST Cache is enabled for the pool.

Pools

Name	FAST Cache	RAID Type	Drive Type	User Capacity (GB)	Free Capacity (...)	Auto-Tiering Status
Infrastructure ...	Off	RAID5	FC	1070.065	25.002	Manual
Pool 1	Off	RAID1/0	Mixed	5548.339	324.020	Scheduled

1 Selected Create Delete Properties Expand 2 items

Last Refreshed: 2010-08-17 16:48:28

Details

Pool LUNs Disks

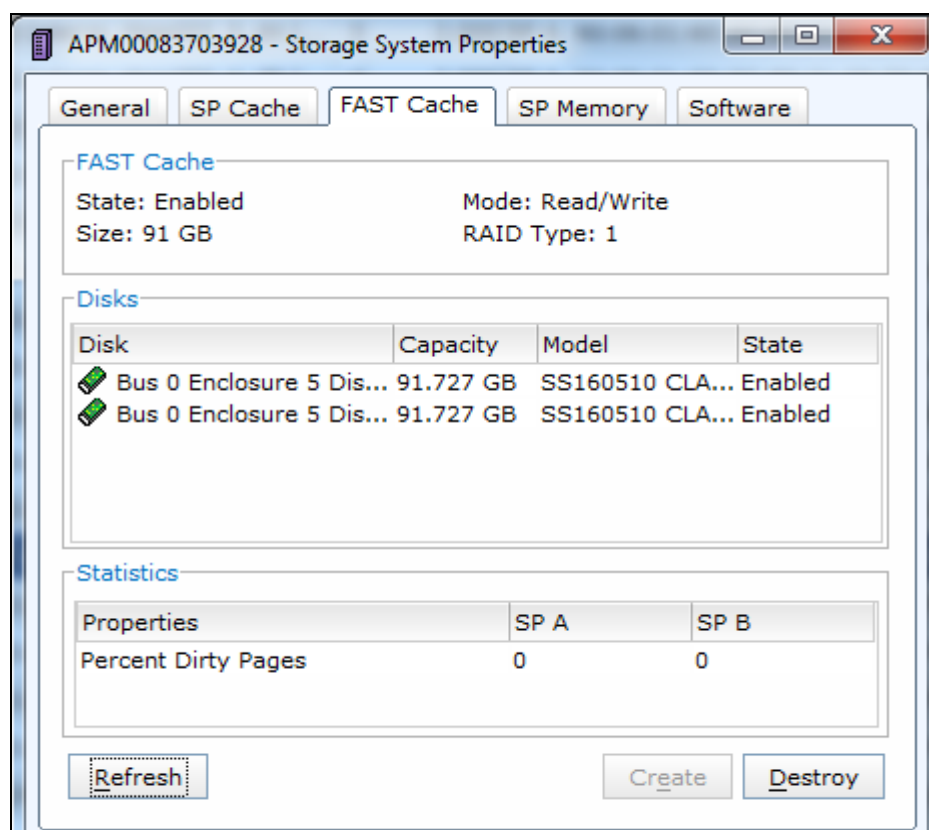
Filter for Usage: ALL User LUNs

Name	ID	User Capacity (GB)	Initial Tier	Tiering Policy
LUN 36 - Pool1_1	36	1280.000	Highest Available Tier	Highest Available Tier
LUN 35 - Pool1_2	35	1280.000	Highest Available Tier	Highest Available Tier
LUN 34 - Pool1_3	34	1280.000	Highest Available Tier	Highest Available Tier
LUN 33 - Pool1_4	33	1280.000	Highest Available Tier	Highest Available Tier

For each LUN in the storage pool, the tiering policy is set to **Highest Available Tier** to ensure that all frequently accessed desktop data remains on the FC disks. As data ages and is used infrequently, it is moved to the SATA drives in the pool.

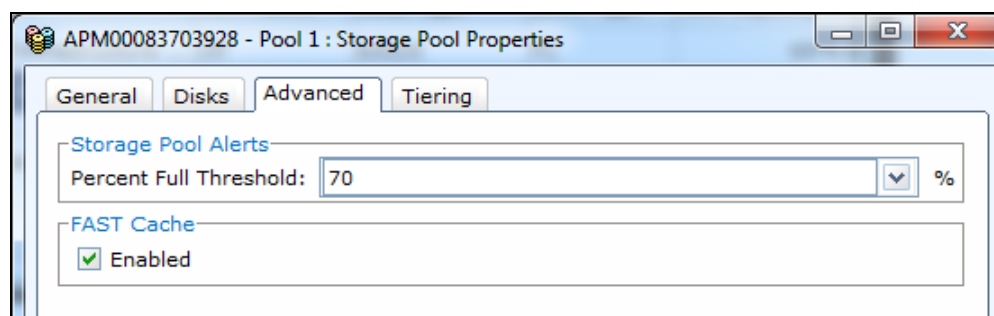
Enable FAST Cache

FAST Cache is enabled as an array-wide feature in the system properties of the array in Unisphere™. Click the **FAST Cache** tab, click the **Create** button, and select the eligible EFDs to create the FAST Cache. There are no user-configurable parameters for the FAST Cache.



FAST Cache is not enabled for the replica storage in this solution. The replica images are serviced from the EFDs. Enabling FAST Cache for these LUNs causes additional overhead without added performance.

If the replica images are stored on FC disks, enable FAST Cache for those LUNs. To enable FAST Cache for any LUNs in a pool, go to the properties of the pool and click the **Advanced** tab. Select **Enabled** to enable FAST Cache, as shown in the following figure.



Configure FAST

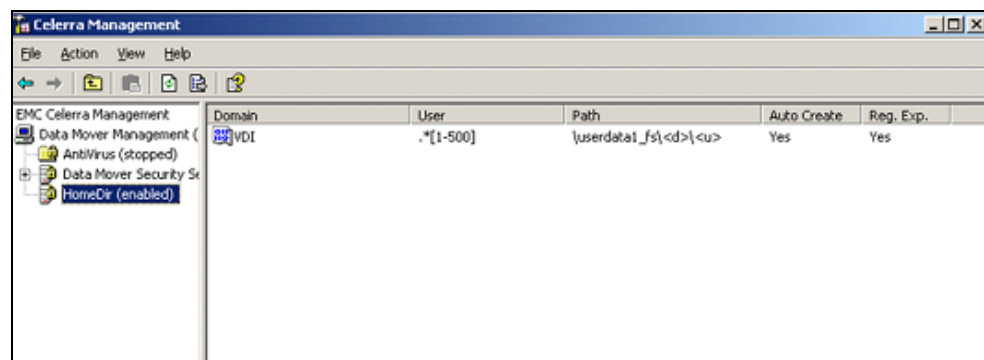
To configure the FAST feature for a pool LUN, go to the properties for a pool LUN and click the **Tiering** tab. From this location, set the tiering policy for the LUN.

Celerra Home Directory feature

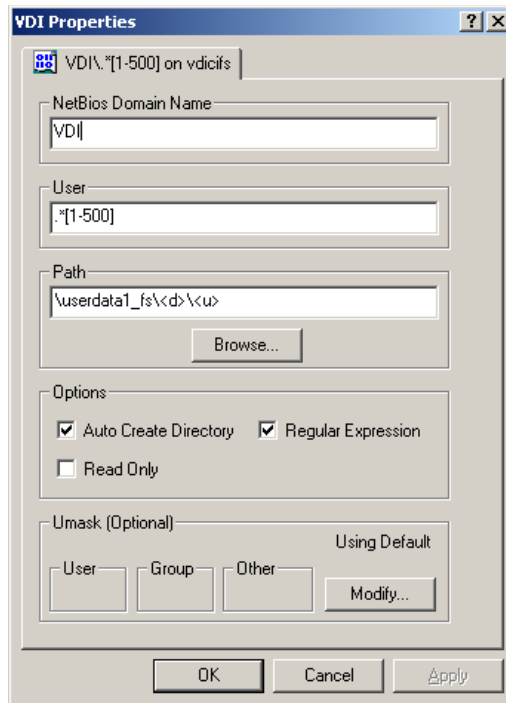
The Celerra Home Directory installer is available on the NAS Tools and Apps CD for each DART release and can be downloaded from Powerlink at the following location: Home > Support > Software Downloads and Licensing > Downloads C > Celerra Software.

Instructions to install the Celerra Home Directory feature are located on the EMC Celerra Network Server Documentation CD available on Powerlink.

After the feature is installed, use the Celerra Management Microsoft Management Console (MMC) snap-in to configure the Home Directory feature. A sample configuration is shown in the following two figures.



For any user account that ends with a suffix between 1 and 500, the sample configuration shown in the following figure automatically creates a user home directory on the \userdata1_fs file system in the format \userdata1_fs\<domain>\<user> and maps the H:\ drive to this path. Each user has exclusive rights to the folder.



The *EMC Celerra Home Directory — A Detailed Review* white paper available on Powerlink provides more details and advanced configuration examples.

Chapter 6: Testing and Validation

Overview

Introduction This chapter describes the tests that were run to validate the configuration of the solution.

Contents This chapter contains the following topics:

Topic	See Page
Testing overview	37
Boot storm results	40
View refresh results	46
View recompose results	52
Antivirus results	58
Patch install results	65
LoginVSI results	71

Testing overview

Introduction

This chapter provides a summary and characterization of the tests performed to validate the solution. The goal of the testing was to characterize the performance of the solution and component subsystems during the following scenarios:

- Boot storm of all desktops
- View desktop refresh of all desktops
- View recompose of all desktops
- McAfee full scan on all desktops
- Security patch install with Microsoft SCCM
- User workload testing using LoginVSI

The steps used to configure McAfee and SCCM are beyond the scope of this document.

Observed user workload

Using LoginVSI, a task worker benchmark was run with Windows 7 virtual desktops. The following table shows the observed workload that was generated. The results were used as a guide to design the reference architecture of this solution.

Windows 7 workload							
	Committed bytes	Read IOPS	Total IOPS	Write IOPS	Active RAM (MB)	% Processor time	Network bytes/sec
Avg	522349163.5	3.9	8.9	5.3	264.3	7.5	75551.1
95th	589459456.0	4.0	37.0	26.4	453.0	36.6	145559.2
Max	599506944.0	577.0	875.0	875.0	460.0	109.3	5044232.8

Sizing options

There are two traditional ways of sizing the I/O requirements: Average IOPS and 95th percentile IOPS. The following table shows the number of disks that are required to meet the IOPS requirements by sizing for the average and the 95th percentile IOPS.

Windows 7 disk requirements					
Avg IOPS	No. of users	Total IOPS	Read: Write Mix	IOPS	FC disks required
9	500	4,500	45:55	Read: 2000	10
				Write: 2500	13
95th IOPS	No. of users	Total IOPS	Read: Write Mix	IOPS	FC disks required
37	500	18,500	45:55	Read: 8325	42
				Write: 10125	50

Average versus 95th percentile workload

Sizing on the average IOPS can provide good performance for the virtual desktops in steady state. However, this leaves insufficient headroom in the array to absorb high peaks in I/O without causing very high response times for the clients. As a result, the desktop performance will suffer during boot storms, desktop recompose or refresh tasks, antivirus DAT updates, and similar events. Change management becomes the most important focus of the View administrator because all tasks must be carefully balanced across the desktops to avoid I/O storms.

To address the issue of I/O storms, the disk I/O requirements can be sized based on the 95th percentile load. Sizing to the 95th percentile ensures that 95 percent of all the values measured for IOPS fall below that value. Sizing by this method ensures great performance in all scenarios except during the most demanding of mass I/O events. However, the disadvantage of this method is cost. In this example, it takes 92 disks to satisfy the I/O requirements instead of 23 disks. This leads to high capital and operational costs.

Comparison

Based on the observed workload, two configurations were built and tested for comparison. The following figure shows the disk layout that was used based on the average IOPS requirement. It will be referred to as the baseline configuration.

Baseline - Storage Pool Layout - View 4.5 - Release 30															
Slot	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Bus0 Enc0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 1/0 Replica 450GB FC RG - 1	RAID 1/0 Replica 450GB FC RG - 1	RAID 1/0 Replica 450GB FC RG - 1	RAID 1/0 Replica 450GB FC RG - 1	RAID 1/0 Replica 450GB FC RG - 1	RAID 1/0 Replica 450GB FC RG - 1	RAID 1/0 Replica 450GB FC RG - 1	RAID 1/0 Replica 450GB FC RG - 1	RAID 1/0 Replica 450GB FC RG - 1	RAID 1/0 Replica 450GB FC RG - 1
	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	HOT SP FC RG - 200
Bus0 Enc2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	HOT SP SATA RG - 201

The configuration shown in the reference architecture takes this baseline architecture and replaces the 10 FC disks for the replica with two 100 GB EFDs and adds two more 100 GB EFDs for FAST Cache, as shown in the following figure.

Reference Architecture - Storage Pool Layout - View 4.5 - Release 30															
Slot	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Bus0 Enc0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 5 SYSTEM 146 GB FC RG - 0	RAID 5 SYSTEM 146 GB FC RG - 0	Free	Free	Free	Free	Free	RAID 1/0 Replica 100 GB EFD	RAID 1/0 Replica 100 GB EFD	RAID 1 FAST C 100 GB EFD	RAID 1 FAST C 100 GB EFD	HOT SP EFD RG - 201
	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	RAID 1/0 Link Cln 450 GB FC Pool1	HOT SP FC RG - 202
Bus0 Enc2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA RG - 2	RAID 1/0 Usr Data 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	RAID 1/0 Link Cln 1TB SATA Pool1	HOT SP SATA RG - 203

These two architectures are compared against each other to show how the combination of EFDs and FAST Cache can provide much better performance with fewer disks as compared to the baseline configuration. It will also be shown that the reference architecture can achieve the performance required to satisfy the 95th percentile load requirements.

Use cases This solution ran six common use cases to validate whether the solution performed as expected under heavy load situations. The use cases tested are:

- Simultaneous boot of all desktops
- View refresh operation on all desktops
- View recompose operation on all desktops
- Full antivirus scan of all virtual machines
- Installation of five security updates using SCCM on all desktops
- Login and steady state user load simulated using the LoginVSI medium workload

In each use case, a number of key metrics are presented showing the overall performance of the solution. Additionally, for each use case, a set of comparison graphs are presented to show the impact of FAST Cache on the pool LUNs holding the View linked clone images.

LoginVSI To run a user load against the desktops, the Virtual Session Index (VSI) tool was used. VSI provided guidance to gauge the maximum number of users a desktop environment can support. The LoginVSI workload can be categorized as light, medium, heavy, and custom. A medium workload was chosen for testing and had the following characteristics:

- The workload emulated a medium knowledge worker who uses Microsoft Office, Internet Explorer, and PDF.
- After a session had started, the medium workload repeated every 12 minutes.
- The response time was measured every 2 minutes during each loop.
- The medium workload opened up to five applications simultaneously.
- The type rate was 160 ms for each character.
- The medium workload in VSI 2.0 was approximately 35 percent more resource-intensive than VSI 1.0.
- Approximately 2 minutes of idle time were included to simulate real-world users.

Each loop of the medium workload opened and used the following:

- Outlook 2007: Browsed 10 messages.
 - Internet Explorer: One instance was left open (BBC.co.uk), one instance browsed Wired.com, Lonelyplanet.com and a heavy Flash application gettheglass.com (not used with MediumNoFlash workload).
 - Word 2007: One instance to measure the response time and one instance to review and edit the document.
 - Bullzip PDF Printer and Acrobat Reader: The Word document was printed and the PDF was reviewed.
 - Excel 2007: A very large sheet was opened and random operations were performed.
 - PowerPoint 2007: A presentation was reviewed and edited.
 - 7-zip: Using the command line version, the output of the session was zipped.
-

LoginVSI launcher A LoginVSI launcher is a Windows system that launches desktop sessions on target virtual desktop machines. There are two types of launchers — master and slave. There is only one master in a given test bed and there can be as many slave launchers as required.

The number of desktop sessions a launcher can run is typically limited by the CPU or memory resources. Login consultants recommend using a maximum of 45 sessions per launcher with two CPU cores (or two dedicated vCPUs) and 2 GB RAM, when the GDI limit has not been tuned (default). However with the GDI limit tuned, this limit extends to 60 sessions per two-core machine.

In this validated testing, 500 desktop sessions were launched from 12 launcher virtual machines, resulting in approximately 42 sessions established per launcher. Each launcher virtual machine is allocated two vCPUs and 4 GB of RAM. There were no bottlenecks observed on the launchers during the VSI-based tests.

FAST Cache configuration

For all tests, FAST Cache is enabled for the storage pool holding the four Pool1_x datastores. FAST Cache is not enabled for the EFD-based replica image.

Boot storm results


Test methodology

This test was conducted by selecting all the desktops in vCenter Server and selecting **Power On**. The desktops were powered on as soon as the tasks were cleared from vCenter Server.

Overlays are added to the graphs to show when the last power on task was completed and when the I/O to the pool LUNs achieved a steady state. For the boot storm test, all the desktops were powered on within 5 minutes and achieved a steady state approximately 9 minutes later. The total start to finish time was approximately 14 minutes.

EFD replica disk load


The following graph shows the I/O and response time metrics from one of the EFDs in the replica datastore.



Each EFD serviced 5,000 IOPS at peak load, but the response time for the EFDs remained below 1.2 ms, which indicates that the disks were not driven to saturation.

EFD replica LUN load


The following graph shows the I/O and response time metrics from the replica LUN.



The replica LUN serviced nearly 10,000 IOPS during peak load with a maximum response time of less than 2 ms, which indicates that the LUN was not driven to saturation.

Pool individual disk load

The following graph shows the disk I/O and response time for a single FC drive in the storage pool housing the four Pool1_x datastores. Because the statistics from all the drives in the pool were similar, a single drive is reported for the purpose of clarity and readability of the graph.




The number of I/Os serviced by the individual drives in the pool was extremely low. This was because nearly all the I/O requests were serviced by FAST Cache. The disk response time appeared relatively high against the load, but this was misleading.

FAST Cache serviced all the I/Os with a good locality of reference. Therefore, only I/Os that were very far apart made it to the actual disk drives. Because these I/Os caused large disk seeks, the response time appeared high. However, because the I/Os were a small number, they did not adversely affect the performance of the desktops.

Pool LUN load


The following graph shows the LUN I/O and response time from the Pool1_3 datastore. Because the statistics from all the pools are similar, a single pool is reported for the purpose of clarity and readability of the graph.



During peak load, the LUN response time did not exceed 3 ms while the datastore serviced over 1,000 IOPS.


Storage processor I/O

The following graph shows the total I/Os served by the storage processor during the test. The replicas were owned by SPB, which caused all the reads to be handled by SP B. Therefore, the throughput and utilization for SP B were higher than those for SP A.



Storage processor utilization


The following graph shows the storage processor utilization during the test. The replicas were owned by SPB, which caused all the reads to be handled by SP B. Therefore, the throughput and utilization for SP B were higher than those for SP A.



The replica traffic caused high SP utilization during the peak load of the test. For a better distribution of the load, place the replica images on separate storage processors.

FAST Cache I/O


The following graph shows the number of I/Os serviced from FAST Cache during the boot storm test.



At peak load, FAST Cache serviced over 3,700 IOPS from the linked clone datastores, which is the equivalent of 18 FC disks serving 200 IOPS each.

Comparison of pool disk I/O

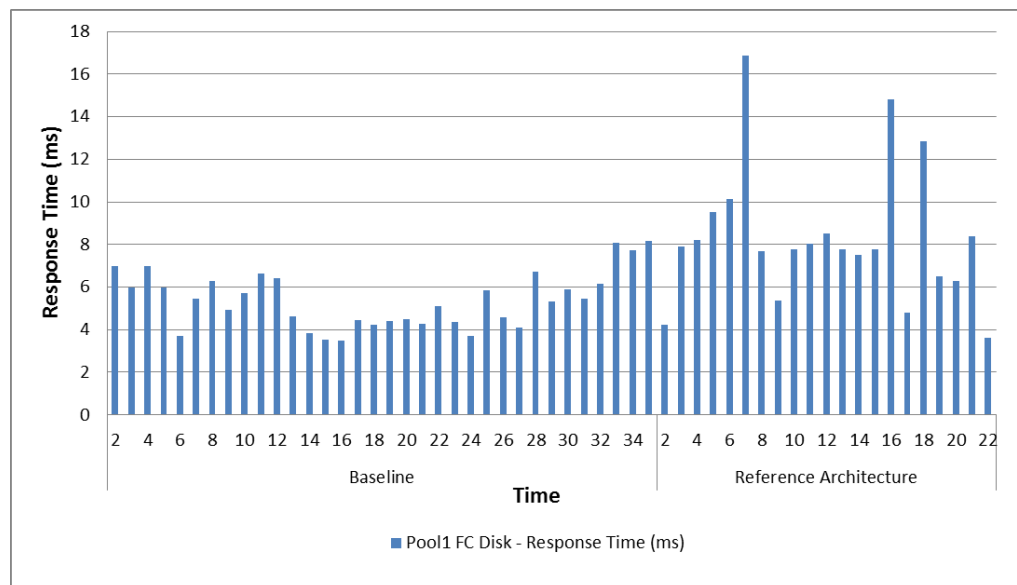
The following graph shows a comparison of the disk I/Os in the reference architecture solution and the baseline configuration.



FAST Cache reduced the peak load on the disks from 232 IOPS to 13 IOPS, nearly an 18 times decrease in the I/O required to be serviced from the pool disks. Additionally, the time to boot all desktops to steady state was decreased from 21 minutes to 13 minutes in the reference architecture solution.

Comparison of pool disk response time

The following graph shows a comparison of the disk response time in the reference architecture solution and the baseline configuration.



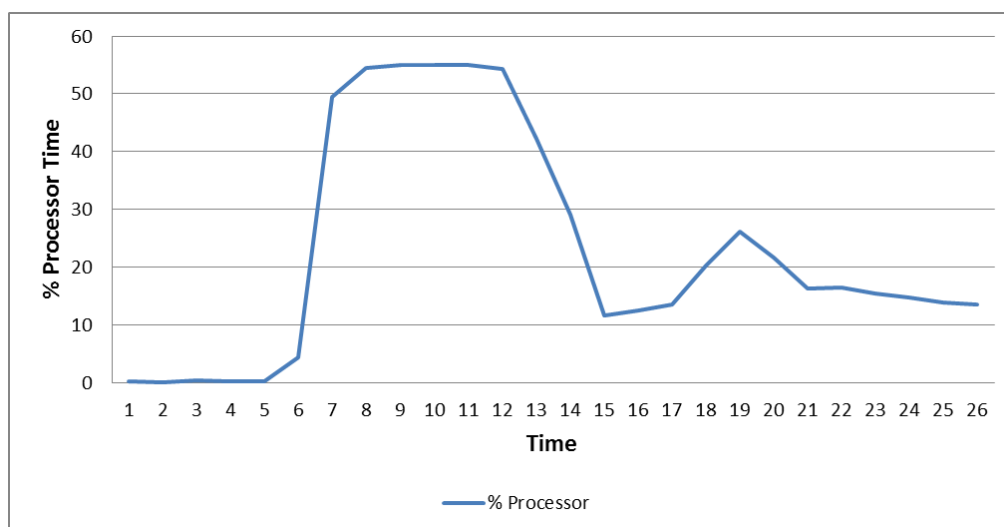
The number of I/Os serviced by the individual drives in the pool was extremely low because nearly all the I/O requests were serviced by FAST Cache. The requests that were serviced by the drives tend to have a very poor locality of reference, which

caused very large disk seeks. These large seek distances caused high response times.

Due to the low number of I/Os, there were very few samples to measure. All the samples had relatively high response times, which caused the high response time values. However, when combined with the I/Os serviced from FAST Cache, which had a very low response time, the overall response time from the client remained very low.

ESX CPU load


The following graph shows the CPU load from the ESX servers in the View cluster. All servers had similar results. Therefore, a single server is reported.



The ESX server briefly achieved CPU utilization of approximately 55 percent during peak load in this test.

ESX disk response time

The following graph shows the Average Guest Millisecond/Command counter, which is shown as GAVG in esxtop. This counter represents the response time for I/Os issued to the storage array.



GAVG for the EFD replica storage and the linked clone storage on the Pool1_x datastores was less than 2.5 ms. This indicates excellent performance under this load.

View refresh results


Test methodology

This test was conducted by selecting a refresh operation for all desktops from the View Manager administration console. No users were logged in during the test. Overlays are added to the graphs to show when the last power on task completed and when the I/O to the pool LUNs achieved a steady state.

For the refresh test, all vCenter Server tasks completed within 83 minutes and achieved a steady state approximately 23 minutes later. The start to finish time was approximately 106 minutes for all desktops.

EFD replica disk load


The following graph shows the I/O and response time metrics from one of the EFDs in the replica datastore.



Each EFD was servicing nearly 3,000 IOPS at peak load, but the response time for the EFDs remained below 10 ms, which indicates that the disks were not driven to saturation.

EFD replica LUN load


The following graph shows the I/O and response time metrics from the replica LUN.



The replica LUN serviced nearly 5,500 IOPS during peak load with a maximum response time of less than 2.2 ms, which indicates that the LUN was not driven to saturation.

Pool individual disk load

The following graph shows the disk I/O and response time for a single FC drive in the storage pool housing the four Pool1_x datastores. Because the statistics for all the drives in the pool were similar, only a single disk is reported.




During the peak I/O load, the disk response time was consistently below 20 ms. This indicates that the disk was not saturated. At the beginning and at the end of the test, I/Os serviced by the individual drives in the pool were extremely low. This was because all I/O requests were serviced by FAST Cache.

FAST Cache serviced all I/Os with a good locality of reference. Only I/Os that were very far apart made it to the actual disk drives. Because these I/Os caused large disk

seeks, the response time appeared high. However, because the I/Os were a small number, they did not adversely affect the performance of the desktops.

Pool LUN load


The following graph shows the LUN I/O and response time from the Pool1_3 datastore. Because the statistics from all the pools are similar, only a single pool is reported for clarity and readability of the graph.



During peak load, the datastore serviced over 2,800 IOPS and the LUN response time remained within 13 ms.


Storage processor I/O

The following graph shows the total I/Os served by the storage processor during the test. The replicas were owned by SPB, which caused all the reads to be handled by SP B. Therefore, the throughput and utilization for SPB were higher than those for SP A.



Storage processor utilization


The following graph shows the storage processor utilization during the test. The replicas were owned by SP B, which caused all reads to be handled by SP B. Therefore, the throughput and utilization for SP B were higher than those for SP A.



The refresh operation caused high CPU utilization during peak load, but headroom was still available.

FAST Cache I/O


The following graph shows the number of I/Os serviced from FAST Cache during the test.



At peak load, FAST Cache serviced over 10,000 IOPS from the linked clone datastores, which is the equivalent of 50 FC disks serving 200 IOPS each.

Comparison of pool disk I/O


The following graph shows a comparison of the disk I/Os in the reference architecture solution and the disk I/Os in the baseline configuration.



FAST Cache reduced the peak load on the disks from 295 IOPS in the baseline configuration to 109 IOPS in the reference architecture solution, which is a decrease of nearly three times in the I/Os required to be serviced from the pool disks.

Comparison of pool disk response time


The following graph shows a comparison of the disk response time in the reference architecture solution and the disk response time in the baseline configuration.



The peak disk response time was reduced by approximately 50 percent in the reference architecture solution when compared to the baseline configuration.

ESX CPU load


The following graph shows the CPU load from the ESX servers in the View cluster. A single server is reported because all servers had similar results.



The CPU load of the ESX server was well within the acceptable limits during this test.

ESX disk response time

The following graph shows the Average Guest Millisecond/Command counter, which is shown as GAVG in esxtop. This counter represents the response time for I/Os issued to the storage array.



GAVG for the EFD replica storage and the linked clone storage on the Pool1_x datastores was well below 10 ms, which indicates a very good performance under this load.

View recompose results

Test methodology

This test was conducted by selecting a recompose operation for all desktops from the View Manager console. No users were logged in during the test.

Overlays are added to the graphs to show when the last power-on task completed and when the I/O to the pool LUNs achieved a steady state.

In the recompose test, the graphs do not show the deletion of the existing desktops and the last stage of the desktop reconfiguration tasks. These events are not significant I/O drivers. Omission of these events enhances the readability of the charts during high I/O periods that occur during desktop creation from the new replica.


The timeline for the test is as follows:

- 1 to 93 minutes — Delete the existing desktops (not shown in the graphs)
- 93 to 98 minutes — Copy the new replica image
- 99 to 189 minutes — Create new desktops
- 190 to 280 minutes — Run the desktop reconfiguration tasks (low I/O not shown in graphs)

In all the graphs, the first highlighted spike of the I/O is the replica copy operation. This lasted approximately 7 minutes.

EFD replica disk load

The following graph shows the I/O and response time metrics from one of the EFDs in the replica datastore.




The copy of the new replica image caused a heavy sequential write workload on the EFDs, thus causing a small spike in the response time, up to a peak of 6.5 ms for the EFDs.

Each EFD was servicing 2,700 IOPS at peak load, but the response time for the EFDs remained below 1 ms. This indicates that the disks were very lightly loaded during this test except for the initial replica copy operation.

EFD replica LUN load


The following graph shows the I/O and response time metrics from the replica LUN.



The replica LUN serviced nearly 5,000 IOPS during peak load with a maximum response time of less than 2 ms. This indicates that the LUN was not driven to saturation during the actual deployment of new desktops.

Pool individual disk load


The following graph shows the disk I/O and response time for a single FC drive in the storage pool housing the four Pool1_x datastores. Because the statistics from all the drives in the pools were similar, only a single drive is reported for clarity and readability of the graph.



Each drive was servicing fewer than 60 IOPS at peak load and the response time for the drives remained below 10 ms. This indicates that the disks were not driven to saturation.

Pool LUN load


The following graph shows the LUN I/O and response time from the Pool1_3 datastore. Because the statistics from all the pools were similar, only a single pool is reported for clarity and readability of the graph.



During peak load, the LUN response time remained within 3.5 ms and the datastore serviced over 2,200 IOPS.


Storage processor I/O

The following graph shows the total I/Os served by the storage processor during the test. The replicas were owned by SP B, which caused all the reads to be handled by SP B. Therefore, the throughput and utilization for SP B were higher than those for SP A.



Storage processor utilization


The following graph shows the storage processor utilization during the test. The replicas were owned by SP B, which caused all the reads to be handled by SP B. Therefore, the throughput and utilization for SP B were higher than those for SP A.



The recompose operation caused high CPU utilization during peak load, but headroom was still available. Running a recompose operation on several virtual desktops will not cause any performance issues because the task will take longer, and will not increase the load when compared to 500 desktops.

FAST Cache I/O


The following graph shows the number of I/Os serviced from FAST Cache during the test.



At peak load, FAST Cache serviced over 5,000 IOPS from the linked clone datastores, which is the equivalent of 25 FC disks serving 200 IOPS each.

Comparison of pool disk I/O

The following graph shows a comparison of the disk I/Os in the reference architecture solution and the disk I/Os in the baseline configuration.




FAST Cache reduced the peak load on the disks from 224 IOPS in the baseline configuration to 59 IOPS in the reference architecture solution, which was a decrease of more than three times in the number of I/Os to be serviced from the pool disks.

The graph shows the period while deleting the current desktops in the baseline. The actual time required to deploy the new desktops was approximately 160 minutes for the baseline configuration compared to 96 minutes for the reference architecture.

Comparison of pool disk response time


The following graph shows a comparison of the disk response time in the reference architecture solution as compared to the baseline configuration.



The disk response time stayed well below 15 ms for the duration of the test in the reference architecture when compared to a peak of 39 ms in the baseline configuration.

ESX CPU load


The following graph shows the CPU load from the ESX servers in the View cluster. A single server is reported because all servers had similar results.



The CPU load on the ESX server was well within acceptable limits during this test.

ESX disk response time

The following graph shows the Average Guest Millisecond/Command counter, which is shown as GAVG in esxtop. This counter represents the response time for I/Os issued to the storage array.



GAVG for both the EFD replica storage and the linked clone storage on the Pool1_x datastores was below 5 ms. This indicates a very good performance under this load.


Antivirus results

Test methodology

This test was conducted by scheduling a full scan of all desktops through a custom script using McAfee 8.7. The full scans were started over the course of an hour on all desktops. The entire test took three hours to complete.

EFD replica disk load


The following graph shows the I/O and response time metrics from one of the EFDs in the replica datastore.



Each EFD was servicing over 2,700 IOPS at peak load, but the response time remained below 2 ms. This indicates that the disks were not driven to saturation.

EFD replica LUN load


The following graph shows the I/O and response time metrics from the replica LUN.



The replica LUN serviced nearly 5,300 IOPS during peak load with a maximum response time of less than 3.25 ms.

Pool individual disk load

The following graph shows the disk I/O and response time for a single FC drive in the storage pool housing the four Pool1_x datastores. Because the statistics from all the drives in the pool were similar, only a single drive is reported for clarity and readability of the graph.




I/Os serviced by the individual drives in the pool were extremely low because nearly all I/O requests were serviced by FAST Cache. The response time appeared relatively high given the disk load, but this is misleading.

FAST Cache serviced all the I/Os with a good locality of reference. Only I/Os that were very far apart made it to the actual disk drives. Because these I/Os caused large disk seeks, the response time appeared high. However, because the number of I/Os was small, they did not adversely affect the performance of the desktops.

Pool LUN load


The following graph shows the LUN I/O and response time from the Pool1_3 datastore. Because the statistics from all the pools were similar, only a single pool is reported for clarity and readability of the graph.



During peak load, the LUN response time remained within 4 ms and the datastore serviced over 500 IOPS.

Storage processor I/O


The following graph shows the total I/Os served by the storage processor during the test. The replicas were owned by SP B, which caused all the reads to be handled by SP B. Therefore, the throughput and utilization for SP B were higher than those for SP A.



SP B owned all replica images, thus causing its throughput to be high compared to SP A.


Storage processor utilization

The following graph shows the storage processor utilization during the test. The replicas were owned by SP B, which caused all the reads to be handled by SP B. Therefore, the throughput and utilization for SP B were higher than those for SP A.



SP B owned all replica images, thus causing its utilization to be high compared to that of SP A, but there was still headroom available on the storage processor. Performance can be improved by splitting the replicas between both storage processors.


FAST Cache I/O The following graph shows the number of I/Os serviced from FAST Cache during the test.



At peak load, the FAST Cache serviced over 1,700 IOPS from the linked clone datastores, which is the equivalent of eight FC disks serving 200 IOPS each. The majority of the read operations are serviced from the replica during this test.

Comparison of pool disk I/O


The following graph shows a comparison of the disk I/Os in the reference architecture solution and the disk I/Os in the baseline configuration.



FAST Cache reduced the peak load on the disks from 125 IOPS in the baseline configuration to 23 IOPS in the reference architecture solution, which was a decrease of more than five times in the I/Os to be serviced from the pool disks.

Comparison of pool disk response time

The following graph shows a comparison of the disk response time in the reference architecture solution and the disk response time in the baseline configuration.




The disk response time stayed below 15 ms for the duration of the test in the reference architecture.

The number of I/Os serviced by the individual drives in the pool were extremely low because nearly all I/O requests were serviced by FAST Cache. Requests serviced by the drives tend to have a very poor locality of reference, which causes very large disk seeks. These large seek distances cause high response times.

Due to the low number of I/Os, there are very few samples to measure and the samples have relatively high response times. However, when combined with the I/Os serviced from FAST Cache, which have a very low response time, the overall response time from the client remains very low.

ESX CPU load


The following graph shows the CPU load from the ESX servers in the View cluster. A single server is reported because all servers had similar results.



The CPU load on the ESX server was well within the acceptable limits during this test.

ESX disk response time

The following graph shows the Average Guest Millisecond/Command counter, which is shown as GAVG in esxtop. This counter represents the response time for I/Os issued to the storage array.



GAVG for both the EFD replica storage and the linked clone storage on the Pool1_x datastores was well below 10 ms for the majority of the test, with a single spike to 10 ms.

Patch install results


Test methodology

This test was performed by pushing five security updates to all desktops using Microsoft System Center Configuration Manager (SCCM). The desktops were divided into five collections, each containing 100 desktops.

The collections were configured to install updates in a 12-minute staggered schedule. This caused all patches to be installed in one hour.

EFD replica disk load


The following graph shows the I/O and response time metrics from one of the EFDs that contain the replica datastore.



The replica EFDs showed good response time during the entire test.

**EFD replica
LUN load**


The following graph shows the I/O and response time metrics from the replica LUN.




The replica LUN serviced over 1,000 IOPS during peak load with a maximum response time below 10 ms, which indicates that the LUN was not driven to saturation.

**Pool individual
disk load**

The following graph shows the disk I/O and response time for a single FC drive that consists of the storage pool housing the four Pool1_x datastores. Because the statistics from all the drives in the pool are similar, the statistics of a single drive is shown in the graphs for clarity and readability.






The I/Os that were serviced by the individual drives in the pool did not peak above 150. The disk response time briefly spiked above 20 ms when each desktop collection installs the security updates.

Pool LUN load


The following graph shows the LUN I/O and the response time from the Pool1_3 datastore. Because the statistics from all the drives in the pool were similar, the statistics of a single drive are shown in the graphs for clarity and readability.



During the peak load, the LUN response time did not move above 10.5 ms when the datastore serviced over 2,000 IOPS during the peak load.

Storage processor I/O


The following graph shows the total I/Os serviced by the storage processor during the test.



The replicas were owned by SP B, which caused all reads to be handled by SP B. This resulted in throughput and utilization of SP B to be higher than those of SP A.

Storage processor utilization


The following graph shows the storage processor utilization during the test.



The replicas were owned by SP B, which caused all reads to be handled by SP B. This resulted in throughput and utilization of SP B to be higher than those of SP A. The patch install for each desktop collection caused high CPU utilization briefly during peak load. However, the headroom was still available.

FAST Cache I/O


The following graph shows the number of I/Os serviced from FAST Cache during the test.



The FAST Cache serviced over 6,200 IOPS at the peak load from the linked clone datastores, which is the equivalent to 31 FC disks servicing 200 IOPS each.

Comparison of pool disk I/O


The following graph shows the comparison of the disk I/Os in the reference architecture solution when compared to the baseline configuration.



FAST Cache reduced the peak load on the disks from 391 IOPS to 154 IOPS, which is more than two times decrease in the I/Os required to be serviced from the pool disks.

Comparison of pool disk response time


The following graph shows a comparison of the disk response time in this reference architecture solution when compared to the baseline configuration.



Disk response is significantly reduced during the patch test with FAST Cache. The response time remained below 34 ms for the duration of the test for this reference architecture, and for the baseline configuration, it peaked over 78 ms.

ESX CPU load


The following graph shows the CPU load from the ESX servers in the View cluster. Because all servers had similar results, the results from a single server are shown in the graph.



The ESX server CPU load was well within the acceptable limits during the duration of this test.

ESX disk response time

The following graph shows the Average Guest Millisecond/Command counter, which is shown as GAVG in esxtop. This counter represents the response time for the I/Os issued to the storage array.



GAVG for both the EFDs, replica storage and linked clone storage on the Pool1_x datastores were well below 20 ms for the entire test.


LoginVSI results

Test methodology

This test was conducted by scheduling all 500 users to connect over a Remote Desktop Protocol (RDP) connection within a 30-minute window and start the LoginVSI-medium workload. This workload was run for two hours in a steady state to observe the load on the system.

EFD replica disk load


The following graph shows the I/O and response time metrics from one of the EFDs that contain the replica datastore.



The response time during this test was excellent.

EFD replica LUN load


The following graph shows the I/O and response time metrics from the replica LUN.



The replica LUN serviced nearly 2,500 IOPS during peak load with a maximum response time of less than 2 ms, which indicates that the LUN was not driven to saturation.

Pool individual disk load

The following graph shows the disk I/O and response time for a single FC drive that contains the storage pool housing the four Pool1_x datastores. Because the statistics are similar between the drives in the pool, only a single disk is reported for clarity and readability of the graphs.




I/Os serviced by the individual drives in the pool were extremely low because nearly all I/O requests were serviced by FAST Cache. The response time appeared relatively high given the disk load. But, this was misleading.

FAST Cache serviced all I/Os with a good locality of reference. Only I/Os that were very far apart made it to the actual disk drives. Because these I/Os caused large disk seeks, the response time appeared high. However, a small number of these I/Os do not adversely affect the performance of the desktops.

Pool LUN load


The following graph shows the LUN I/O and response time from the Pool1_3 datastore. Because the statistics from all the pools are similar, only a single pool is reported for clarity and readability of the graphs.



During peak load, the LUN response time never moved above 4 ms while the datastore serviced over 900 IOPS during peak load.


Storage processor I/O

The following graph shows the total I/Os served by the storage processor during the test. The replicas were owned by SP B, which caused all the reads to be handled by SP B. Therefore, the throughput and utilization for SP B are higher than those for SP A.




Storage processor utilization

The following graph shows the storage processor utilization during the test. The replicas were owned by SP B, which caused all the reads to be handled by SP B. Therefore, the throughput and utilization for SP B are higher than those for SP A.



The storage processor utilization peaks at 58 percent during the logon storm. The steady state performance was stable with a utilization of 40 percent on both storage processors.


FAST Cache I/O The following graph shows the number of I/Os serviced from FAST Cache during the test.



FAST Cache serviced over 3,500 IOPS at peak load from the linked clone datastores, which is the equivalent of 18 FC disks serving 200 IOPS each.

Comparison of pool disk I/O


The following graph shows a comparison of the disk I/Os in the reference architecture solution and the disk response time in the baseline configuration.



FAST Cache reduced peak load on the disks from 245 IOPS in the baseline configuration to 16 IOPS in the reference architecture solution — which is more than 15 times reduction in the I/O required to be serviced from the pool disks.

Comparison of pool disk response time


The following graph shows a comparison of the disk response time in the reference architecture solution and the disk response time in the baseline configuration.



The disk response time stayed below 15 ms for the duration of the test on the reference architecture solution.

ESX CPU load


The following graph shows the CPU load from the ESX servers in the View cluster. A single server is reported because all servers had similar results.



The CPU load on the ESX server was well within the acceptable limits during this test.

ESX disk response time

The following graph shows the Average Guest Millisecond/Command counter, which is shown as GAVG in esxtop. This counter represents the response time for I/Os issued to the storage array.



GAVG for both the EFD replica storage and linked clone storage on the Pool1_x datastores were well below 10 ms during this test.