DEPLOYMENT GUIDE

# EMC Celerra Unified Storage

*A Guide to deploying EMC Celerra NS20 storage with VMware View*

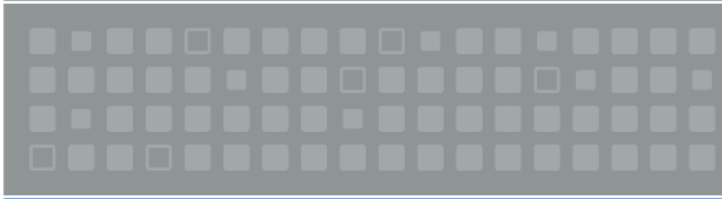

vmware®

EMC²®

where information lives®

# Contents


**Introduction to VMware View ..... 2**

**About This Guide ..... 2**

**Hardware and Software requirements ..... 3**

Terminology ..... 3

Hardware Resources ..... 4

Software Resources ..... 4

**Solution Configuration ..... 5**

Network Architecture ..... 5

EMC Celerra NS20FC ..... 7

EMC Celerra Configuration ..... 7

VMware ESX Configuration ..... 12

Storage Architecture ..... 13

**Validation Results ..... 16**

Storage System Utilization ..... 16

Application Response Time ..... 17

**References ..... 18**

# Introduction to VMware View

Built on VMware's industry-leading 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 customers to extend the value of VMware Infrastructure and virtual desktop infrastructure (VDI) environments 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 customers use and manage desktop operating systems. Desktop instances can be rapidly deployed in secure data centers 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.

# About This Guide

This document provides a detailed summary and characterization of designing and configuring the EMC Celerra NS20FC for use with VMware View. It describes a validated configuration for a 1000 user workload – a "1000 user building block". Larger configurations can be deployed using multiples of the server and storage configuration. This configuration was used as part of the VMware View reference architecture validation. This guide is intended to offer Architects and Desktop Administrators guidance specific to the use of EMC Celerra NS20FC storage for use with a VMware View based solution. Although this configuration was used specifically with the VMware View reference architecture for large scale VMware View deployments the information provided in this guide also can be helpful to anyone looking to deploy VMware View using an EMC Celerra NS20FC. In addition these guidelines also directly apply to the entry-level NS20 system and can be further extrapolated to other larger scale Celerra systems. The intended audiences of this document should have a working knowledge of VMware View as well as EMC Clariion and EMC Celerra technologies.

# Hardware and Software requirements

## Terminology

| Term | Definition |
|---|---|
| DAE | Disk Array Enclosure |
| Data Mover | X-Blade 20. The NAS Blade. Networked storage server that runs the Data Access in Real Time (DART) operating system to process storage requests. |
| iSCSI | Internet SCSI protocol |
| iSCSI target | An iSCSI endpoint, identified by a unique iSCSI name, which executes commands issued by the iSCSI initiator. |
| Link Aggregation | A high-availability featured based on the IEEE 802.3ad Link Aggregation Control Protocol (LACP) standard allowing Ethernet ports with similar characteristics on the same switch to combine into a single virtual device/link with a single MAC address and potentially multiple IP addresses. |
| LUN | Logical Unit: For iSCSI on a Celerra Network Server, a logical unit is an iSCSI software feature that processes SCSI commands, such as reading from and writing to storage media. From an iSCSI host perspective, a logical unit appears as a block-based device. |
| RAID | Redundant Array of Independent Disks, designed for fault tolerance and performance |
| Temporary Writable Snaps | A writeable point-in-time copy of an iSCSI LUN or Filesystem on a Celerra that can be promoted to be accessed by a host. The Celerra snapshot technology consumes a small fraction of the space of the primary object (LUN or Filesystem). |
| VMware View Manager | VMware View Manager: Manages secure access to virtual desktops, works with VMware vCenter to provide advanced management capabilities |
| VC | VMware vCenter |
| VMware View | A set of software products that provide services and management infrastructure for centralization of desktop operating environments using virtual machine technology. |
| VMFS | VMware's Virtual Machine File System |
| Virtual(thin) provisioning | Consolidation and automated process of allocating just "the exact required amount" of server space at the time it is required. It operates by allocating disk storage space in a flexible manner among multiple users, based on the minimum space required by each user at any given time. As more storage is needed additional volumes can be automatically added. |

Table 1: Glossary of terms

## Hardware Resources

The following equipment was used in this configuration:

| Description | Minimum Revision |
|---|---|
| Celerra NS20FC with a Clariion CX3-10F backend array | Celerra Shared Storage; Filesystems, iSCSI LUNS and Snaps |
| CLARiiON Write Cache | 259 MB (default) |
| Two X-Blade 20 configurations: | Configured in a Primary/Standby configuration for High Availability |
| Dual 2.8 GHz Pentium IV CPUs | |
| 4 GB Double Data Rate RAM (266 MHz) | |
| 2 Fibre Channel ports for back-end storage connectivity | |
| 2 Fibre Channel ports for tape connectivity | |
| 4 10/100/1000 BaseT Ethernet ports | Configured for two 2 port link aggregation devices running LACP protocol |
| 1 10/100/1000 management port | |
| BIOS | Release 3.58 |
| POST | Release 1.50 |
| 30 300 GB /15K 2/4 Gb Fibre Channel disks | supports 6 to 90 FC or SATA disks in an all or mixed configuration |
| 8 additional Fibre Channel ports for host connectivity | |

Table 2: Hardware Configuration

## Software Resources

The following software was used in the configuration:

| Description | Minimum Revision |
|---|---|
| NS20FC: | Celerra Shared Storage; Filesystems, iSCSI LUNS |
| NAS/DART | 5.6 Maintenance Update 3 (5.6.40.3) |
| CLARiiON Flare | Release 26 (3.26.10.5.010) |
| CLARiiON Navisphere | 6.26.9 (0.58) |
| VMware ESX Servers | ESX 3.5 Update 2 |
| VMware vCenter | |

| | |
|---|---|
| OS | MS Windows Server 2003 Enterprise Edition SP2 (32-bit) |
| VMware vCenter | 2.5 Update 3 |
| Desktops/Virtual Machines | |
| OS | MS Windows XP Professional Version SP3 (32-bit) |
| VMware Tools | 3.5.0000 |

Table 3: Software Resources

# Solution Configuration

## Network Architecture

The networks used were dedicated 1 Gb Ethernet. All virtual desktops were assigned an IP address using a DHCP server. The VMware ESX Servers consisted of six Gb Ethernet Controllers. Two were configured as NIC Teaming ports for iSCSI traffic. The two NS20FC X-Blade 20s were configured in an Active/Passive state where the second X-Blade 20 serves as a standby in case of a failover event.

The network configuration on the X-Blade 20 is as follows. It employs link aggregation devices using the LACP protocol with VLAN tagging enabled. There are two network interfaces per device, one subnet for iSCSI storage traffic and the other for CIFS traffic.

```
# /nas/bin/server_ifconfig server_2 –a server_2 : ns20a-lnk2b-
server2 protocol=IP device=lnk2 inet=172.16.10.26
netmask=255.255.0.0 broadcast=172.16.255.255 UP, ethernet,
mtu=1500, vlan=16, macaddr=0:60:16:c:4f:de

ns20a-lnk2a-server2 protocol=IP device=lnk2 inet=172.23.11.27
netmask=255.255.0.0 broadcast=172.23.255.255 UP, ethernet,
mtu=1500, vlan=23, macaddr=0:60:16:c:4f:de ns20a-lnk1b-server2
protocol=IP device=lnk1 inet=172.16.10.25 netmask=255.255.0.0
broadcast=172.16.255.255 UP, ethernet, mtu=1500, vlan=16,
macaddr=0:60:16:c:4f:e0

ns20a-lnk1a-server2 protocol=IP device=lnk1 inet=172.23.11.26
netmask=255.255.0.0 broadcast=172.23.255.255 UP, ethernet,
mtu=1500, vlan=23, macaddr=0:60:16:c:4f:e0
```

This can be seen in the Celerra Manager as well:

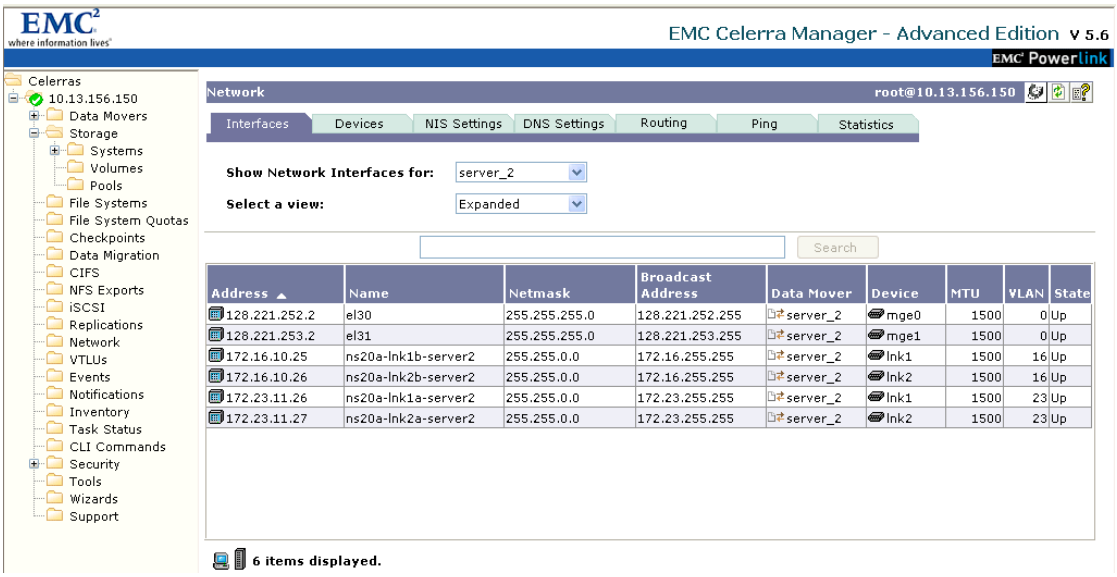


It's recommended that the switches support Gigabit Ethernet (GbE) connections and the ports on the switches support copper-based media. In this configuration, the VMware virtual switches are associated with physical network cards.

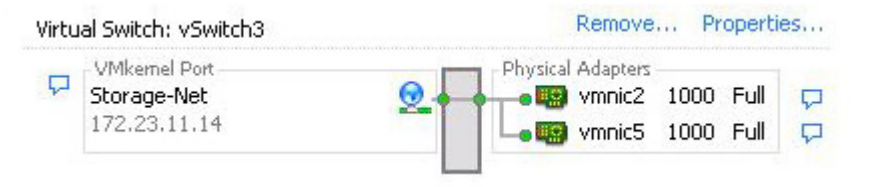


Figure 1: VMware ESX server NIC configuration

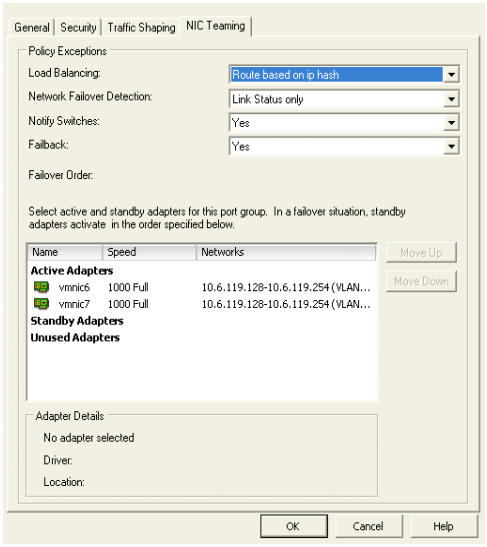


Figure 2: Example VMware ESX NIC teaming Configuration

## EMC Celerra NS20FC

The NS20FC contains two X-Blade 20s. The X-Blade 20s can operate independently, or they can operate in the active/passive mode, with the passive X-Blade 20 serving as a failover device for the active X-Blade 20. In this solution, the X-Blade 20s operate in active/passive mode.

The NS20FC X-Blade 20 consists of four Gb Ethernet Controller ports. These four ports were configured as two 2-port link aggregation devices. Each link aggregation device was placed on a different subnet in order to create multiple paths for the iSCSI objects. Multiple iSCSI targets were created and iSCSI sessions were distributed across both logical network interfaces.

The following figure shows the ports on the rear of an EMC Celerra NS20FC X-Blade 20.

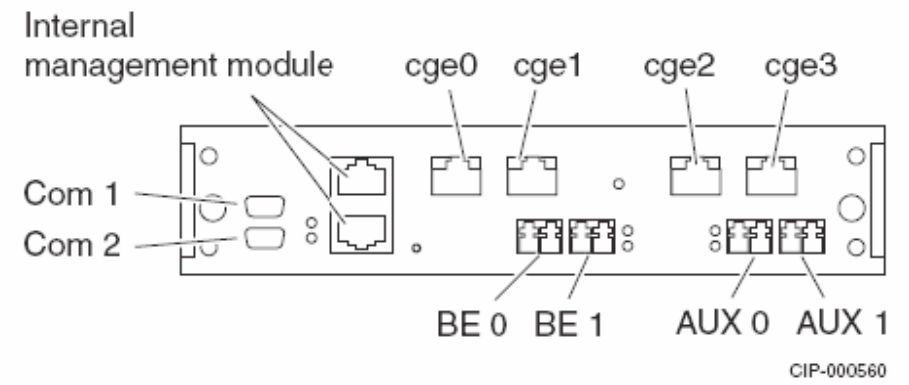


Figure 3: EMC Celerra NS20FC X-Blade 20 Ports

Ports cge0 and cge1 are set up for link aggregation and support the iSCSI storage traffic. Ports cge2 and cge3 are used for a second link aggregation device.

Note: As a best practice, the X-Blade 20 network ports connected to the storage network (as shown in the figure above) should be dedicated to storage traffic. However, if the ports are not heavily used, they can be shared with non-storage network traffic. EMC recommends monitoring the network to avoid bottlenecks.

## EMC Celerra Configuration

The required EMC NS20FC configuration is as follows:

Note: These configurations can be implemented using control station Celerra commands or using the Celerra Manager web-based interface. Both are shown below for completeness.

To generate the list of only iSCSI targets, for which a host has been explicitly granted access to a LUN, use the following command:

```
/nas/bin/server_param <server_name> -facility iscsi –modify SendTargetsMode –
value 1
```

Or from the following page in the Celerra Manager:

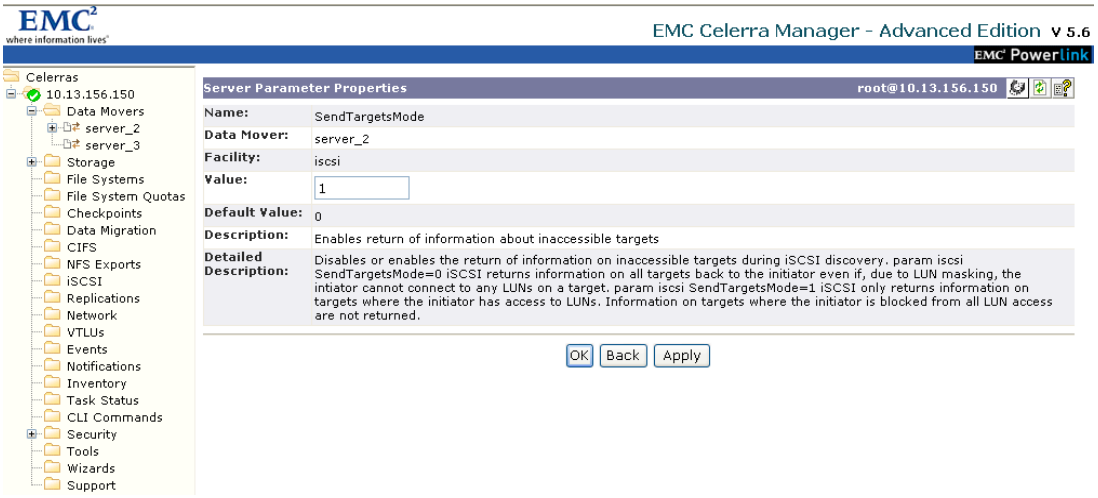


Information will not be returned on targets where the initiator is blocked from all LUN access.

To create a user-defined storage pool, use the command:

```
/nas/bin/nas_pool -create -name <pool name> -description 'Storage Pool' -
volumes <dvol…>, -default_slice_flag y
```

Or from the following page in the Celerra Manager:

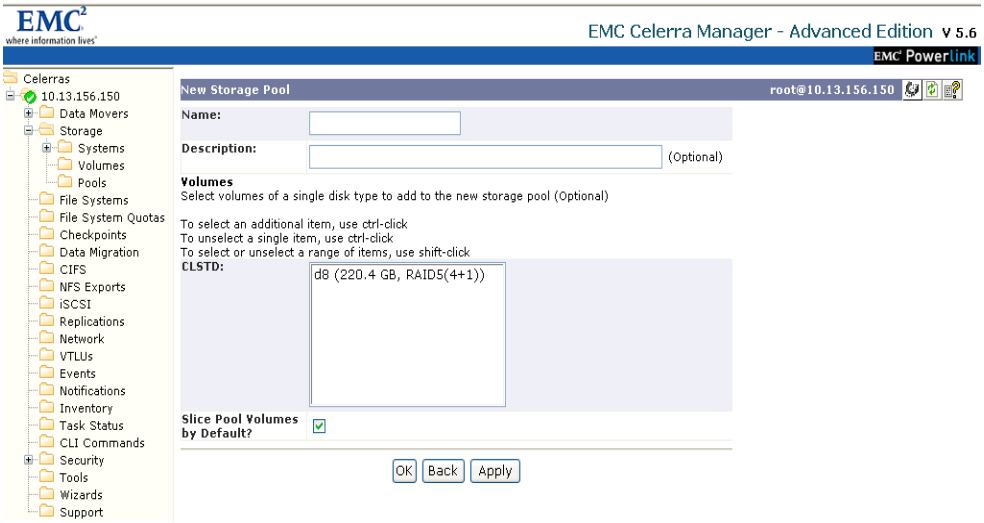


Note: One can specify one or many disk volumes to define the storage pool.

To create a file system with the user-defined storage pool and then mount it on a X-Blade 20, use the following commands:

```
/nas/bin/nas_fs -name <fs name> -type uxfs -create size=<size> pool=<pool name> -
auto_extend yes -vp yes -hwm <value>% -option mover=<server_name>, slice=y

/nas/bin/server_mount <server_name> <fs name> <fs pathname>
```

Or in the Celerra Manager with:

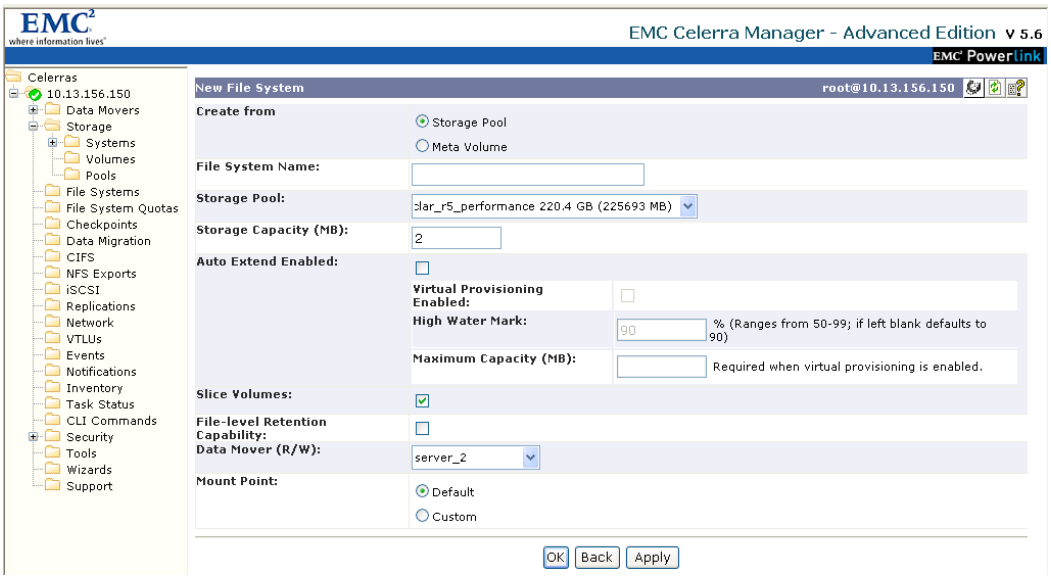


To create an iSCSI LUN serving as a data store to the ESX server, use the following command:

```
/nas/bin/server_iscsi <server_name> -lun -number <lun #> -create <target alias
name> -size <lun size> -fs <fs name> -vp yes
```

Or in the Celerra Manager with:

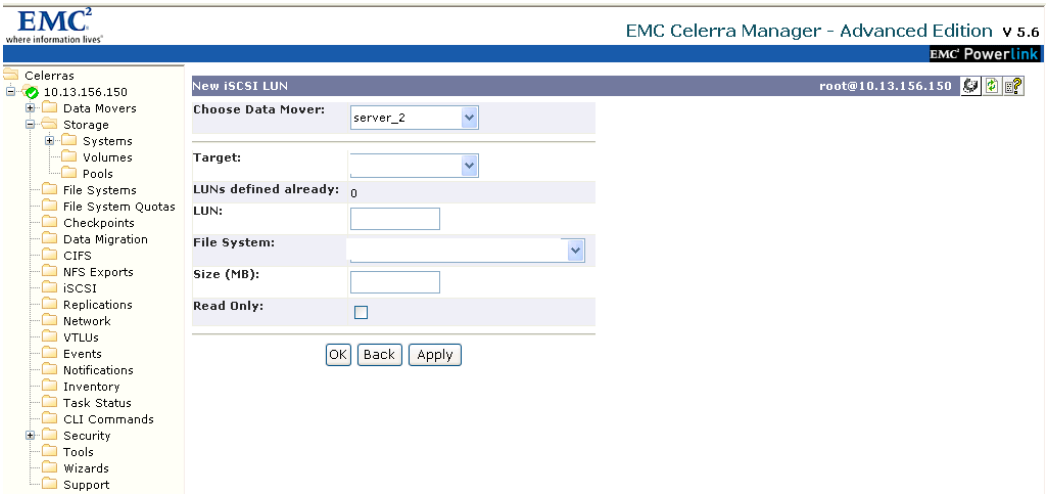


The **-vp yes** option allows for the creation of a virtually provisioned iSCSI LUN instead of a regular (thick) LUN. When using virtual provisioning, closely monitor file system space that contains virtually provisioned iSCSI LUNs. You can use the Celerra Manager or use the following commands to determine used or available space in a file system:

```
/nas/bin/server_df
```

```
/nas/bin/nas_fs
```

Below is an example of using both the Celerra storage commands and the Celerra Manager for creating the iSCSI LUNS for the VMware View environment.

To create storage pools containing one 4+1 RAID 5 disk group, perform the following:

```
/nas/bin/nas_pool -create -name Pool1  -description "Pool for
linked clones"   -volumes d11,d14  -default_slice_flag y
```

Or in the Celerra Manager with:

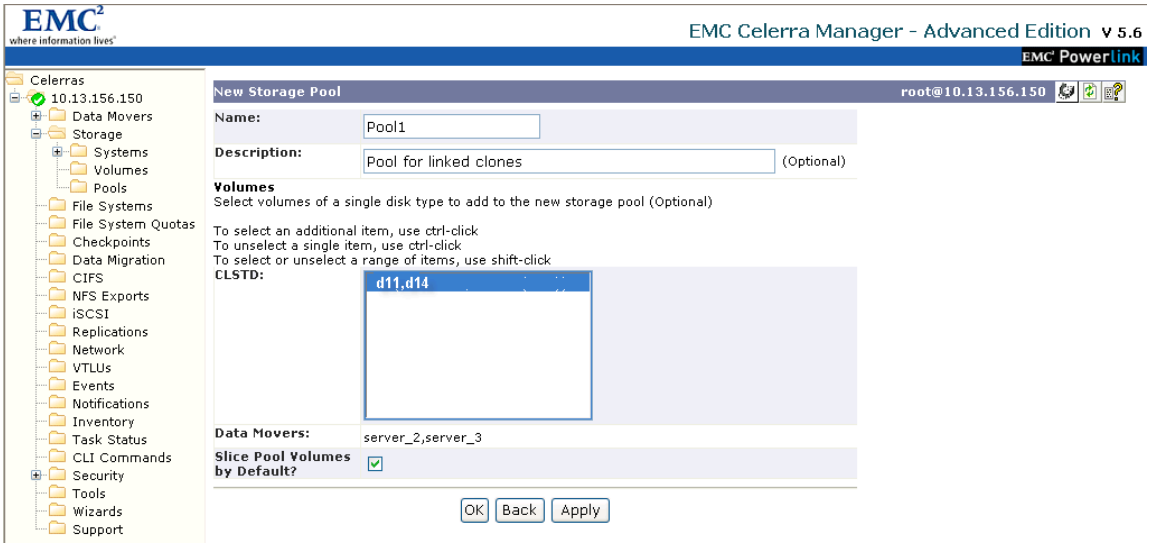

Create the file system and mount on the X-Blade:

```
/nas/bin/nas_fs -name fs1 -type uxfs -create size=82G pool=Pool1
-auto_extend yes -vp yes -hwm 75% -max_size 84G -option
mover=server_2,slice=y
```

```
/nas/bin/server_mount server_2 fs1 /fs1
```

Or in the Celerra Manager with:

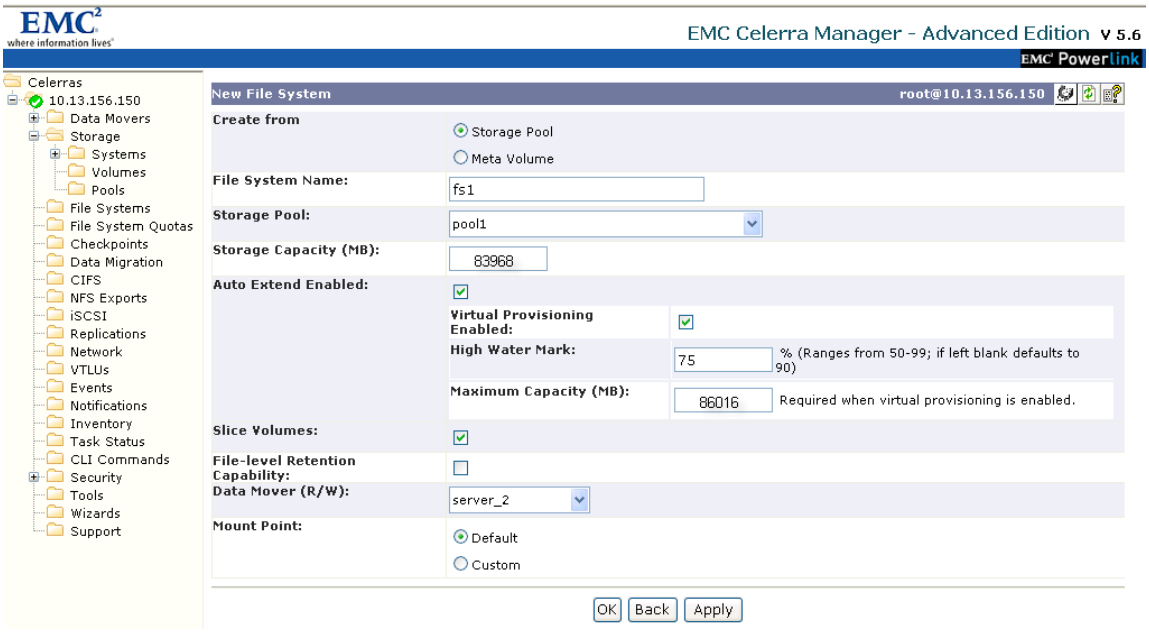


After creating an iSCSI target on the X-Blade so the initiator can establish an iSCSI session, create the required iSCSI LUN that will serve as the ESX Data Store:

```
/nas/bin/server_iscsi server_2 -lun -number 0  -create <target alias
name>   -size 80G -fs fs1 -vp yes
```

Or in the Celerra Manager with:

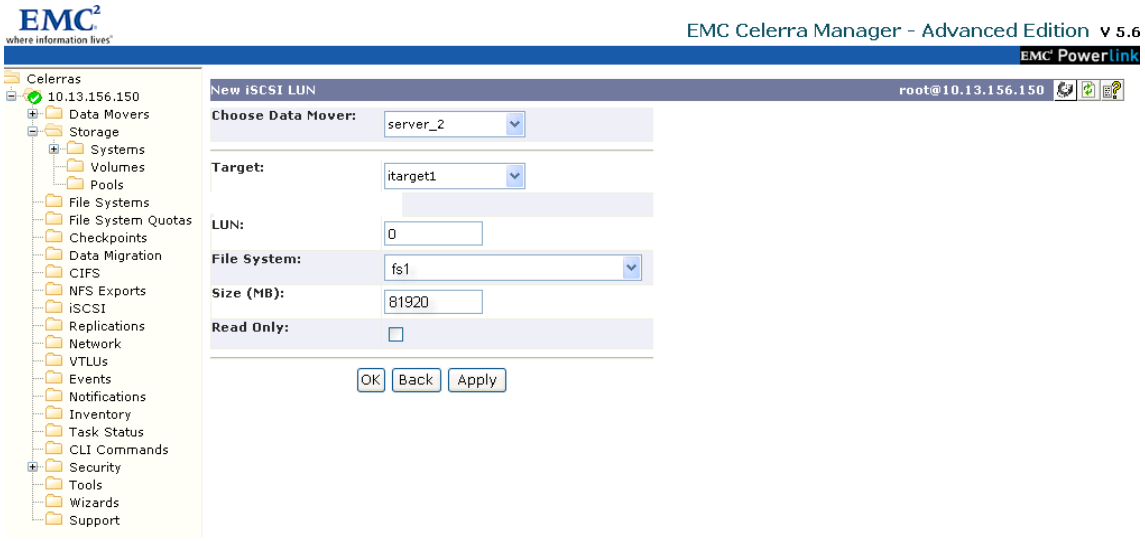


## VMware ESX Configuration

Each VMware ESX server must be configured to allow for iSCSI access To provide iSCSI access, connect to VMware VirtualCenter using the VI client (VIC) and for each host set the following parameters:

- **Configuration Tab** > S**ecurity Profile** > **Properties** > (check) **Software iSCSI clien**t
- Configuration Tab > Networking > Add networking > VmKernel > Create a New switch
    - Provide a Network Label – Example: Storage-Net
    - Set a VLAN tag (Optional)
    - Provide 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 it should also be configured using the CHAP Authentication tab.

## Storage Architecture

The EMC Celerra NS20FC backend array was configured as illustrated in the figure on the next page. A CX3-10f is the backend array for a Celerra NS20FC. This CX3-10f had two DAEs, each containing 15 Fibre Channel 300GB/15K 2/4 Gb disks. All testing used a 4+1 RAID 5 disk grouping only. According to the standard NAS template, 2 LUNS were created per 4+1 RAID group and each LUN was owned by a different SP for load balancing. File systems were created with a user defined storage pool in order to have direct control over file system location space allocation and load balancing.

| CLARiiON array objects | Configuration required |
|---|---|
| **File system for Master VM Images – Full Clone Images** | |
| Storage Capacity | 220 GB |
| iSCSI LUN Capacity | 100 GB |
| Number disks used | 5 |
| Number of disk volumes used | 1 |
| **File systems for Linked Clone Images** | |
| Storage Capacity | 1 TB |
| iSCSI LUN Capacity | 80GB |
| Number of disks used | 5 |
| Number of disk volumes used | 2 |
| **File systems for CIFS** | |
| Storage Capacity | 536GB |
| iSCSI LUN Capacity | 536GB |
| Number of disks used | 4 |
| Number of disk volumes used | 4 |

Table 4: File System Configuration

The storage array consisting of 2 DAEs of 30 300GB/15K Fibre Channel disk drives is configured with the following RAID disk groups and layout:

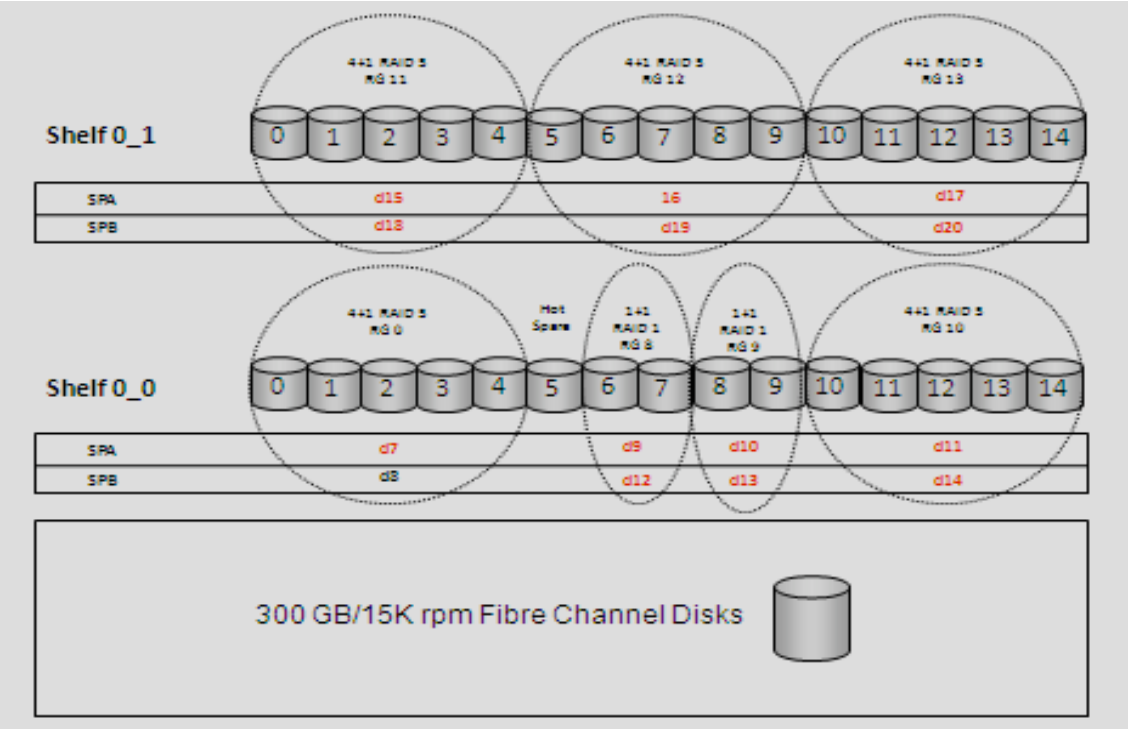


Figure 4: Example EMC NS20FC RAID Group configuration

| RAID Group ID | RAID Protection | Uses Disk Volumes | Storage Pool Name | Capacity | Used by File Systems | Description | iSCSI LUN Size |
|---|---|---|---|---|---|---|---|
| 0 | RAID 5 (4+1) | d7 | Golden_pool | 220.4 GB | goldenvm_fs | iSCSI LUN for Golden Image | 100 GB |
| 8 & 9 | RAID 1 | d9,d10,d12,d13 | UserData_pool | 536.7 GB | userdata_fs | File system for CIFS user data | 536 GB |
| 10 | RAID 5 (4+1) | d11,d14 | Pool 1 | 1 TB | fs1,fs5,fs9,fs13 | iSCSI LUNs on each file system | 80 GB |
| 11 | RAID 5 (4+1) | d15,d18 | Pool 2 | 1 TB | fs2,fs6,fs10, fs14 | iSCSI LUNs on each file system | 80 GB |
| 12 | RAID 5 (4+1) | d16,d19 | Pool 3 | 1 TB | fs3,fs7,fs11 | iSCSI LUNs on each filesystems | 80 GB |
| 13 | RAID 5 (4+1) | d17,d20 | Pool 4 | 1 TB | fs4,fs8,fs12 | iSCSI LUN on each filesystem | 80 GB |

Table 5: Storage Description

Note: the file systems and iSCSI LUNs were all virtually provisioned.

The User Data_fs file system was used as storage for the user data for each virtual desktop. It was a mapped network drive via the CIFS protocol on each desktop.

Each of the 14 desktop file systems named fs* has a capacity of 84 GB. On each of the 14 file systems, there was a 80GB iSCSI LUN that was presented via VMFS to one of two VMware ESX server clusters as a datastore where the linked clones resided.

The Golden Pool was used to store 25 virtual machines from two pools residing on Cluster A. These virtual machines were created as full clones.

VMware View Manger 3 was used to create each of the desktop pools. Each desktop pool leveraged the 7 available datastores assigned to its respective cluster. The following formula was used to estimate the size needed for each LUN. Although 25 of the 1000 virtual machines were full clones when planning the datastore size for use with linked clones based virtual machines, 1000 was used for simplicity. In all other calculations numbers were rounded up or down for simplicity.

*(number of clones * 2x memory) + (number of patch replicas * VM disk size) =Total Amount of usable space needed*

1000 * 1 + 4 * 8 = 1032

In order to accommodate linked clone based desktop pools, each pool was configured to use aggressive overcommit. Storage overcommit is a feature of VMware View Composer that allows an administrator to control how aggressively virtual machines are placed each datastore. When using a more aggressive level of storage overcommit their will less free space available for virtual machines to grow over time. For more information about VMware View Composer please see the VMware View Manager Administrator Guide or the VMware View Composer deployment guide.

It's important to note when using this formula to estimate the required storage, an assumption is made that aggressive overcommit will be used. When using aggressive storage over commit, there will be very little additional storage available to provide room for growth, over time, for the linked clones. For persistent pools a refresh policy should be implemented that resets the virtual machines back to their original size. For non-persistent pools a policy to delete after first use will need to be implemented. Alternatively, additional storage can easily be added to the above formula to provide additional room for growth.

Each of the disk groups above provides plenty of additional space to accommodate increasing the size of each of the iSCSI LUNs that were configured, at any time. By increasing the size of each iSCSI LUN this will provide additional room for each of the linked clone based virtual machines to grow over time.

Regardless of the approach taken its best practice to configure your storage arrays notification and alert settings to notify you of any files systems that have exceeded acceptable capacity. By doing so action can be taken before a file system fills up and causing problems.

The following table lists disk volumes per file system for the above storage configuration.

| File system | Disk volumes |
|---|---|
| Master – Full Images | d7 |
| Virtual Machines (Linked Clones) | d11,d14,d15,d16,d17,d18,d19,d20 |

Table 6: Disk Volumes

# Validation Results

This configuration implemented and used in during the validation of the VMware View reference architecture validation for large scale VMware View deployments. Using this configuration the EMC NS20FC was used to validate a 1000 user VMware View building block architecture. For additional information please see the VMware View reference architecture document.

## Storage System Utilization

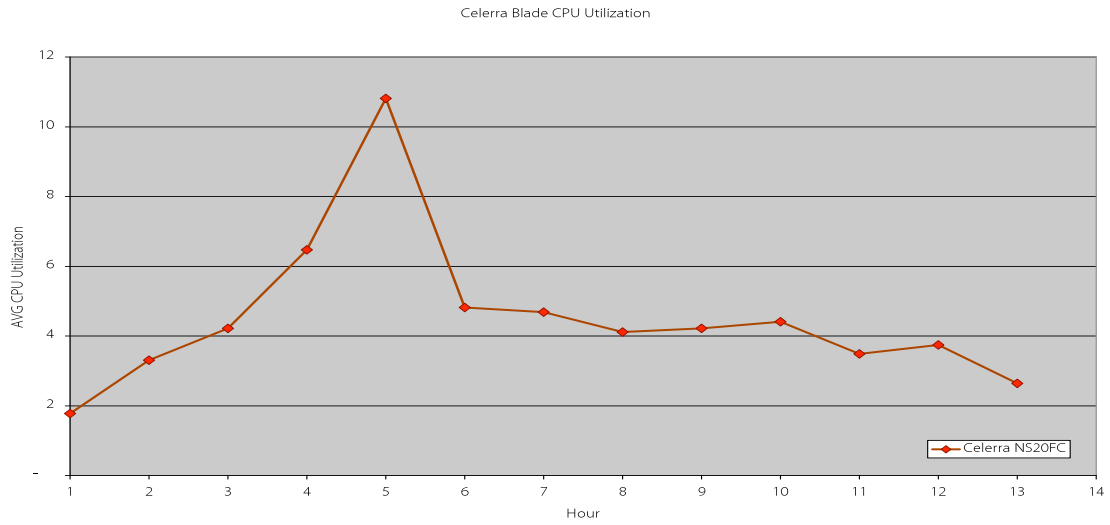


Chart 1: Celerra NS20FC average CPU utilization

The above charts represent the average CPU utilization of the EMC Celerra NS20FC X-Blade 20s. The average CPU utilization of the data movers was 4.5%.

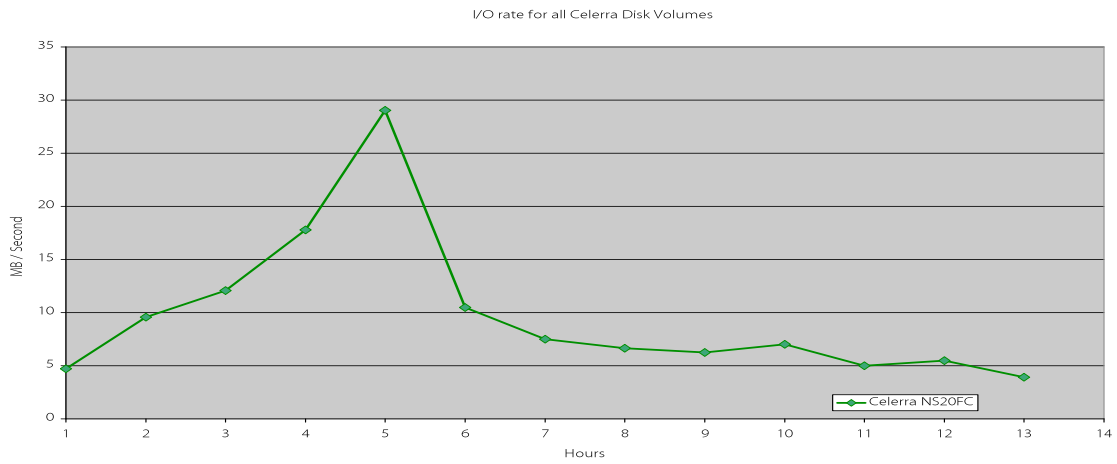


Chart 2: EMC Celerra NS20FC average I/O rate for all disk volumes

Also shown above is the average I/O rate for all the EMC Celerra disk volumes representing plenty of capacity for handling the needed capacity of each building block. The average I/O rate for all the Celerra disk volumes was 9.65 MB/s over the 14-hour time span.

## Application Response Time

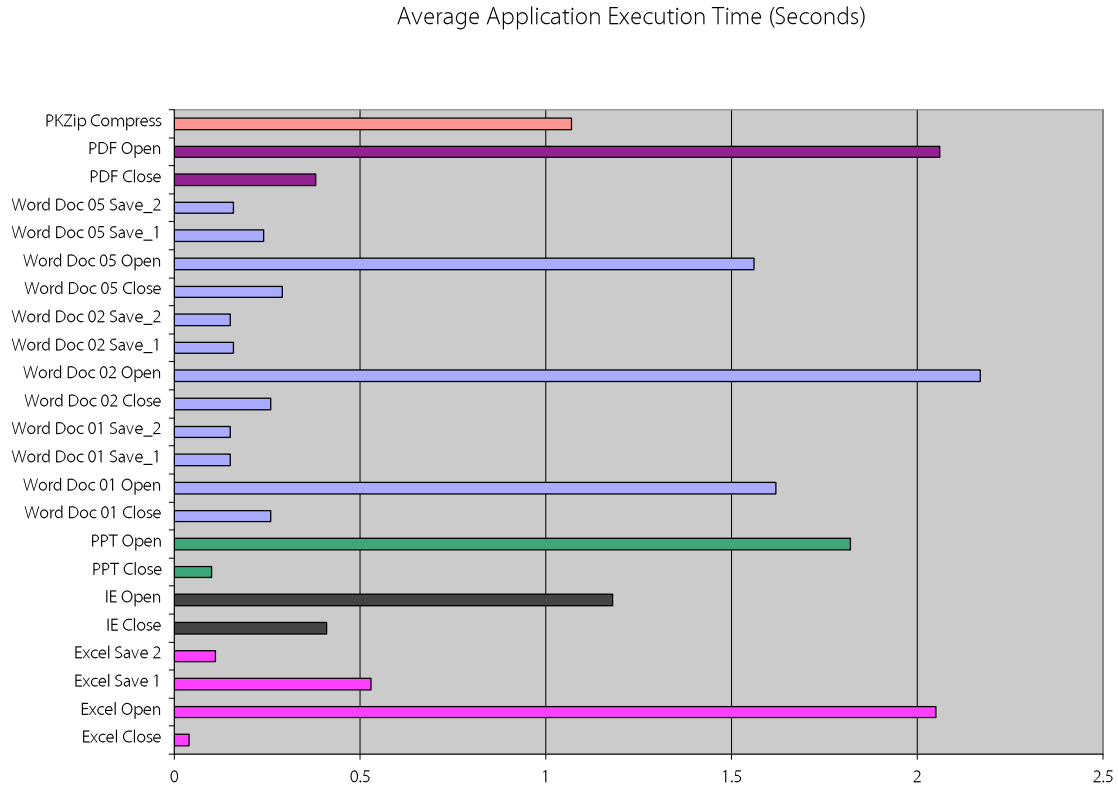


Chart 3: Average application execution time in seconds

The above chart reflects the average application execution time from all virtual desktops in both Cluster A and Cluster B. These application times represent the amount of time it took to open,

close or save a document that was created. This does not represent the amount of time an application is minimized or being worked on. Because of the random nature of the workload, applications being minimized not simply opened, worked on and closed. The randomized workload might start out creating a Microsoft Word document, work on this document, minimize the document and then use Microsoft Internet Explorer. Later the workload will return to the minimized Microsoft Word document work on it some more and close it.

# References


VMware View Administrator Guide

Introduction to View Manager

VMware Infrastructure 3 Documentation

VMware View Windows XP Deployment Guide

VMware View Composer Deployment Guide

VMware View Reference Architecture

Revision: DG-075-INF-01-01

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



vmware®