

**EMC Infrastructure for Virtual Desktops
Enabled by EMC Unified Storage,
VMware vSphere 4.1, VMware View 4.0.1, and
VMware View Composer**

Reference Architecture

*Unified Storage Solutions
Corporate Headquarters*
Hopkinton, MA 01748-9103
1-508-435-1000
www.EMC.com

Copyright © 2010 EMC Corporation. All rights reserved.

Published July, 2010

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.

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

EMC Infrastructure for Virtual Desktops Enabled by EMC Unified Storage, VMware vSphere 4.1, VMware View 4.0.1, and VMware View Composer

Reference Architecture

P/N **h8003**

Chapter 1	Solution Overview.....	11
	Business challenge	12
	Technology solution for VMware View.....	12
	Solution advantages.....	12
Chapter 2	Solution Architecture	13
	Overall architecture	14
	General characteristics	14
	Load Generator	14
	Storage architecture.....	15
	CLARiiON pool and RAID groups	15
	CLARiiON FAST Cache	17
	CLARiiON FAST	18
	CLARiiON LUN layout.....	19
	Celerra file systems.....	20
	CIFS exports	20
	File system checkpoints	21
	vSphere storage.....	21
	Network architecture	22
	Celerra NS-120	23
	View Manager configuration.....	24
	Pool datastore configuration	25
	High-availability and failover.....	25
	Storage layer	25
	Connectivity layer.....	25
	Host layer.....	26
Chapter 3	Hardware and Software Resources.....	27
	Hardware resources	28
	Software resources	28

Figures

Figure 1	Solution architecture	14
Figure 2	Backend configuration	15
Figure 3	Pool and LUN configuration	16
Figure 4	RAID group configuration	17
Figure 5	FAST Cache configuration dialog box	18
Figure 6	FAST configuration	19
Figure 7	LUN configuration	20
Figure 8	CIFS exports	20
Figure 9	Checkpoint schedule	21
Figure 10	vCenter datastores	21
Figure 11	Manage Paths dialog box	22
Figure 12	vSwitch configuration in vCenter	23
Figure 13	EMC Celerra NS-120 blade ports	24
Figure 14	Unisphere network interfaces	24
Figure 15	Select datastores	25

Table 1 Solution advantages.....	12
Table 2 Hardware specifications	28
Table 3 Software specifications.....	28

About this Document

This document describes the reference architecture for an EMC infrastructure for virtual desktops enabled by EMC Celerra unified storage, VMware View 4.0.1, and VMware View Composer.

Purpose

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

Audience

This document is intended for internal EMC personnel, partners, and customers.

Scope

This document describes the reference architecture for VMware View 4.0.1 on EMC Celerra NS-120 storage for 500 users. Implementation instructions and sizing guidelines are beyond the scope of this document.

Related documents

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

- ◆ *VMware Virtual Desktop Infrastructure Planning for EMC Celerra — Best Practices Planning*
- ◆ *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*

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 Reference Architecture*
- ◆ *Storage Deployment Guide for VMware View*
- ◆ *VMware View Windows XP Deployment Guide*
- ◆ *VMware View Guide to Profile Virtualization*

Chapter 1 Solution Overview

This chapter presents these topics:

Business challenge	12
Technology solution for VMware View	12
Solution advantages	12

Business challenge

With limited resources and increasing demands, today's businesses must address the following challenges:

- ◆ Consolidate desktops scattered throughout an enterprise
- ◆ Ensure information access, availability, and continuity
- ◆ Maximize server and storage utilization, and deliver high desktop performance
- ◆ Manage upgrades and migrations quickly and easily
- ◆ Reduce the demands on limited IT resources and budgets
- ◆ Reduce the complexity of technology choices

In addition, businesses must manage IT costs and reduce the risk of business disruption.

Technology solution for VMware View

This reference architecture establishes a configuration of validated hardware and software that permits easy and repeatable deployment of virtual desktops with VMware View Manager and View Composer on a Celerra® NS-120 system utilizing EMC® Fully Automated Storage Tiering (FAST), FAST Cache, and pooled storage.

This document describes the reference architecture to configure ESX servers and the storage provided by a Celerra NS-120 system in a manner that provides performance, recoverability, and protection. In addition, these guidelines can be further extrapolated to other larger-scale Celerra systems.

Solution advantages

Table 1 shows the benefits of the VMware View solution with the Celerra NS-120 IP storage solution.

Table 1 **Solution advantages**

Benefit	Details
Maintains service levels	This solution keeps users' desktops available and running at peak performance.
Reduces support costs	This solution minimizes the cost to upgrade and maintain users' desktops.
Reduces risk	This solution offers a reference architecture that includes tested and proven configurations that improve performance and scalability.
Accelerates implementations	EMC Professional Services and ASN-certified EMC partners provide rapid assessment and efficient implementation.

Chapter 2 Solution Architecture

This chapter presents these topics:

Overall architecture	14
Storage architecture.....	15
Network architecture	22
View Manager configuration	24
High-availability and failover	25

Overall architecture

Figure 1 shows the architecture of the VMware View 4.0.1 solution environment.

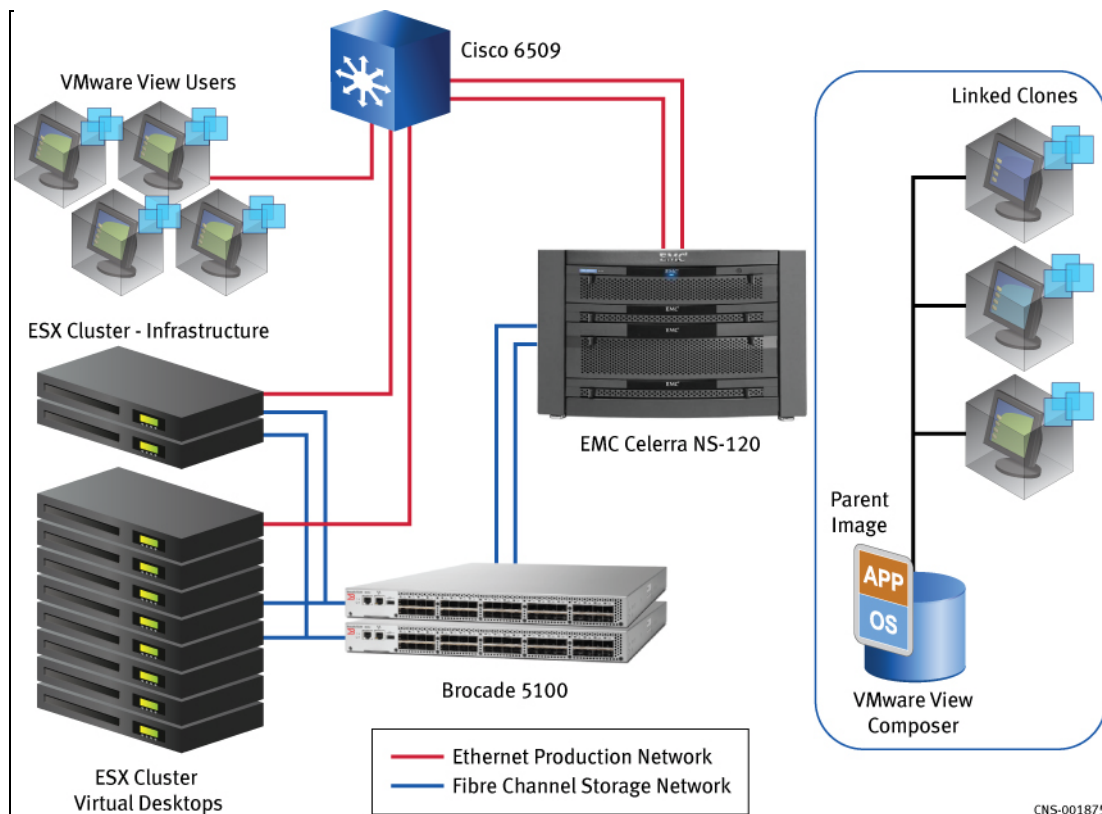


Figure 1 Solution architecture

General characteristics

The general characteristics of the solution architecture are:

- ◆ Management and infrastructure virtual machines are located on a dedicated two-node cluster to eliminate circular dependencies of the management framework.
- ◆ Virtual desktops are created and deployed by using linked clones created by View Composer.
- ◆ Storage allocation for virtual desktops is composed of a storage pool with Fibre Channel (FC) and SATA disks. Two EFDs are used to create a writeable FAST Cache. FAST and FAST Cache are enabled for the virtual desktop storage pool LUNs.
- ◆ All virtual machine files (vmdk, vmx, and log) are stored with the storage provided by the EMC Celerra NS-120 storage system.
- ◆ The ESX cluster uses both VMware HA and DRS to aid quick restart of virtual desktops due to failed hosts and for efficient load balancing of virtual desktop resources in the cluster.

Load Generator

This reference architecture leverages the Login VSI tool from Login Consultants to generate user load on the system. The Login VSI tool uses AutoIT scripts to simulate user interactions with common applications such as Office 2007, Adobe Reader, Internet Explorer, and multimedia type workloads.

Storage architecture

In this solution, Celerra NS-120 provides both Fibre Channel and file services to the ESX cluster. Each ESX server has two 4 Gb links into the fully redundant fabric, one each to SAN A and SAN B, which are used to store the VMDK files of the virtual machines. Additionally, each server has four 1 Gb Ethernet links connected to the production local area network (LAN) for virtual machine traffic, home directories, roaming profile storage, and management purposes.

CLARiiON pool and RAID groups

The backend storage layout to support 500 users is configured as shown in Figure 2. 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.

Figure 2 shows backend RAID group and pool layout in a tabular format.

Storage Pool Layout - View 4.0.1 - Release 30															
Front View - RG Type / (Disk Size in GB) / Storage Pool Type - ID															
Slot	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Bus0 Enc0	RAID5 SYSTEM 300 GB FC RG - 0	RAID5 SYSTEM 300 GB FC RG - 0	RAID5 SYSTEM 300 GB FC RG - 0	RAID5 SYSTEM 300 GB FC RG - 0	RAID5 SYSTEM 300 GB FC RG - 0	HOT SP 300 GB FC RG - 200	RAID1/0 Profiles 450 GB FC RG - 3	RAID1/0 Profiles 450 GB FC RG - 3	RAID1/0 Profiles 450 GB FC RG - 3	RAID1/0 Profiles 450 GB FC RG - 3	FREE	FREE	RAID1/0 FAST C 73 GB EFD	RAID1/0 FAST C 73 GB EFD	HOT SP 73 GB EFD RG - 201
Bus0 Enc1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	RAID1/0 Link Cln 450 GB FC Pool1	HOT SP 450 GB FC RG - 202
Bus0 Enc2	RAID1/0 Usr Data 1TB SATA RG - 2	RAID1/0 Usr Data 1TB SATA RG - 2	RAID1/0 Usr Data 1TB SATA RG - 2	RAID1/0 Usr Data 1TB SATA RG - 2	RAID1/0 Usr Data 1TB SATA RG - 2	RAID1/0 Usr Data 1TB SATA RG - 2	RAID1/0 Usr Data 1TB SATA RG - 2	RAID1/0 Usr Data 1TB SATA RG - 2	RAID1/0 Link Cln 1TB SATA Pool1	RAID1/0 Link Cln 1TB SATA Pool1	RAID1/0 Link Cln 1TB SATA Pool1	RAID1/0 Link Cln 1TB SATA Pool1	RAID1/0 Link Cln 1TB SATA Pool1	RAID1/0 Link Cln 1TB SATA Pool1	HOT SP 1TB SATA RG - 203

Figure 2 Backend configuration

The storage required for 500 virtual desktops is completely contained in enclosures 0-2. To configure for 1,000 users, the configuration is replicated with the exception of the system drives and the two EFDs used for the FAST Cache.

Figure 2 shows the following:

- ◆ FC disks 0_0 – 0_4 are system LUNs for both CLARiiON® and Celerra. During installation of Celerra, the free space on these disks is allocated to a storage pool. This storage pool is used for storing the SavVol for file system checkpoints on the Celerra Common Internet File System (CIFS) exports. No other user storage is present on the disks.
- ◆ Disks 0_5, 0_14, 1_14, and 2_14 are hot spare drives and are denoted in yellow.
- ◆ FC disks 0_6 – 0_9 on RAID 1/0 denoted in dark green are used to store roaming profiles. This RAID group is given to Celerra for a CIFS export.
- ◆ FC disks 0_12 – 0_13 on the RAID 1/0 group are used for FAST Cache. This is configured automatically by the system when FAST Cache is enabled. This is denoted in gray.
- ◆ 450 GB 15k FC disks (1_0 – 1_13) and 1 TB 7.2k SATA disks (2_8 – 2_13) on the RAID 1/0 group are used for linked clone storage. Initial placement of linked clone data is on FC. The mixed drive allocation allows FAST to move old data to the SATA drives, which allows doubling the pool size for the linked clones with cheaper and lower power storage. FAST Cache is enabled for the entire pool. These disks are denoted in dark blue.
- ◆ 1 TB 7.2k rpm SATA disks (2_0 – 2_7) on the RAID 1/0 group are used for user data storage. These disks are denoted in light green. This RAID group is given to Celerra for a CIFS export.
- ◆ Additional utility storage to support the Login VSI tool is not shown because this is not part of a production deployment.

Figure 3 shows the pool and LUN layout from Unisphere™.

The screenshot displays the Unisphere management interface. At the top, there are tabs for 'Pools' and 'RAID Groups'. The 'Pools' section shows a table with the following data:

Name	FAST Cache	RAID Type	Drive Type	User Capacity (GB)	Auto-Tiering Status
pool1	On	RAID1/0	Mixed	5548.339	Scheduled
vmstore_pool	Off	RAID5	FC	1070.065	Manual

Below the table, it indicates '1 Selected' and provides buttons for 'Create', 'Delete', 'Properties', and 'Expand'. The status '2 items' and 'Last Refreshed: 2010-06-05 12:44:02' are also visible.

The 'Details' section is open to 'Pool LUNs'. It shows a table of LUN configurations:

Name	ID	User Capacity (GB)	Initial Tier	Tiering Policy
LUN 11 - pool1_1	11	1024.000	Optimize for Pool Performance	Auto-Tier
LUN 12 - pool1_2	12	1024.000	Optimize for Pool Performance	Auto-Tier
LUN 13 - pool1_3	13	1024.000	Optimize for Pool Performance	Auto-Tier
LUN 14 - pool1_4	14	1024.000	Optimize for Pool Performance	Auto-Tier
LUN 15 - pool1_5	15	1024.000	Optimize for Pool Performance	Auto-Tier

At the bottom of the 'Pool LUNs' section, it shows '0 Selected' and buttons for 'Delete', 'Properties', 'Add to Storage Group', and 'Add'. The status 'Filtered: 5 of 5' is also present.

Figure 3 Pool and LUN configuration

Figure 4 shows the RAID group layout from Unisphere.

The screenshot displays the RAID Groups configuration interface. The main table lists four RAID groups:

ID	Drive Type	RAID Type	User Capacity (GB)	% Full
RAID Group 0	FC	RAID5	825.617	[Progress Bar]
RAID Group 1	FC	RAID1/0	536.742	[Progress Bar]
RAID Group 2	SATA II	RAID1/0	3668.597	[Progress Bar]
RAID Group 3	FC	RAID1/0	1207.719	[Progress Bar]

Below the table, the 'Details' section is active, showing the 'LUNs' tab. It contains a table with the following data:

Name	ID	State	User Capacity (GB)	Host Information
LUN 7 - UserData1	7	Ready	1024.000	RTPSOL22-SPA

Figure 4 RAID group configuration

CLARiiON FAST Cache

EMC CLARiiON FAST Cache is a new feature introduced in FLARE® release 30 that enables EFDs to be utilized as an expanded cache layer for the array. Celerra NS-120 supports using two 73 GB STEC SLC drives in a RAID 1 configuration for a 66 GB read/write capable cache. Larger arrays support FAST Cache sizes up to 2 TB.

FAST Cache is an array-wide feature that can be enabled for any LUN or a storage pool. FAST Cache works by examining 64 KB chunks of data in FAST Cache enabled objects on the array. Any 64 KB chunk that has data accessed frequently will be copied to the FAST Cache. Subsequent accesses to that data chunk are serviced from the Flash drives backing the FAST Cache. This allows promotion of very active data to Flash drives, which dramatically improves response times for very active data and reduces data hot spots that can occur within the LUN.

The FAST Cache is both an extended read and write cache, which can have tremendous benefits for absorbing read-heavy activities such as boot storms and antivirus scans, and write-heavy workloads such as patch and application updates.

Figure 5 shows the FAST Cache dialog box.

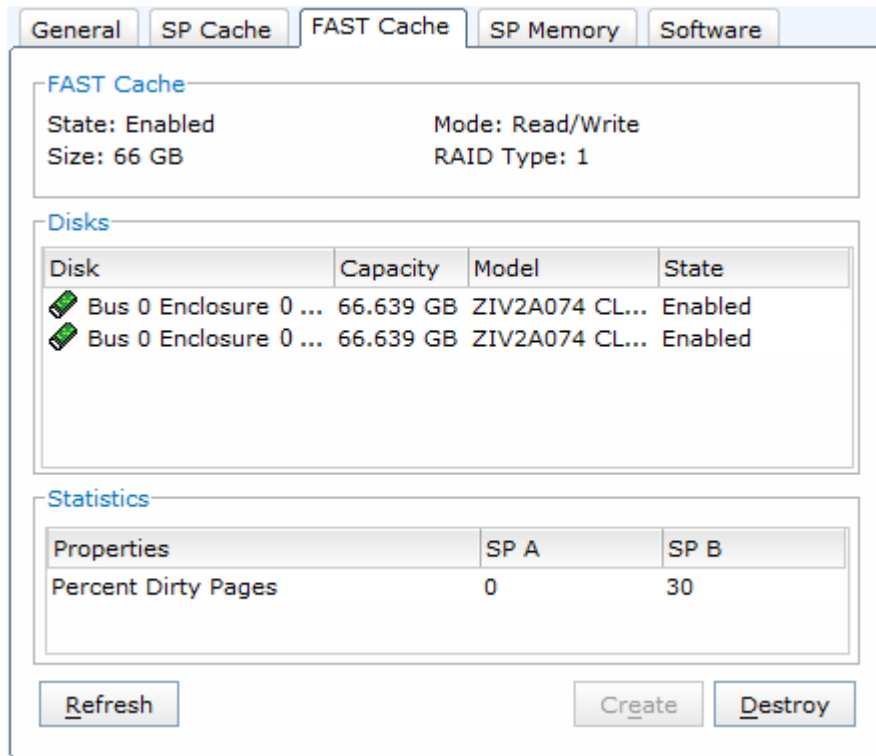


Figure 5 FAST Cache configuration dialog box

CLARiiON FAST

FAST is a pool-based feature of FLARE 30 that supports scheduled migration of data to different storage tiers based on performance requirements of individual 1 GB chunks of data.

The “pool1” storage pool is a FAST-enabled pool configured with a mix of FC and SATA disks. Initially, linked clones are placed on the FC tier. Any data created by the linked clones that are not frequently accessed is automatically migrated to the SATA storage tier releasing the space in the faster FC tier for more active data. The FAST data migration is configured to start at 11 P.M. each evening and run for a maximum of eight hours, finishing at 7 A.M. to minimize the impact on production.

Figure 6 shows the FAST configuration.

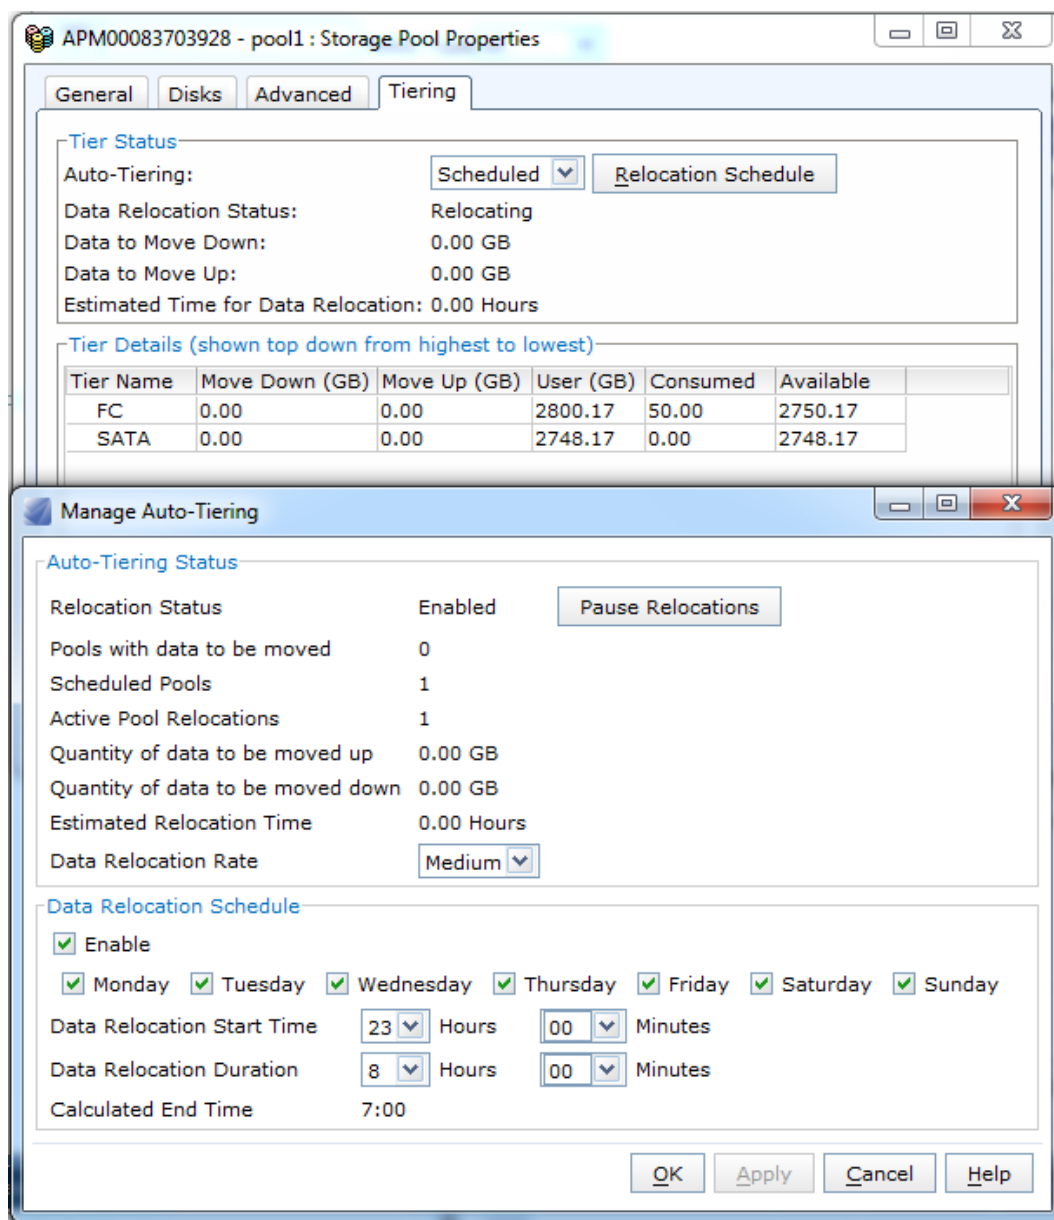


Figure 6 FAST configuration

CLARiiON LUN layout

The CLARiiON LUN configurations required to support 500 users as shown in Figure 7 are:

- ◆ LUNs 6, 7, and 8 are provided to Celerra for the creation of CIFS exports to support the virtual desktop environment.
- ◆ LUNs 11-15 are allocated out of the "pool1" storage pool. Each LUN is 1 TB in size and has a linked clone pool consisting of 100 desktops deployed on it. This allows enough capacity for each desktop to grow approximately 10 GB in size. Initially, desktop data is placed on FC. FAST moves inactive data to the SATA tier, which releases space on the higher performance FC tier for frequently accessed data. FAST Cache is also enabled for the pool that places very frequently accessed data of the desktops in the memory on the array.

Figure 7 shows the LUN layout.

Name	ID	State	User Capacity (GB)	Host Information
LUN 6 - testharness	6	Ready	536.742	RTPSOL22-SPA
LUN 7 - UserData1	7	Ready	1024.000	RTPSOL22-SPA
LUN 8 - profiles1	8	Ready	1024.000	RTPSOL22-SPA
LUN 11 - pool1_1	11	Ready	1024.000	rtpsol224.vdi.local - Datastc
LUN 12 - pool1_2	12	Ready	1024.000	rtpsol224.vdi.local - Datastc
LUN 13 - pool1_3	13	Ready	1024.000	rtpsol224.vdi.local - Datastc
LUN 14 - pool1_4	14	Ready	1024.000	rtpsol224.vdi.local - Datastc
LUN 15 - pool1_5	15	Ready	1024.000	rtpsol224.vdi.local - Datastc
LUN 20 - VMSTORE	20	Ready	1024.000	rtpsol260.vdi.local - Datastc

Figure 7 LUN configuration

Celerra file systems

Two file systems are exported for use by the virtual desktops. One file system is exported for use by the Login VSI tool that is not part of a production deployment. In general, redirecting user’s data out of the base image to Celerra enables centralized administration, backup and recovery, and makes the desktops more stateless. The following file systems are exported to the environment through a CIFS share.

- ◆ Userdata1_fs – This file system is used by the HomeDir feature of Celerra to automatically map the H:\ drive of each virtual desktop to their own dedicated subfolder on the share. This ensures that each user has a dedicated home drive share with exclusive rights to that folder. The Documents folder will be redirected to this path. This allows users to recover data in the Documents folder by using the Previous Versions functionality if snapshots are employed on the file system. The file system is set at an initial size of 1 TB but can extend itself automatically as more space is required.
- ◆ Profile1_fs – This file system is used to house users’ roaming profiles.
- ◆ Testharness_fs – Login VSI tool uses this file system, which is not part of a production deployment for administrative use. This is included for the sake of completeness.

CIFS exports

Three CIFS exports were created, one for each Celerra file system, as shown in Figure 8.

Name	File System	Data Mov...	CIFS Serv...	Comment
profiles	profile1_fs	server_2	VDICIFS	Share for roaming profiles
testharness	testharness_fs	server_2	VDICIFS	Utility share for LoginVSI
userdata	userdata1_fs	server_2	VDICIFS	Share for users redirected data

Figure 8 CIFS exports

File system checkpoints

Checkpoints are configured for all file systems to aid in quick recovery of user files. By configuring snapshots on the userdata1 file system, users can recover files quickly without administrator intervention by using the previous versions dialog box when browsing their home directory.

The checkpoint schedule configured is shown in [Figure 9](#).

Name	Next Run	State	Description
profile1_weekly_ckpt	Mon Jun 07 01:15:00 EDT 2010	Pending	Weekly snap of profile1_fs
testharness_daily_ckpt	Mon Jun 07 02:15:00 EDT 2010	Pending	Daily checkpoint of testharness_fs
userdata_daily_ckpt	Mon Jun 07 00:15:00 EDT 2010	Pending	Daily read only checkpoint of userdata_fs

Figure 9 Checkpoint schedule

vSphere storage

The datastore configuration for 500 users is shown in [Figure 10](#). To deploy 1,000 users, the configuration is doubled.

Identification	Status	Device	Capacity	Free	Type
pool1_1	Normal	naa.600601603c80210...	1,023.75 G	839.44 GB	vmfs3
pool1_2	Normal	naa.600601603c80210...	1,023.75 G	1,023.20 G	vmfs3
pool1_3	Normal	naa.600601603c80210...	1,023.75 G	1,023.20 G	vmfs3
pool1_4	Normal	naa.600601603c80210...	1,023.75 G	1,023.20 G	vmfs3
pool1_5	Normal	naa.600601603c80210...	1,023.75 G	1,023.20 G	vmfs3
VMSTORE	Normal	naa.600601603c80210...	1,023.75 G	770.85 GB	vmfs3

Figure 10 vCenter datastores

The following is a datastore configuration for 500 users:

- ◆ Pools 1_1 through 1_5 – Each of these datastores holds 100 users with a 1 TB size. These datastores allow each desktop to grow to a maximum size of 10 GB. The pools of desktops created in View Manager are balanced across all these datastores.
- ◆ Vmstore – This datastore is used for infrastructure virtual machines.

VMware View workloads tend to be volatile in nature and are prone to “micro bursting.” This means that there can be brief periods of extremely high I/O. This type of I/O behavior can saturate paths and front-end ports to the array if they are not planned carefully.

To provide the most efficient use of all 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. The Manage Paths configuration dialog box for a LUN is shown in Figure 11. Notice that all paths are active and capable of servicing I/Os.




Figure 11 Manage Paths dialog box

Network architecture

All network interfaces in this solution utilize 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 utilize the four onboard Broadcom Gb Ethernet Controllers for all network connections.

The vSwitches are configured as shown in Figure 12.

Networking

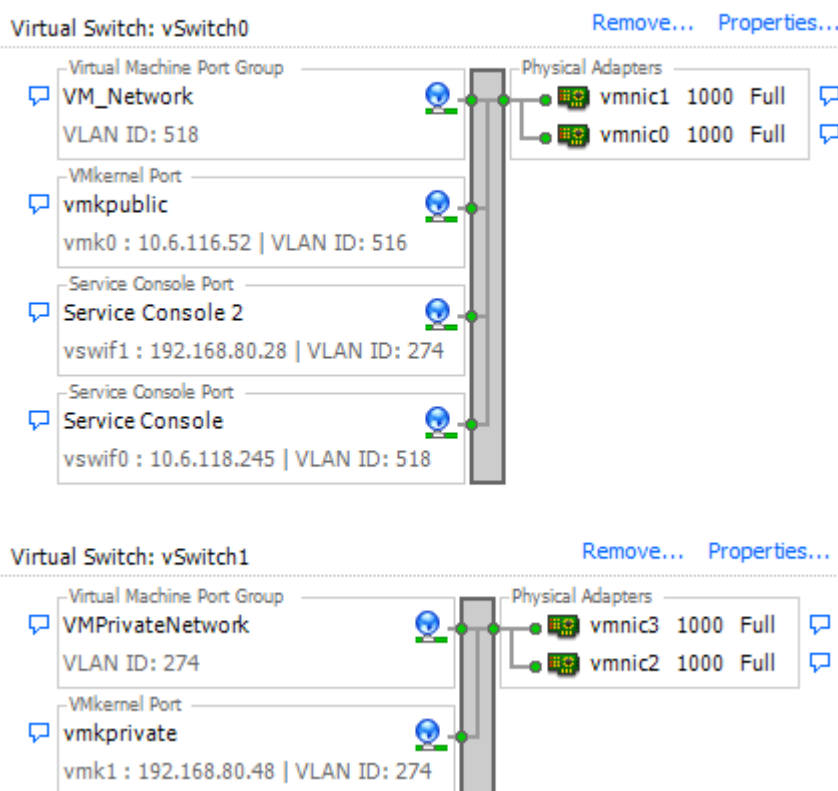


Figure 12 vSwitch configuration in vCenter

All components of the VMware View implementation are present in private VLAN 274 and pass the traffic through vSwitch1. External management connections to the cluster are maintained through vSwitch0.

The VM_Network, vmkpublic, and Service Console port groups are for administrative use for the cluster and their traffic is separated from that of the virtual desktops by using VLANs and a separate vSwitch with dedicated uplink ports to the network.

The VMPrivateNetwork port group in private VLAN 274 handles all the IP traffic from the virtual desktops and infrastructure virtual machines that support the implementation.

The vmkprivate port group handles VMotion duties for virtual machines in private VLAN 274.

Celerra NS-120

Celerra NS-120 contains two blades. These blades can operate either independently or in the active/passive mode, with the passive blade serving as a failover device for the active blade. In this solution, the blades are operated in active/passive mode.

The NS-120 blade consists of 4 Gb Ethernet controller ports. Ports cge0 and cge1 are configured using LACP to support virtual machine traffic for roaming profiles, and home folder and external access for test harness. 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.

Figure 13 shows the ports on the back of an EMC Celerra NS-120 blade.

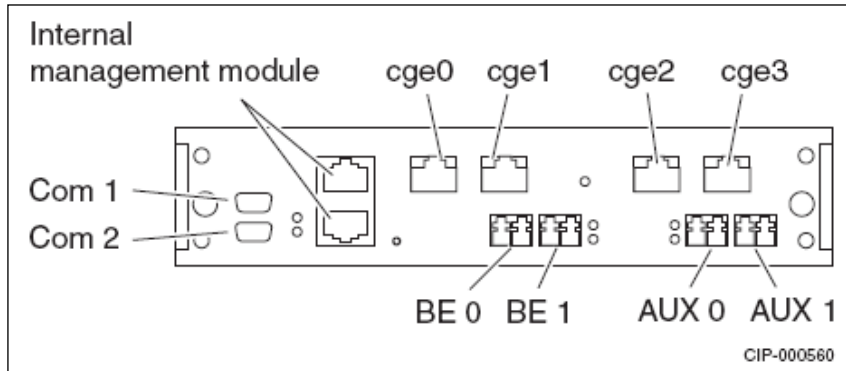


Figure 13 EMC Celerra NS-120 blade ports

The configuration can be seen in the following:

```
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
```

This can also be seen in Unisphere as shown in Figure 14.

Address	Name	Netmask	Data Mover	Device	State
10.6.121.55	external_interface	255.255.255.0	server 2	lacp1	Up
128.221.252.2	el30	255.255.255.0	server 2	mge0	Up
128.221.252.3	el30	255.255.255.0	server 3	mge0	Up
128.221.253.2	el31	255.255.255.0	server 2	mge1	Up
128.221.253.3	el31	255.255.255.0	server 3	mge1	Up
192.168.80.5	lacp1_int	255.255.240.0	server 2	lacp1	Up

Figure 14 Unisphere network interfaces

View Manager configuration

This solution uses VMware View Manager 4.0.1 as the desktop broker. Desktop configuration, deployment, and user management are done from this interface.

Pool datastore configuration

For this reference architecture, five linked clone-based pools of 100 users each are created. Each desktop pool is spread across all five pool datastores. The datastore selection for a desktop pool is shown in [Figure 15](#).

Add Desktop

Datastores
Choose the datastores on which to store new virtual machines. Only the datastores that can be used by the host (or every host in a cluster) are shown.

=Local Datastore =Shared Datastore

<input type="checkbox"/>	Name	Capacity (GB)	Free (GB)	Type	Storage Overcommit
<input checked="" type="checkbox"/>	pool1_1	1023.75	839.34	VMFS	Aggressive▼
<input checked="" type="checkbox"/>	pool1_2	1023.75	1023.2	VMFS	Aggressive▼
<input checked="" type="checkbox"/>	pool1_3	1023.75	1023.2	VMFS	Aggressive▼
<input checked="" type="checkbox"/>	pool1_4	1023.75	1023.2	VMFS	Aggressive▼
<input checked="" type="checkbox"/>	pool1_5	1023.75	1023.2	VMFS	Aggressive▼
<input type="checkbox"/>	solutions2	961.66	335.15	NFS	
<input type="checkbox"/>	VMSTORE	1023.75	745.66	VMFS	

Storage Overcommit determines how the system assigns new VMs to the free space available on a datastore. As the level increases, less space will be reserved for individual VM growth but more VMs will fit on the datastore.

Minimum Recommended Required for 50% Storage Provision Maximum Required

Free space **4932.14 GB** 174 GB 699 GB 1,299 GB

< Back Next > Cancel

Figure 15 Select datastores

High-availability and failover

This reference architecture consists of a highly available virtual desktop infrastructure. Each component is configured to provide a robust and scalable solution throughout the ESX host, connectivity, and storage layers.

Storage layer

The EMC unified storage solutions provide five 9s availability through the use of redundant components through the array. All Data Movers, storage processors, and array components are capable of continued operation in the event of a hardware failure. The RAID disk configuration on the Celerra backend provides protection against data loss due to hard disk failures with available hot spare drives that can be dynamically allocated to replace a failing disk.

Connectivity layer

The advanced networking features of the Celerra, such as failsafe networks 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 the Cisco 6509 to guard against component failure in the switch.

For Fibre Channel 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.

Host layer

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 the event of a complete host failure.

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 HBA, path, or port failures, and also allows automated failback once the paths are restored.

Chapter 3 Hardware and Software Resources

This chapter presents these topics:

Hardware resources	28
Software resources	28

Hardware resources

Table 2 lists the hardware resources required for this solution.

Table 2 Hardware specifications

Hardware	Qty	Configuration	Notes
EMC Celerra NS-120	1	<ul style="list-style-type: none"> NS-120 Three DAEs configured with: <ul style="list-style-type: none"> Six 300 GB 15k FC disks Nineteen 450 GB 15k FC disks Fifteen 1 TB 7.2k SATA disks Two 73 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, 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 5100	2	24 enabled 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
Desktops/VMs	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 for a medium Login VSI workload

Software resources

Table 3 lists the software resources required for this solution.

Table 3 Software specifications

Software	Configuration
Celerra NS-120 (Celerra shared storage, file systems, and snaps)	
NAS/DART	Release 6
CLARiiON FLARE	Release 30
ESX servers	
ESX	ESX 4.1
vCenter Server	
OS	Windows 2008 R2

Software	Configuration
VMware vCenter	4.1
View Manager	4.0.1 (Build 233023)
View Composer	2.5 (Build 255175)
PowerPath Virtual Edition	5.4 SP2
Desktops/virtual machines	
Note: This software is used for generating test load.	
OS	MS Windows 7 Enterprise (32-bit)
VMware Tools	8.3.1
Login VSI	Version 212
Microsoft Office	Office 2007 SP2
Internet Explorer	8.0.7600.16385
Adobe Reader	9.1.0
McAfee virus scan	8.7.0i Enterprise