

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

EMC Unified Storage Solutions



Copyright © 2010, 2011 EMC Corporation. All rights reserved.

Published January, 2011

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

The information in this publication is provided “as is”. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

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.

VMware, ESX, vMotion, VMware View, VMware vCenter, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part number: h8027.1

Table of Contents

Reference architecture overview.....	4
Solution architecture	6
Key components	8
VMware View architecture	10
Storage architecture	11
Network configuration	12
High availability and failover.....	14
Validated environment profile.....	15
Hardware and software resources	17
Conclusion.....	19

Reference architecture overview

Document purpose

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 realworld deployments in which TCE use cases are developed and executed. These use cases provide EMC with an insight into the challenges currently facing its customers.

The document describes the reference architecture 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, which was tested and validated by the EMC Unified Storage Solutions group.

Solution purpose

The purpose of this reference architecture is to build and demonstrate the functionality, performance, and scalability of virtual desktops enabled by EMC unified storage, VMware View 4.5, and VMware View Composer 2.5. This solution is built on an EMC Celerra NS-120 platform with multi-protocol support, which enables Fibre Channel (FC) block-based storage for the VMware vStorage Virtual Machine File System (VMFS) and CIFS-based storage for user data.

This reference architecture validates the performance of the solution and provides guidelines to build similar solutions.

This document is not intended to be a comprehensive guide to every aspect 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.

The technology solution

This solution demonstrates how an EMC unified storage platform can be used to provide the storage resources for a robust VMware View 4.5 environment by using Windows 7 virtual desktops.

Planning and designing the storage infrastructure for VMware View is a critical step because the shared storage must be able to absorb large bursts of input/output (I/O) that occur during the course of a day, which can lead to periods of erratic and unpredictable virtual desktop performance. Users can adapt to slow performance, but unpredictable performance is sure to quickly frustrate them.

To provide a predictable performance to a virtual desktop infrastructure, the storage must be able to handle the peak I/O load from the clients without resulting in high response time. Designing for this workload involves deploying several disks to handle brief periods of extreme I/O pressure, which is expensive to implement.

This reference architecture shows a solution that can handle peak workloads by using around 25 percent of the disks and the new features of FLARE[®] 30 such as FAST Cache and FAST with VMware View 4.5.

The solution benefits

This solution aids in the design and implementation steps required for the successful implementation of virtual desktops on VMware View 4.5. This solution balances performance requirements and cost by using the new features in FLARE 30 such as EMC Fully Automated Storage Tiering (FAST), FAST Cache, and storage pools with Enterprise Flash Drives (EFDs).

Additionally, desktop virtualization allows organizations to exploit benefits such as:

- Increased security by centralization of business-critical information
- Increased compliance as information is moved from end points into the data center
- Simplified and centralized management of desktops

Related documents

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

- *EMC Performance Optimization for Microsoft Windows XP for the Virtual Desktop Infrastructure — Applied Best Practices*
- *Deploying Microsoft Windows 7 Virtual Desktops with VMware View — Applied Best Practices Guide*
- *EMC Infrastructure for Deploying VMware View in the Enterprise EMC Celerra Unified Storage Platforms — Solution Guide*
- *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*

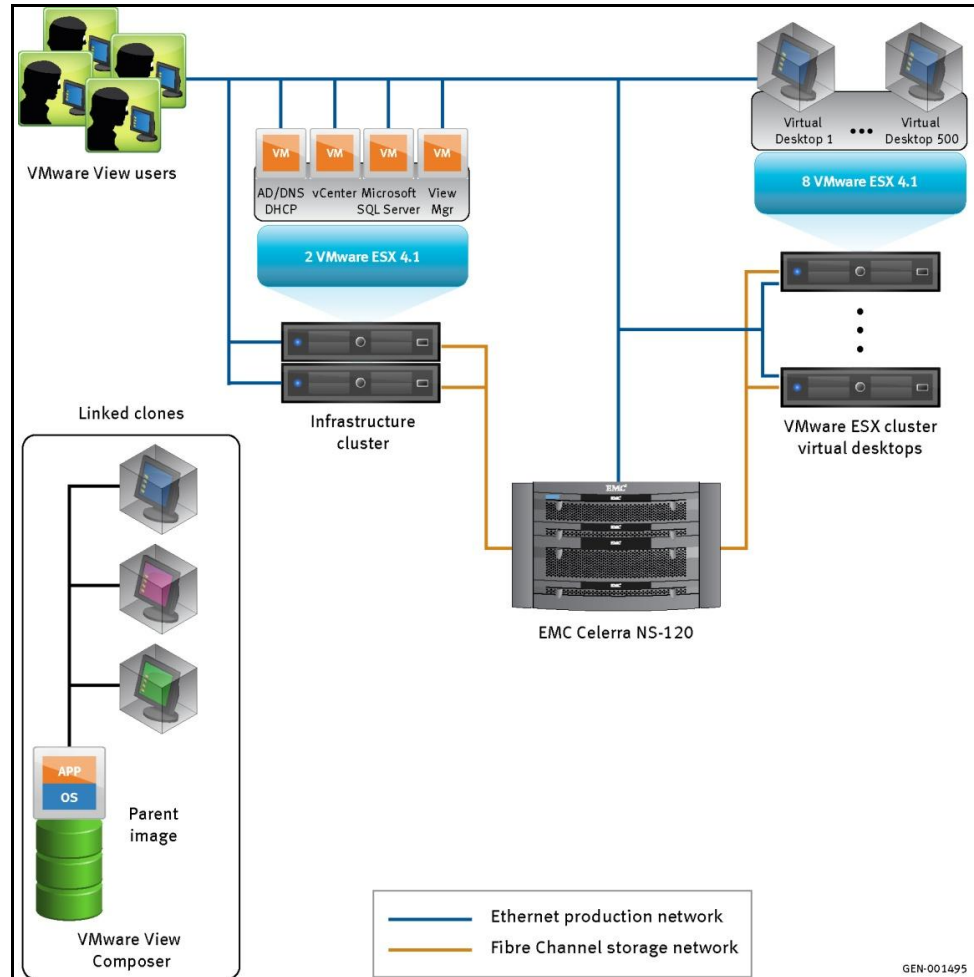
The following VMware documents, located on the VMware website, also provide useful information:

- *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*
-

Solution architecture

Architecture diagram

The following diagram shows the overall architecture of the solution.



Reference architecture overview

The validated solution is built in a VMware View 4.5 environment on a Celerra unified storage platform.

The key components of the physical architecture are:

- A two-node VMware ESX[®] 4.1 cluster to host infrastructure virtual machines
- An eight-node VMware ESX 4.1 cluster to host the virtual desktops
- An EMC Celerra unified storage platform

VMware View Manager 4.5, View Composer 2.5, VMware vCenter™ Server and all the supporting services are installed as virtual machines hosted on the infrastructure cluster.

The following are the details of the virtual architecture:

- Virtual desktops are created by using View Composer 2.5 and are deployed as linked clones.
 - The View Composer 2.5 tiered storage feature is used to store desktop replicas on dedicated LUNs that are separate from the linked clones.
 - Storage for the read-only replica images is provided by EFDs.
 - Windows folder redirection is used to redirect user data to a CIFS network share on a unified storage platform.
 - Storage pools with FC and SATA drives are used for the linked clones.
 - An EMC Celerra NS-120 storage system stores all virtual machine files (VMDK, VMX, and log).
 - VMware High Availability (HA) is used to quickly restart the virtual desktops when the hosts fail.
 - VMware Distributed Resource Scheduler (DRS) is used to load-balance virtual desktops in the ESX cluster.
-

Key components

Introduction

This section briefly describes the key components of this solution. For details on all the components that make up the reference architecture, refer to the [Hardware and software resources](#) section.

EMC Celerra unified storage platform

The EMC Celerra unified storage platform is a dedicated network server optimized for files and block access, delivering high-end features in a scalable and easy-to-use package.

For high scalability, Celerra unified storage platforms use:

- The innovative EMC CLARiiON® Fibre Channel RAID storage, delivering best-in-class availability, and data protection.
- The availability, performance, and ease of management of EMC Celerra.

Celerra unified storage systems deliver a single-box block and file solution, which offers 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 the NFS and CIFS protocols by allowing Windows and Linux/UNIX clients to share files by using the Celerra system's sophisticated file-locking mechanisms and 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.

VMware View 4.5

VMware View 4.5 is the leading desktop virtualization solution that enables desktops to deliver cloud computing services to the users. VMware View 4.5 integrates effectively with vSphere 4.1 to provide:

- View Composer 2.5 performance optimization — Optimizes storage utilization and performance by reducing the footprint of virtual desktops and using tiered storage.
 - Tiered storage support — View Composer 2.5 supports the usage of different tiers of storage to maximize performance and reduce cost.
 - Thin provisioning support — Enables efficient allocation of storage resources when virtual desktops are provisioned. This results in better utilization of the storage infrastructure and reduced CAPEX/OPEX.
-

VMware vSphere 4.1

VMware vSphere 4.1 is the market-leading virtualization platform that is 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 systems and applications just like a physical computer.

The high-availability features of VMware vSphere 4.1 coupled with DRS and Storage vMotion® enable the seamless migration of virtual desktops from one ESX server to another with minimal or no impact to the customer's usage.

**EMC FAST
Cache**

EMC FAST Cache is a new feature introduced in FLARE release 30 that 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 Flash drives. This dramatically improves response times for very active data and reduces 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 systems patches and application updates.

EMC FAST

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

Storage pools are configured with a mix of FC and SATA drives. Initially, the linked clones are placed on the FC tier. Any data created by the linked clones that is not frequently accessed is automatically migrated to the SATA storage tier releasing space in the faster FC tier for more active data.

**EMC
PowerPath
Virtual Edition**

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

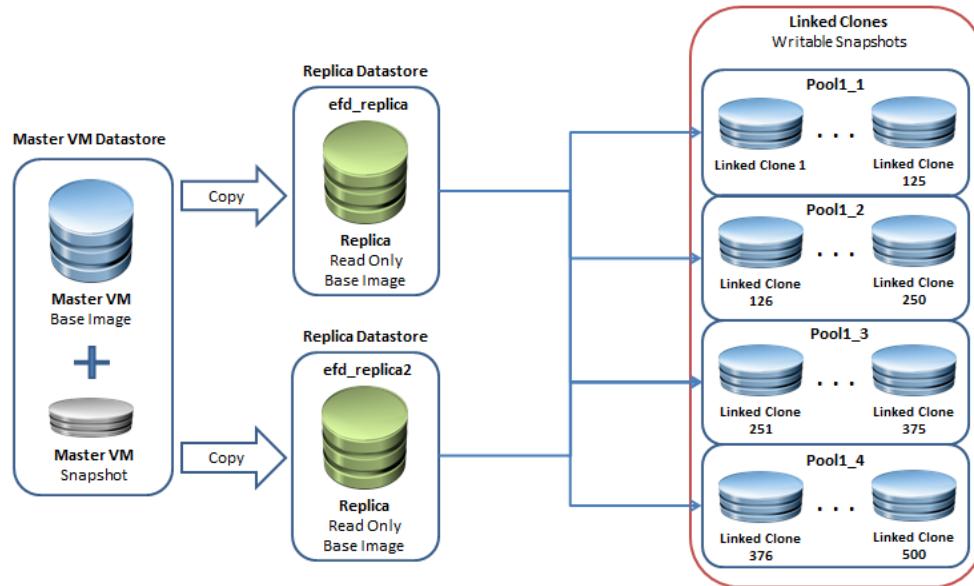
**AppSense
Environment
Manager**

AppSense Environment Manager is a component product of the AppSense Management Center. In this solution, AppSense Environment Manager is used to optimize the performance and manageability of user profiles. AppSense also helps to minimize I/O load on the storage array during user logon and logoff activities by streaming only the required data to the virtual desktops.


VMware View architecture

Linked clone overview

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



The operating system reads all the common data from the read-only replica and the unique data that is created by the operating system or user, which is stored on 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 using a common Windows 7 master image. Dedicated data stores are used for the replica images and the linked clones are spread across four data stores.

Storage architecture

Storage layout The layout of the disks is shown in the following table. Celerra NS-120 has a single backend bus and all the drives are on bus 0. Therefore, the disk numbers are given in the format of ENCLOSURE_DISK.

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	RAID 5 SYSTEM 146 GB FC	RAID 5 SYSTEM 146 GB FC	RAID 5 SYSTEM 146 GB FC	RAID 5 SYSTEM 146 GB FC	RAID 5 SYSTEM 146 GB FC	Free	Free	Free	Free	Free	RAID 1/0 Replica 100 GB	RAID 1/0 Replica 100 GB	RAID 1 FAST C 100 GB	RAID 1 FAST C 100 GB	HOT SP EFD	
Enc0	RG-0	RG-0	RG-0	RG-0	RG-0						RG-3	RG-3	EFD	EFD	RG-201	
Bus0	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	RAID 1/0 Link Cln 450 GB FC	HOT SP FC	
Enc1	Pool1	Pool1	Pool1	Pool1	Pool1	Pool1	Pool1	Pool1	Pool1	Pool1	Pool1	Pool1	Pool1	Pool1	Pool1	RG-202
Bus0	RAID 1/0 Usr Data 1TB SATA	RAID 1/0 Usr Data 1TB SATA	RAID 1/0 Usr Data 1TB SATA	RAID 1/0 Usr Data 1TB SATA	RAID 1/0 Usr Data 1TB SATA	RAID 1/0 Usr Data 1TB SATA	RAID 1/0 Usr Data 1TB SATA	RAID 1/0 Usr Data 1TB SATA	RAID 1/0 Usr Data 1TB SATA	RAID 1/0 Link Cln 1TB SATA	RAID 1/0 Link Cln 1TB SATA	RAID 1/0 Link Cln 1TB SATA	RAID 1/0 Link Cln 1TB SATA	RAID 1/0 Link Cln 1TB SATA	RAID 1/0 Link Cln 1TB SATA	HOT SP SATA
Enc2	RG-2	RG-2	RG-4	RG-4	RG-5	RG-5	RG-6	RG-6	Pool1	Pool1	Pool1	Pool1	Pool1	Pool1	Pool1	RG-203

Storage layout overview The following configuration is used in the reference architecture:

- FC disks (0_0 – 0_4) are system LUNs for both CLARiON 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 spare drives. These disks are denoted in yellow in the storage layout diagram.
- EFDs (0_10 - 0_11) on the RAID 1/0 group are used for the storage of the linked clone replicas. The EFDs are denoted in purple.
- EFDs (0_12 - 0_13) are used for EMC FAST Cache. There are no user-configurable LUNs on these drives. These disks are denoted in red.
- FC disks (1_0 – 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 for linked clone storage. 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 dark blue. Four LUNs of 1.25 TB each are carved out of the pool and presented to the ESX servers.
- SATA disks (2_0 – 2_7) with 1 TB and 7.2k rpm on RAID 1/0 groups are used for user data storage and roaming profiles. These disks are denoted in green. Two LUNs per RAID group are created and given to Celerra for the profiles and user data file systems.

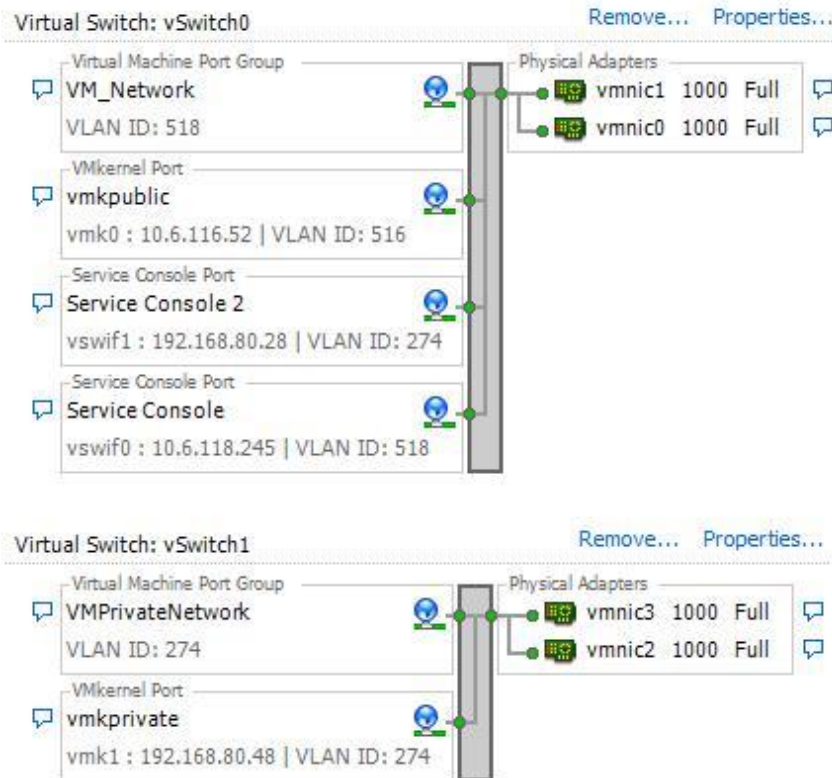
Celerra file systems

There are two file systems used by the virtual desktops — one for user profiles and the other to redirect user storage. In general, redirecting users' data out of the base image onto Celerra enables centralized administration, backup, and recovery, and makes the desktops more stateless. Each file system is exported to the environment through a CIFS share.

Network configuration

Network layout All network interfaces in this solution use 1 Gb Ethernet connections. All virtual desktops are assigned an IP address by using a Dynamic Host Configuration Protocol (DHCP) server. The Dell R710 servers use four onboard Broadcom Gb Ethernet Controllers for all the network connections.

The following diagram shows the vSwitch configuration in vCenter Server.



vSwitch0 and vSwitch1 each use two physical NICs. 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 data stores 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

Celerra NS-120 network configuration

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 Gb Ethernet controller ports. Link Aggregation Control Protocol (LACP) is used to configure ports cge0 and cge1 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 following shows the configuration of the ports:

```
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
```

High availability and failover

Introduction	This solution provides a high-availability virtual desktop infrastructure. Each component is configured to provide a robust and scalable solution for the host, connectivity, and storage layers.
Storage layer	EMC unified storage solutions provide five 9s availability by using redundant components through the array. All blades, storage processors, and array components are capable of continued operation in case of hardware failure. The RAID-disk configuration on the Celerra backend provides protection against data loss due to hard disk failures. The available hot spare drives can be dynamically allocated to replace a failing disk.
Connectivity layer	<p>The advanced networking features of Celerra, such as Fail-Safe Network (FSN) and link aggregation, provide protection against network connection failures at the array. Each ESX host has multiple connections to both Ethernet networks to guard against link failures. These connections are spread across multiple blades in a Cisco 6509 to guard against component failure in the switch.</p> <p>For FC connectivity, each host has a connection to two independent fabrics in a SAN A/B configuration. This allows complete failure of one of the SANs while still maintaining connectivity to the array.</p>
Host layer	<p>The application hosts have redundant power supplies and network connections to reduce the impact of component failures in the ESX servers. Additionally, VMware High Availability (HA) is configured on the cluster to help recover virtual desktops quickly in case of a complete host failure.</p> <p>Additionally, PowerPath Virtual Edition is configured on each ESX host, which allows dynamic load balancing of I/O requests from the server through the fabric to the array. This configuration guards against host bus adapter (HBA), path, or port failures, and also enables automated failback after the paths are restored.</p>

Validated environment profile

Observed workload

A commercial desktop workload generator was used to run an example task worker benchmark with the Windows 7 virtual desktops. The following table shows the observed workload that was used to size this reference architecture:

Windows 7 workload							
	Committed bytes	Read IOPS	Total IOPS	Write IOPS	Active RAM (MB)	% Processor time	Network bytes/s
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

Traditional sizing

From the observed workload 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 required to meet the IOPS requirements by sizing for both the average and the 95th percentile IOPS:

Windows 7 disk requirements					
Avg IOPS	# Users	Total IOPS	Read: Write Mix	IOPS	FC disks required
9	500	4500	45:55	Read: 2000	10
				Write: 2500	13
95th IOPS	# Users	Total IOPS	Read: Write Mix	IOPS	FC disks required
37	500	18500	45:55	Read: 8325	42
				Write: 10125	50

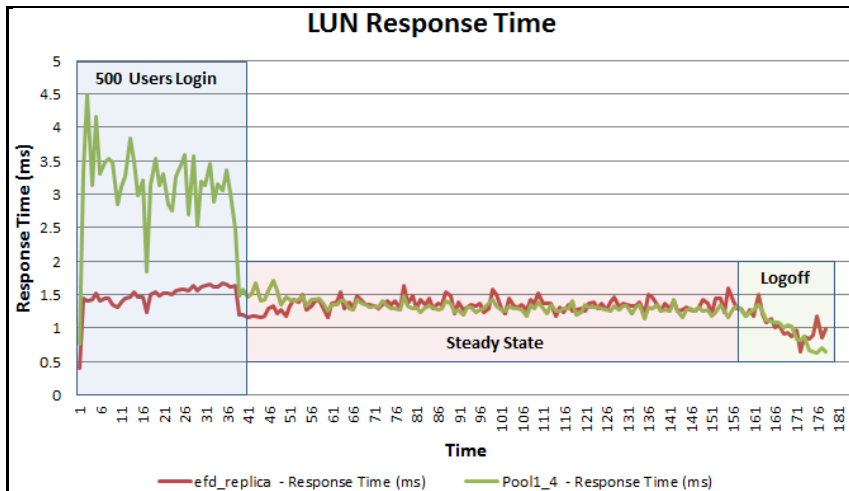
Sizing on the average IOPS can yield good performance for the virtual desktops in steady state. However, this leaves insufficient headroom in the array to absorb high peaks in I/O and the performance of the desktops will suffer during boot storms, desktop recompose or refresh tasks, antivirus DAT updates 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 combat 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 because it takes 92 disks to satisfy the I/O requirements instead of 23 disks. This leads to higher capital and operational costs.

A better sizing solution

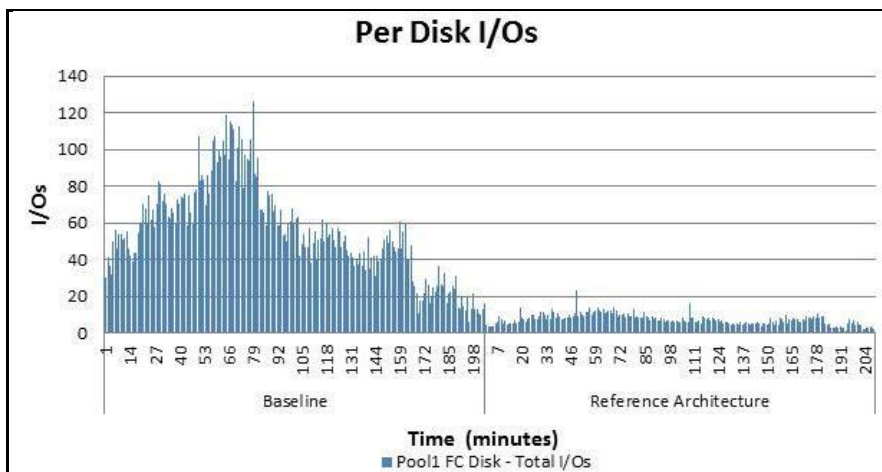
By using EFDs and the new features in vSphere 4.1, VMware View 4.5, View Composer 2.5 (tiered storage), and FLARE 30 features (FAST Cache and FAST), it is possible to design a Virtual Desktop Infrastructure (VDI) solution that reaches new levels of performance, scalability and efficiency than were possible previously.

The following graph shows the peak user load during a logon storm of 500 users over 30 minutes followed by two hours of steady state user workload, and then logoff of all 500 users over a 15-minute period. This is the typical workload that is observed on a Monday morning as users log in to their desktops for the first time.



By using the new features in FLARE 30, this reference architecture is able to satisfy peak load while keeping response time well within acceptable limits. This configuration has the potential to scale to even higher user counts if additional disks are added to increase the capacity for the additional users.

The following graph shows the average number of I/O requests serviced per disk in the Pool1_x data stores during a full virus scan of the linked clones. The peak disk load is reduced from 125 I/Os in the baseline configuration to 23 I/Os in this reference architecture.



Hardware and software resources

Hardware

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, AD, RRAS)
Cisco 6509	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

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

Conclusion

Summary

The new features introduced in FLARE 30 enable EMC unified storage arrays to drive higher storage consolidation ratios at a lower cost than otherwise possible. This reduces the capital expenditure on equipment and lowers the operational costs required to support the placement, power, and cooling of the storage arrays.

The following table compares the configuration of the reference architecture with FAST and FAST Cache with the reference architecture sized to meet the I/O requirements without the new features:

Reference architecture	Disk requirements
Reference architecture using: <ul style="list-style-type: none">• FAST• FAST Cache• EFDs, FC, and SATA disks	The required disks are: <ul style="list-style-type: none">• Five 300 GB 15k rpm disks• Fifteen 450 GB 15k rpm disks• Fifteen 1 TB 7.2k rpm disks• Five 100 GB EFDs Total: 40 disks
Reference architecture using only FC and SATA disks (traditional configuration)	The required disks are: <ul style="list-style-type: none">• Five 300 GB 15k disks• Ninety six 450 GB 15k disks• Fifteen 1 TB 7.2k rpm disks Total: 116 disks

This reference architecture is able to provide the required I/O for 500 concurrent users while reducing the number of disks by 75 percent leading to a 60 percent reduction in storage costs when compared to a solution without FAST and FAST Cache.

Next steps

EMC can help accelerate the assessment, design, implementation, and management of a virtual desktop solution while lowering the implementation risks and costs based on VMware View 4.5.

To learn more about this and other solutions, contact an EMC representative.
