DEPLOYMENT GUIDE

# EMC CLARiiON Storage

*A Guide to Deploying EMC CLARiiON CX4-240 iSCSI with VMware View*

# Contents

# Introduction

Built on the industry-leading VMware® virtualization platform, VMware View™ 3 is a universal client solution that lets you manage operating systems, hardware, applications, and users independently of each other, wherever they may reside. VMware View streamlines desktop and application management, reduces costs, and increases data security through centralization, resulting in greater end user flexibility and IT control. VMware View enables you to extend the value of VMware Infrastructure to encompass not only desktops in the datacenter but also applications and the delivery of these environments securely to remote clients, online or off, anywhere.

VMware View transforms the way you use and manage desktop operating systems. You can deploy desktop instances rapidly in secure datacenters to facilitate high availability and disaster recovery, protect the integrity of enterprise information, and remove data from local devices that are susceptible to theft or loss. Isolating each desktop instance in its own virtual machine eliminates typical application compatibility issues and improves and delivers a more personal computing environment.

The EMC® CLARiiON® CX4 series with UltraFlex™ technology is the fourth-generation CX series and continues EMC's commitment to maximizing customers' investments in CLARiiON technology by ensuring that existing resources and capital assets are optimally utilized as customers adopt new technology.

New UltraFlex technology enables online expansion via hot-pluggable I/O modules so you can meet current Fiber Channel and iSCSI needs and accommodate new technologies, too. You have the option to populate these flexible slots with either Fibre Channel interfaces or iSCSI interfaces.

This document provides a detailed summary of the design and configuration of an environment incorporating the EMC CLARiiON CX4-240 using iSCSI for use with VMware View. VMware and EMC used this configuration as part of the VMware View reference architecture validation. This guide provides guidance specific to the use of EMC CLARiiON CX4-240 storage with VMware View. Although this configuration was used specifically with the VMware View reference architecture for large-scale VMware View deployments, the information provided in this guide can be helpful to anyone planning to deploy VMware View using an EMC CLARiiON CX4-240. In addition, you can extrapolate from these guidelines, which apply directly to the CX4 system, to other CLARiiON systems. Before using this document, you should have a working knowledge of VMware View as well as CLARiiON technologies.

# Hardware and Software requirements

The following sections provide details of the hardware and software used in our validation tests.

## Hardware Resources

The configuration described in this paper uses the following equipment:

| Description | Minimum Requirement |
|---|---|
| CLARiiON CX4-240 storage array | CLARiiON storage, iSCSI, FC LUNs and Snaps |
| CLARiiON write cache | 480MB |
| 2 Fibre Channel ports for back-end storage connectivity | |
| 4 10/100/1000 BaseT Ethernet ports | 2 per storage processor |
| 4 10/100/1000 management ports | 2 per storage processor |
| PROM revision | Release 1.81.00 |
| 45 300GB 15K 4Gb Fibre Channel disks | Supports up to 240 FC or SATA disks in an all or mixed configuration |
| 8 additional Fibre Channel ports for host connectivity | |

**Table 1: Hardware configuration**

## Software Resources

The configuration described in this paper uses the following software:

| Description | Minimum Requirement |
|---|---|
| **CX4-240** | |
| CLARiiON | Release 28 (04.28.000.005) |
| CLARiiON Navisphere® | 6.28.0.4.31 |
| **VMware ESX hosts** | ESX 3.5 Update 2 |
| **VMware vCenter** | |
| Operating system | Microsoft Windows Server 2003 Enterprise Edition SP2 (32-bit) |
| VMware vCenter | 2.5 Update 3 |
| **Desktops (virtual machines)** | |
| Operating system | Microsoft Windows XP Professional SP3 (32-bit) |
| **VMware Tools** | 3.5.0000 |

**Table 2: Software resources**

# Solution Configuration

The following sections provide details of the configuration of the environment used in our validation tests.

## Network Architecture

The networks in the configuration we tested were dedicated 1Gb Ethernet. We used a DHCP server to assign IP addresses to all VMware View desktops. Each ESX host used six 1Gb Ethernet controllers. We configured two of them as NIC teaming ports for iSCSI traffic.

We recommend that the switches support Gigabit Ethernet connections and the ports on the switches support copper-based media. In this configuration, the VMware virtual switches are associated with physical network cards.

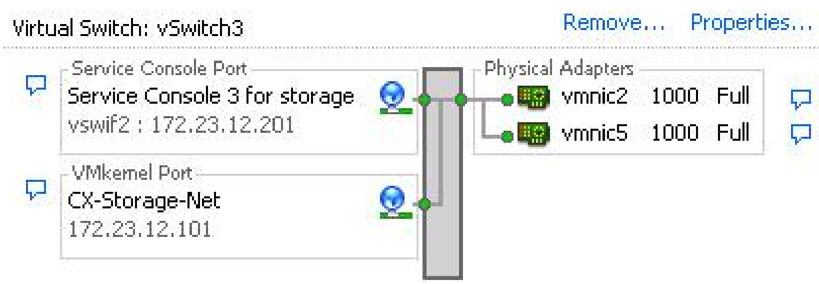


**Figure 1: VMware ESX NIC configuration**

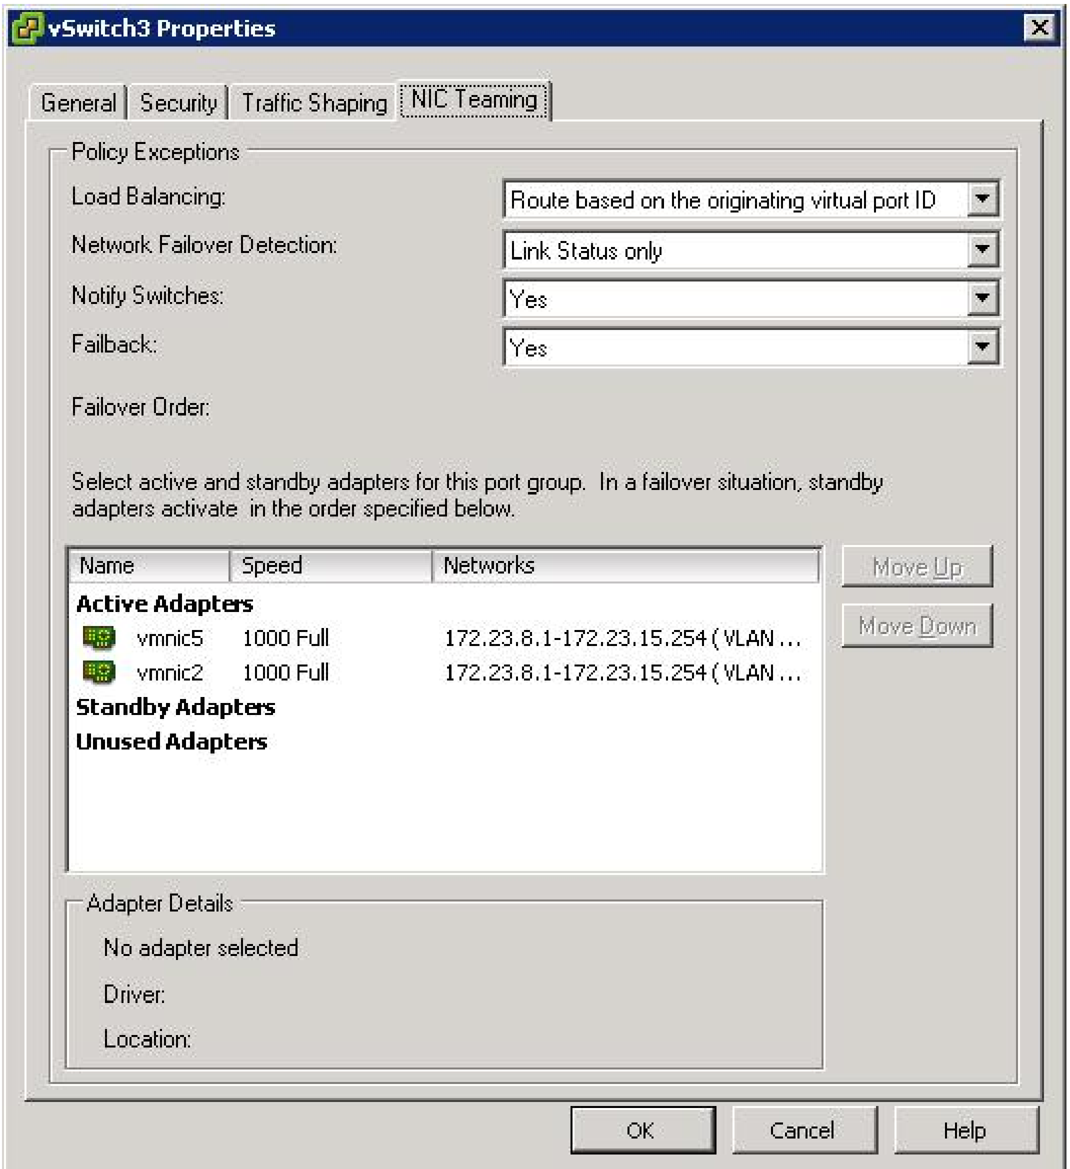


**Figure 2: VMware ESX NIC teaming configuration**

## EMC CLARiiON CX4-240

Figure 3 shows the ports on the rear of an EMC CLARiiON CX4-240.

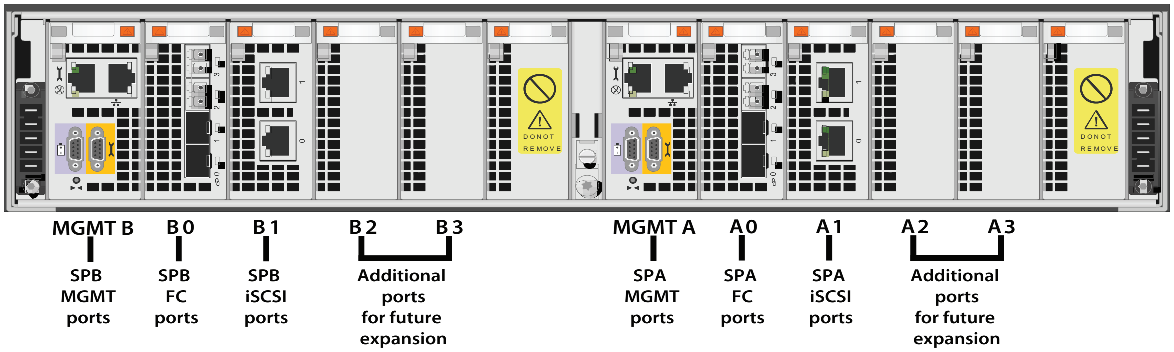


**Figure 3: EMC CLARiiON CX4-240 ports**

## EMC CLARiiON Configuration

The CLARiiON storage array used in this test has two Gigabit Ethernet ports on each storage processor. Currently, CLARiiON storage arrays do not support NIC teaming, so we split the workload manually by pointing some of the ESX hosts to each NIC. ESX handles the failover automatically because it identifies multiple paths to the storage array. Each ESX host identifies a total four paths to the CLARiiON storage array, two into each storage processor.

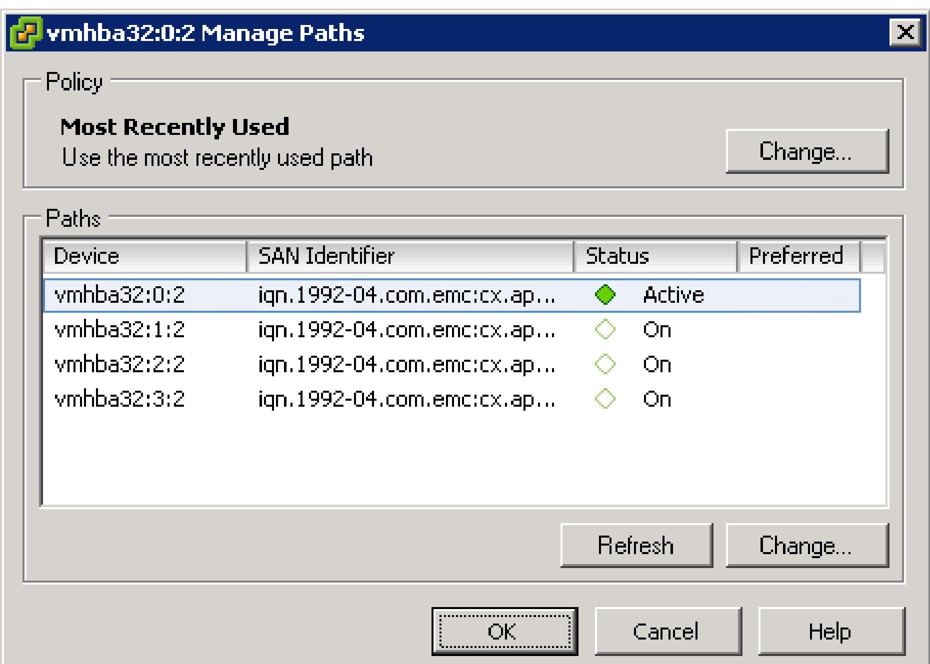


**Figure 4: Multipath configuration on the ESX Host**

## VMware ESX Configuration

You must configure each ESX host to allow for iSCSI access using the procedure below. Connect to VMware vCenter using the VI Client and for each host set the following parameters:

- **Configuration** tab > **Security Profile** > **Properties** > select **Software iSCSI client**
- **Configuration** tab > **Networking** > **Add networking** > **VmKernel** > **Create a new switch**
    - Enter a Network Label — Example: `Storage-Net`
    - Deselect the VLAN tag; CLARiiON does not support VLANs
    - Enter the IP address and subnet mask
- **Configuration** tab > **Storage Adapters** > **Select the iSCSI Adapter** > **Properties**
    - **Configure** > **Enable**
    - **Dynamic Discovery** tab > Add the IP address and port for each iSCSI target

**Note:** If CHAP authentication is enabled on the iSCSI target, you should also configure it using the **CHAP Authentication** tab.

## Storage Architecture

We configured the CLARiiON CX4-240 array as illustrated in Figure 5. This CX4 array had four disk array enclosures, each containing 15 Fibre Channel 300GB 15K disks. We used three of the disk array enclosures for this testing. All testing used a 4+1 RAID 5 disk grouping only.

| CLARiiON array objects | Configuration required |
|---|---|
| Fibre Channel disk capacity | 300GB |
| Number disks used | 35 |
| **VMFS for linked clone Images** | |
| Storage capacity | 14 × 300GB |
| Number of iSCSI LUNs | 14 |
| iSCSI LUN capacity | 300GB |
| Number RAID groups used | 7 × (4+1, RAID-5) |

**Table 4: Storage layout**

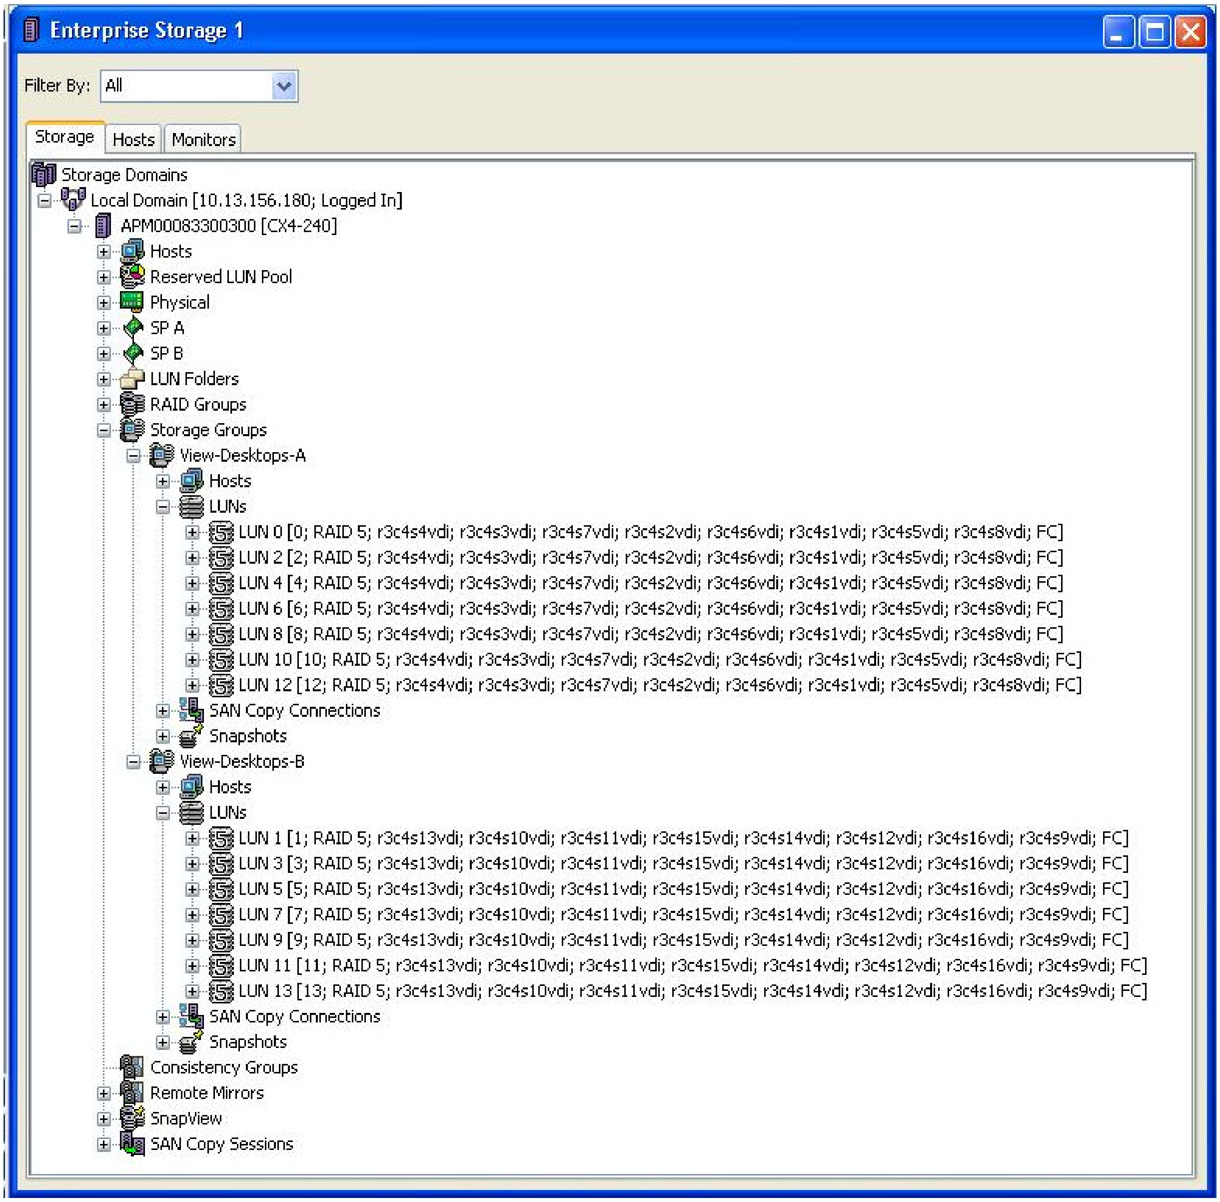


**Figure 5: Configuration of storage groups on CLARiiON**

The set of 35 300GB/15K Fibre Channel disk drives was configured with the RAID disk groups and layout shown in Figure 6.

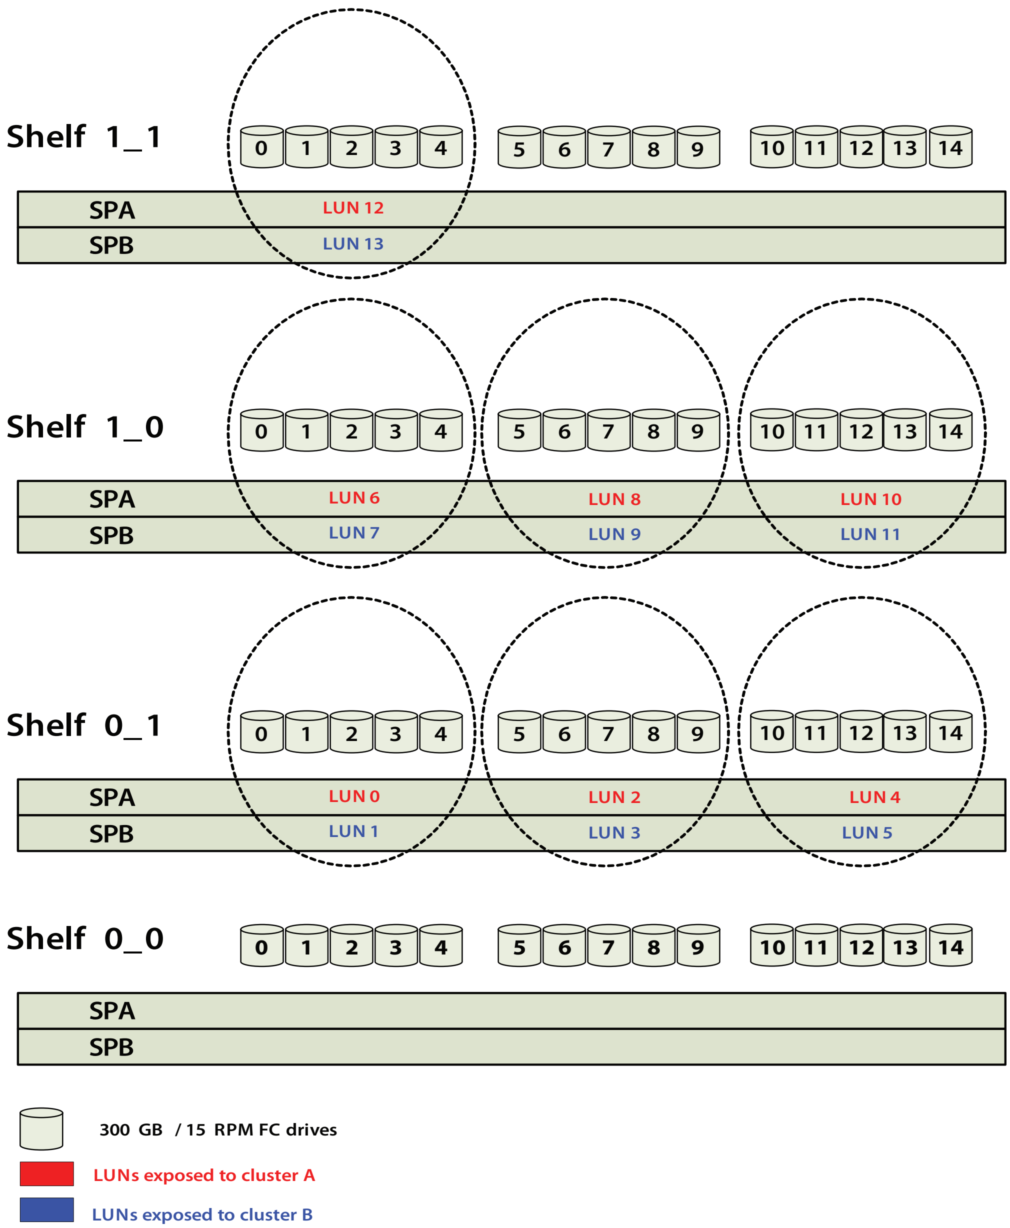


**Figure 6: EMC CX4-240 RAID group configuration**

We presented all the odd numbered LUNs of the 14 available via VMFS to one of two VMware ESX clusters as a datastore where the linked clones resided. We exposed the even numbered LUNs to the other ESX cluster.

We used VMware View Manager 3 to create each of the desktop pools. Each desktop pool used the seven available datastores assigned to its respective cluster. We used the following formula to estimate the size needed for each LUN. In all other calculations we rounded numbers up or down for simplicity.

(Number of clones * 2x memory) + (Number of patch replicas * virtual machine disk size) =Total amount of usable space needed

(1000 * 1) + (4 * 8) = 1032

Even though we could create this amount of storage — approximately 1TB — on fewer spindles, for appropriate storage sizing, we needed to take performance requirements into consideration. Rather than just considering the performance requirements during the system's steady state, we considered the boot-storm performance requirements when we sized the appropriate number of spindles.

In order to accommodate the desktop pools (which were based on linked clones) we configured each pool to use aggressive storage overcommitment. Storage overcommitment is a feature of VMware View Composer that allows you to control how aggressively View places virtual machines on each datastore. When you use a more aggressive level of storage overcommitment, you have less free space available for virtual machines to grow over time. For more information about VMware View Composer, see the "VMware View Manager Administration Guide" or "Design Considerations for VMware View Composer" (see the Resources section for links).

This formula for estimating the required storage assumes that you are using aggressive storage overcommitment. When you use aggressive storage overcommitment, you have very little additional storage available to provide room for growth, over time, for linked clones. For persistent pools, you should implement a refresh policy that resets the virtual machines to their original size. For nonpersistent pools, implement a policy of deleting after first use. As an alternative, you can add additional storage to the above formula to provide additional room for growth.

Each of the disk groups in the configuration described in this paper provides more than enough additional space to accommodate increasing the size of each of the iSCSI LUNs at any time. By increasing the size of each iSCSI LUN, you can provide additional room for each of the virtual machines based on a linked clone to grow over time.

Regardless of the approach you take, it is a best practice to monitor the storage array for space usage. Table 5 lists disk volumes per file system for the storage configuration described in this paper.

| File System | LUNs |
|---|---|
| Virtual machines (linked clones) cluster A | LUN0, LUN2, LUN4, LUN6, LUN8, LUN10, LUN12 |
| Virtual machines (linked clones) cluster B | LUN1, LUN3, LUN5, LUN7, LUN9, LUN11, LUN13 |

**Table 5: LUNs**

# Validation Results

The configuration we implemented and used during the validation of the VMware View reference architecture is suitable for large-scale VMware View deployments. In this configuration, we used the EMC CLARiiON CX4-240 to validate a 1000-user VMware View building block architecture. For additional information, see "VMware View Reference Architecture" (see the Resources section for a link).

## Storage System Utilization

Figures 7 and 8 show the average CPU utilization of the EMC CLARiiON CX4-240 and the average I/O rate for both of the EMC CLARiiON storage processors. As the graphs show, this configuration provides more than enough capacity to handle the resource needs of each building block. The average I/O rate into both of the storage processors was less than 13MB/s during the steady-state stage of testing.

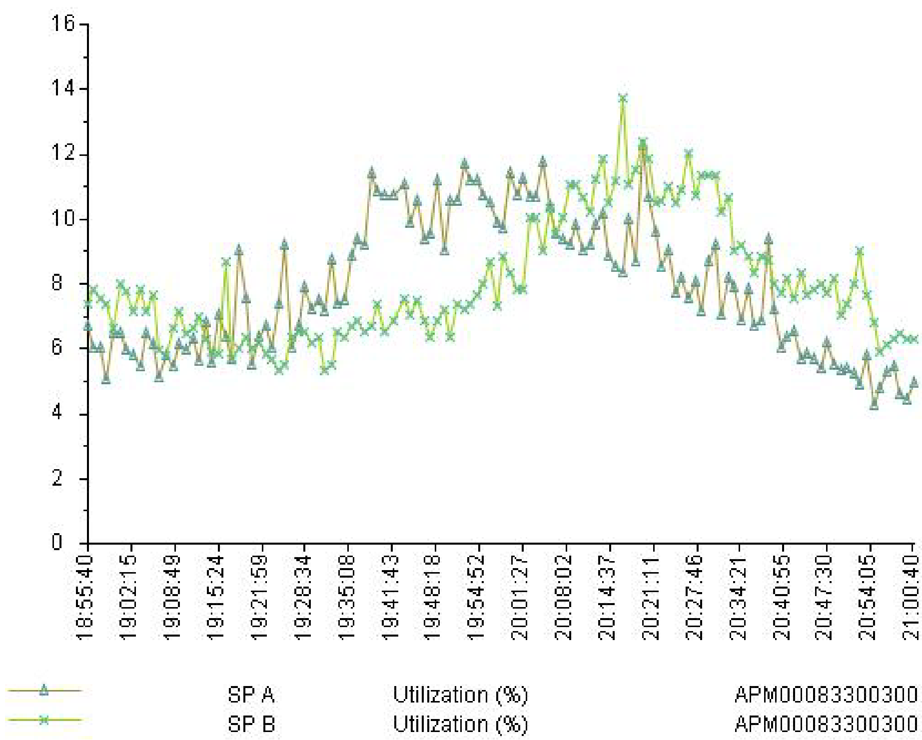


**Figure 7: CLARiiON CX4-240 average CPU utilization**

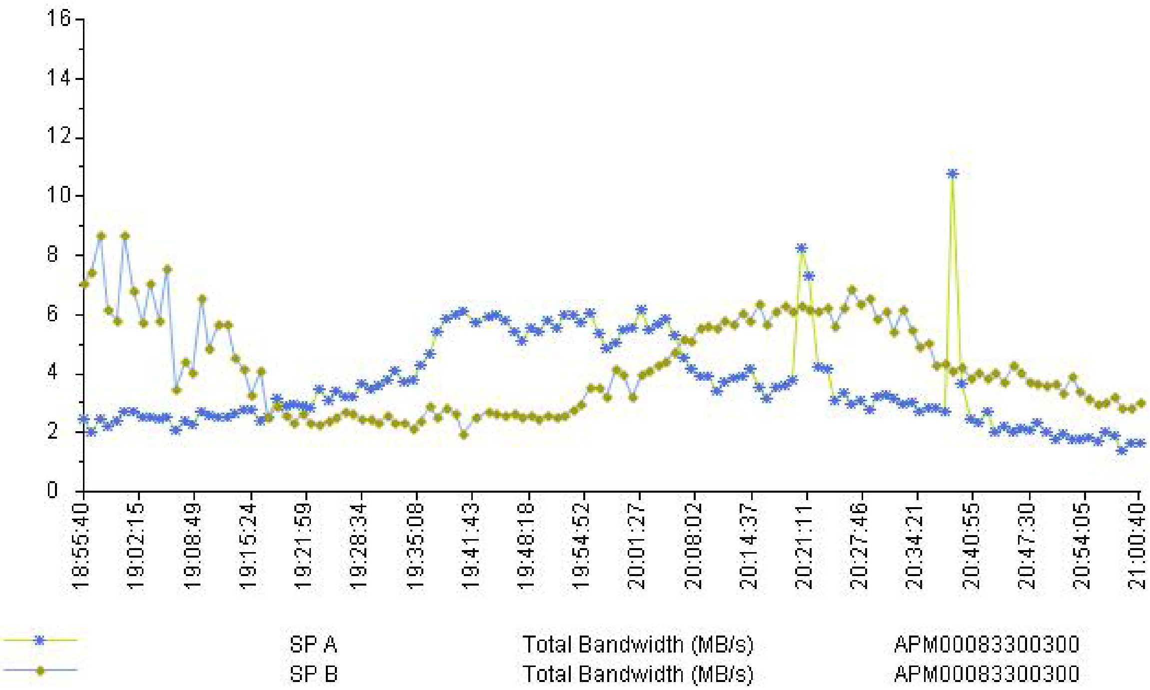


**Figure 8: EMC CLARiiON CX4-240 average I/O rate for all disk volumes**

The workload consisted of 1000 concurrent Windows desktop users performing the normal operations listed below. These users can be classified as high-end task workers or low-end knowledge workers. The workload includes the following tasks:

- Microsoft Word

  Open, minimize, close, write random words and numbers, save modifications

- Microsoft Excel

  Open, minimize, close, write random numbers, insert and delete columns and rows, copy and paste formulas, save modifications

- Microsoft PowerPoint

  Open, minimize, close, conduct a slide show presentation

- Adobe Acrobat Reader

  Open, minimize, close, browse pages in PDF document

- Internet Explorer

  Open, minimize, close, browse page

- McAfee Active VirusScan

  Real-time scanning

- PKZIP

  Open, close, compress a large file

## Application Response Time

Figure 9 shows the average application execution time for all virtual desktops in both cluster A and cluster B. These application times represent the amount of time it took to open a document, close a document, or save a document that was created. Figure 7 does not show the amount of time an application is minimized or being worked on. Because of the random nature of the workload, applications are not just opened, worked on, and then closed. Rather, the workload might create a Microsoft Word document, work on the document, minimize the document, and then use Microsoft Internet Explorer. At some later point, the workload returns to the Microsoft Word document (which had been minimized), reopens the document, works on it again, and then closes it.

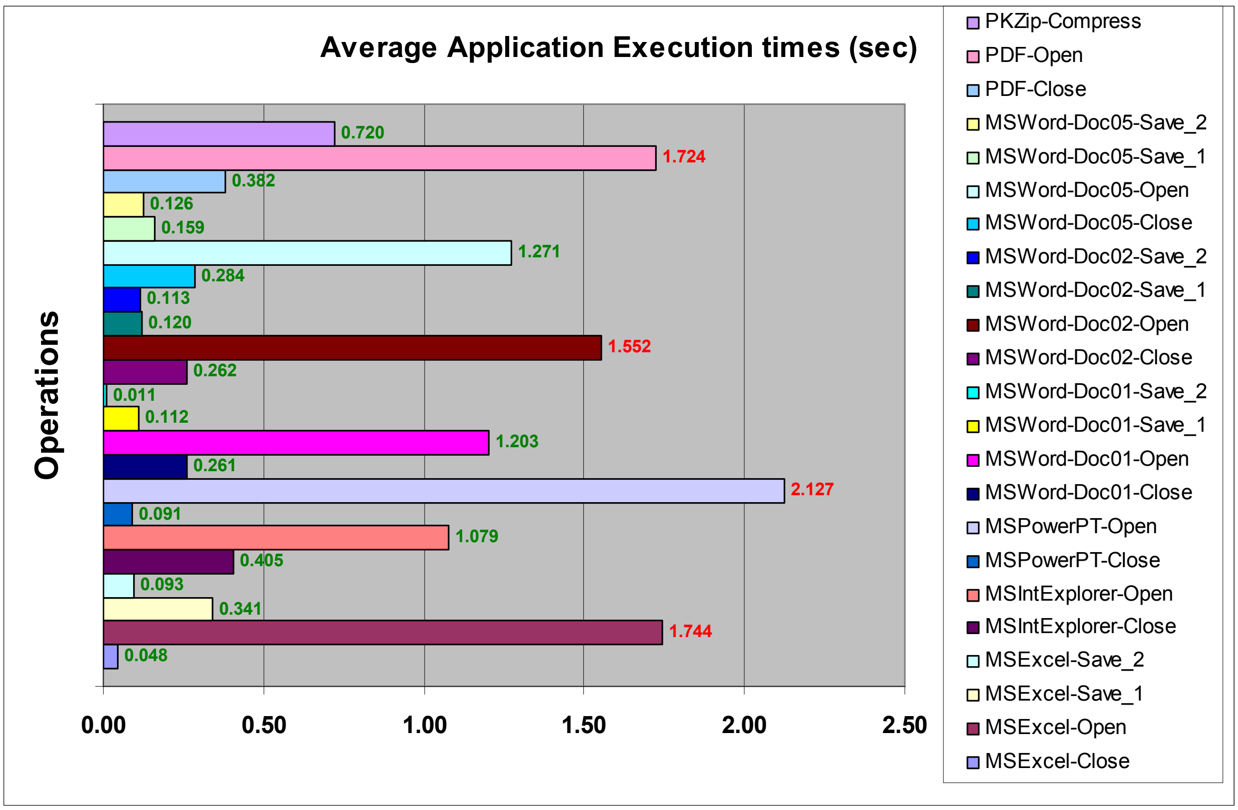


**Figure 9: Average application execution time in seconds**

Figure 10 shows the corresponding latency at the storage system level during the steady-state stage of testing.

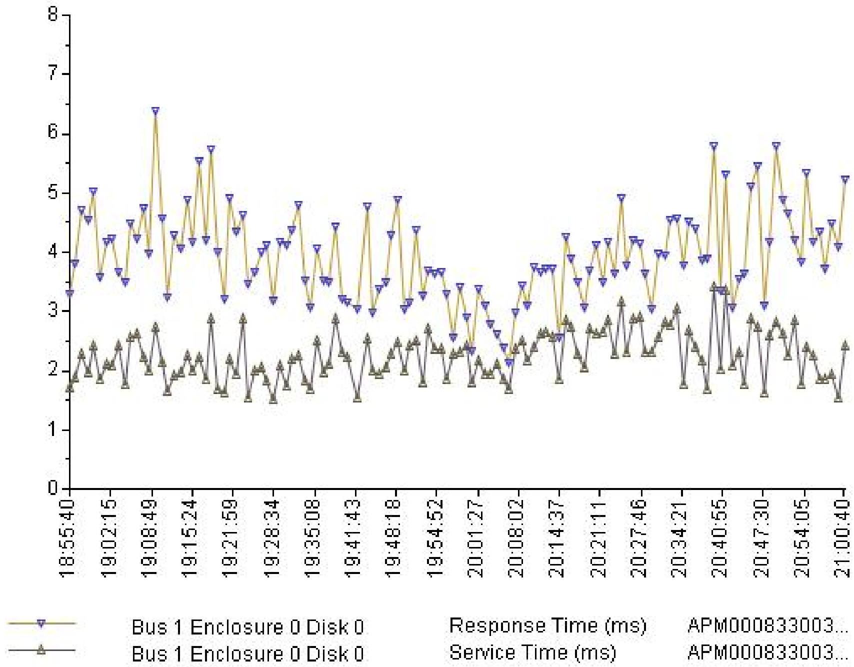


**Figure 10: Average storage latency**

# Summary

The storage configuration used in our validation testing can easily support 1000 virtual desktops that can be used by users who fit the high-end task worker profile. The application response times were well within the acceptable limits. The steady storage latencies were well below 6ms with some headroom to absorb an unexpected boot storm surge.

# Resources

- "Design Considerations for VMware View Composer"
  http://www.vmware.com/resources/techresources/10004
- VMware Infrastructure 3 documentation
  http://www.vmware.com/support/pubs/vi_pubs.html
- "VMware View Manager Administration Guide"
  http://www.vmware.com/pdf/viewmanager3_admin_guide.pdf
- "VMware View Reference Architecture"
  http://www.vmware.com/files/pdf/resources/vmware-view-reference-architecture.pdf
- "Windows XP Deployment Guide"
  http://www.vmware.com/files/pdf/vdi-xp-guide.pdf

# Appendix: Glossary

- iSCSI — Internet SCSI protocol
- iSCSI target — An iSCSI end point, identified by a unique iSCSI name, which executes commands issued by the iSCSI initiator.
- LUN — Logical unit: For Fibre Channel on a CLARiiON storage array, a logical unit is a software feature that processes SCSI commands, such as reading from and writing to storage media. From a Fibre Channel host perspective, a logical unit appears as a block-based device.
- RAID — Redundant array of independent disks, designed for fault tolerance and performance.
- VMFS — VMware Virtual Machine File System.
- VMware View — A set of software products that provide services and management infrastructure for centralization of desktop operating environments using virtual machine technology.
- VMware View Manager — A software product that manages secure access to virtual desktops. It works with vCenter Server to provide advanced management capabilities.

Revision: 20090415 Item: DG-0904NF-01-01

VMware, Inc. 3401 Hillview Ave. Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com