Demartek 16Gb Fibre Channel Deployment Guide

Overview

The current generation of Fibre Channel (FC), which is 16Gb Fibre Channel (16GFC), was first introduced in the second half of calendar year 2011, and the industry is beginning to see increased interest and adoption of 16Gb Fibre Channel. Therefore, Demartek has produced this Demartek Fibre Channel Deployment Guide, one in a series of technology deployment guides. This guide can be found on our website in our FC Zone or by searching the Internet for "***Demartek Fibre Channel Deployment Guide***" using any well-known Internet search engine.

Audience

This guide is designed for managers and technical professionals within IT departments who are exploring the benefits of deploying, or upgrading to, Fibre Channel technology in general, or 16Gb Fibre Channel specifically, or who are looking for actual deployment examples of 16Gb Fibre Channel solutions.

Objectives of this Guide

This guide is designed to provide basic information about 16Gb Fibre Channel and practical guidance for planning and deploying 16Gb Fibre Channel technology and products. The focus is 16Gb Fibre Channel in primarily, but not exclusively, virtualized environments.

Because 16Gb Fibre Channel includes server adapter and switching technologies, this guide provides information and guidance in each area. A basic understanding of each of these areas is needed to successfully deploy 16Gb Fibre Channel technology.

This guide is intended to be used as a reference and is divided into sections including marketplace data, technology areas, and deployment guidance for virtual servers. There are screen shots and information from actual deployments of some products. At the time of the testing conducted for this guide, there were few true 16Gb Fibre Channel storage targets. Multiple 8Gb paths to storage targets were used to provide suitable bandwidth to demonstrate the capability of 16Gb Fibre Channel.

All of the work was performed in the Demartek lab in Colorado, USA.

About Demartek

Demartek is an industry analyst organization with its own ISO 17025 accredited EPA ENERGY STAR certified testing laboratory for Data Center Storage. The vast majority of the research work we do involves hardware and software solutions tested in our lab by our staff. The Demartek lab is equipped with servers, network infrastructure, and storage, and supports multiple speeds of Ethernet, iSCSI and Fibre Channel, along with SSDs and a variety of other technologies. Many public lab validation and test reports highlighting Fibre Channel and other technologies are available on our website.

Demartek Videos

Demartek produces highlight videos of public evaluations and deployment guides. Links to these videos are available on our web site in the Demartek Video Library and are posted on the Demartek channel on YouTube.

Demartek Lab Notes

To be notified when new Deployment Guides and lab validation reports become available, you can subscribe to our free monthly newsletter, Demartek Lab Notes, available on our website. We do not give out, rent or sell our email list.

Table of Contents

Basic Storage Architectures

Direct Attached Storage (DAS)

Direct Attached Storage (DAS) is probably the most well-known form of computer storage. In a DAS implementation, the host computer has a private connection to the storage and almost always has exclusive ownership of the storage. The host computer accesses the storage in a "block" fashion, which means that it directly addresses blocks on the storage device. This implementation is relatively simple and usually inexpensive. Potential disadvantages are that the distance between the host computer and the storage are frequently short, such as inside a computer chassis or within a rack or adjacent rack. Some DAS implementations require that the host computer be taken offline when adding or removing storage devices, such as a boot drive directly connected to a motherboard storage interface. SATA is a common DAS interface.

Network Attached Storage (NAS)

Network Attached Storage (NAS) devices, also known as file servers, share their storage resources with clients on the network in the form of "file shares" or "mount points." The clients use network file access protocols such as SMB (formerly known as CIFS) or NFS to request files from the file server. The file server then uses block protocols to access its internal storage to satisfy the requests. Because NAS operates on a network, the storage can be very far away from the clients. Many NAS solutions provide advanced features such as snapshot technologies, global namespace, SSD caching, and more.

Storage Area Network (SAN)

SAN architecture provides a way to use block access methods over a network such as Ethernet or Fibre Channel to provide storage for host computers. The storage in a SAN is not owned by one server but is accessible by all of the servers on the network. This SAN storage can be carved into logical storage pools or volumes that can be assigned to particular host servers. These logical volumes are independent of the geometries or components of the storage hardware. The storage appears to host servers and applications in the same way that DAS storage appears, but because SAN storage uses a network, storage can be a long distance away from the host servers.

SAN architectures currently use block Small Computer System Interface (SCSI) protocol for sending and receiving storage data over their respective networks. Fibre Channel (FC) SANs implement the SCSI protocol within the FC frames. Internet SCSI (iSCSI) SANs implement the same SCSI protocol within TCP/IP packets. Fibre Channel over Ethernet (FCoE) encapsulates the Fibre Channel protocol within Ethernet packets using a relatively new technology called Data Center Bridging (DCB), which is a set of enhancements to traditional Ethernet and is currently implemented with some 10GbE infrastructure. Because each of these technologies allow applications to access storage using the same SCSI command protocol, it is possible to use all of these technologies in the same enterprise or to move from one to the other. Generally speaking, applications running on a host server cannot tell the difference between Fibre Channel SAN storage, FCoE SAN storage, and iSCSI SAN storage. In fact, applications generally cannot tell the difference between DAS storage and SAN storage.

There has been much debate over Fibre Channel vs. iSCSI. Some people focus on the lower entry price points available for iSCSI SANs, while others focus on high reliability, robustness and availability of Fibre Channel SANs. The conventional wisdom in this debate no longer holds true in many cases. While it is true that at the low end iSCSI tends to be less expensive than Fibre Channel, as performance, reliability, and high availability features and designs are included in iSCSI solutions, the iSCSI price advantage diminishes. Over the past few months we tested various solutions for ease of use and found that some implementations of Fibre Channel are just as easy (or easier) to use than some iSCSI implementations. We have also found several iSCSI solutions that are designed for performance, high reliability, and high availability. Much work has been done to increase performance, reliability, and ease-of-use with both of these SAN technologies. The real issues are satisfying the needs of the business or application, working with existing infrastructure and expertise, maintaining service-level agreements, and staying within budgets. Fibre Channel and iSCSI technology can meet these requirements, and there is room for both in current and future IT environments.

Most of the larger enterprises that have implemented SAN technology have implemented Fibre Channel technology. These enterprises typically demand proven technology, have the need for high bandwidth storage solutions, have the budgets to pay for more expensive hardware to meet their performance and reliability needs, and typically have full-time staff dedicated to storage management. Some of these enterprises continue to invest in Fibre Channel storage solutions and plan to do so for the foreseeable future. However, some of these enterprises are also investing in iSCSI storage solutions, especially with 10GbE technology, for their virtualized server environments.

Smaller enterprises are often attracted to iSCSI technology because of its lower entry price point, allowing them to grow their iSCSI SAN as their needs change. The lower entry price typically uses 1GbE technology, but when 10GbE technology is used, the price for the equipment increases and it is no longer considered "entry level."

There is more to choosing a storage system than selecting the host interface. Regardless of the type of interface, several other factors need to be considered, including the number and type of disk drives, amount and type of SSD technology, management software, advanced features, support from the vendor, service level agreements (SLAs) and several other factors. Advanced features of modern storage systems may include various forms of replication, thin provisioning, compression, data de-duplication, caching, automated storage tiering, and more.

Unified Storage

Unified storage combines NAS and SAN technologies into a single, integrated solution. These unified storage solutions provide both block and file access to the shared storage environment. These often provide simplified management by combining the management of all storage, regardless of the transport or "plumbing," into a single management console.

Demartek Storage Interface Comparison Reference Page

We have compiled a vendor-neutral reference page on our website that provides technical information about Fibre Channel and other storage interfaces used for storage applications. We update this page periodically, and this page includes a downloadable, interactive PDF version with the same contents as the web page. Detailed information regarding history, roadmaps, transfer rates, encoding schemes, cabling, connectors and more is available on this page.

This page can be found by entering "***Storage Interface Comparison***" in any major Internet search engine or by following the link below:
www.demartek.com/Demartek_Interface_Comparison.html on the Demartek website.

Fibre Channel Marketplace

Deployed in 90 percent of Fortune 1000 datacenters, Fibre Channel is the de facto standard for storage networking in the datacenter. There are literally millions of Fibre Channel ports installed in datacenters worldwide.

Fibre Channel History

Past and Current Versions

Fibre Channel is an industry standard storage solution. Development began in 1988, receiving ANSI standard approval in 1994. Fibre Channel technology was introduced to the market in 1997 at 1 gigabit per second (1Gb) and has doubled in speed every three or four years through the current 16Gb Fibre Channel technology. Although each generation of Fibre Channel was introduced to the market in the years indicated in the chart below, widespread industry adoption sometimes began later.

Generally speaking, Fibre Channel is backward compatible with the two previous generations. This means that 16GFC switches, host bus adapters (HBAs) and optics (transceivers) are backward compatible with 8GFC and 4GFC equipment.

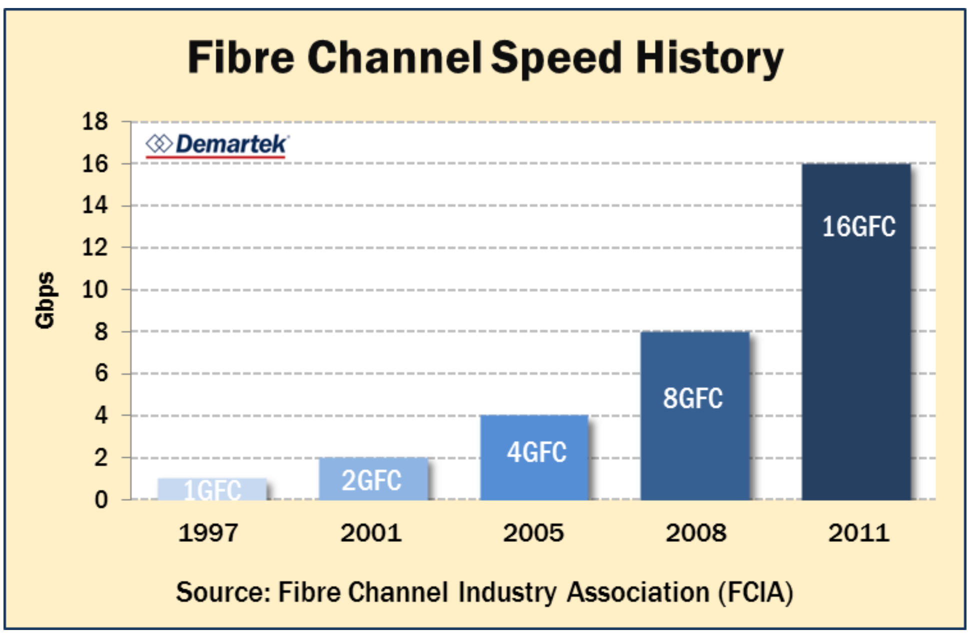

Fibre Channel over Ethernet (FCoE) was introduced into the market in 2009, and works by tunneling Fibre Channel through Ethernet. FCoE currently runs at 10Gbps. Other than the speed and Ethernet carrier, FCoE behaves the same as native FC.

Generation 6 Fibre Channel – the Next Version

In February 2014, the Fibre Channel Industry Association (FCIA) officially announced "Gen 6" Fibre Channel. This generational naming is an effort to get away from focusing only on the speed of the interface, but to acknowledge that other important features are included with each new generation of Fibre Channel.

This generational naming philosophy can be found in other technologies such as mobile telephones, where the industry has moved from "1G" to "4G" today. As the mobile telephone industry has moved through its generations, the speeds have improved, but other things such as talk time, the physical size of the devices and other features have also improved. In fact, few consumers know the actual transmission speeds for these generations of mobile telephone technology.

Gen 6 Fibre Channel includes a single-lane specification running at 32Gbps and a parallel, four-lane (4 x 32Gbps) specification known as 128GFC. Products that use the new Gen 6 Fibre Channel are expected to become widely available in calendar year 2016. It is our belief that the 32GFC products will probably become available before the 128GFC products.

Gen 6 Fibre Channel includes energy efficient features that allow the optical connectors to operate in a stand-by mode (or "nap") multiple times each second.

According to the specification, Gen 6 Fibre Channel (both 32GFC and 128GFC) is backward compatible with 16GFC and 8GFC.

16Gb Fibre Channel Specifics

16Gb Fibre Channel provides not only doubles the throughput from the previous generation, but has other benefits.

Some of these benefits are directly related to the speed increase, such as a reduced number of links needed to achieve the same bandwidth, reduced power consumption needed to achieve the same bandwidth and fewer cables to manage. In addition, the higher speed allows fabrics to be connected with fewer inter-switch links (ISLs), which is especially helpful in large fabrics.

16Gb Fibre Channel includes retimers in the optical modules and transmitter training, features that improve link performance characteristics, electronic dispersion compensation and backplane links.

Table 1 – Fibre Channel Speed Characteristics

Speed	Throughput (Mbps)	Line Rate (Gbps)	Encoding	Retimers in the module	Transmitter training
1 GFC	100	1.0625	8b/10b	No	No
2 GFC	200	2.125	8b/10b	No	No
4 GFC	400	4.25	8b/10b	No	No
8 GFC	800	8.5	8b/10b	No	No
16 GFC	1600	14.025	64b/66b	Yes	Yes

Table 2 – Fiber Optic Cable Link Distance

Speed	Multi-Mode				Single-Mode
	OM1	OM2	OM3	OM4	OS1
1 GFC	300	500	860	*	10,000
2 GFC	150	300	500	*	10,000
4 GFC	50	150	380	400	10,000
8 GFC	21	50	150	190	10,000
16 GFC	15	35	100	125	10,000

* The link distance for OM4 fiber optic cable has not been defined for these speeds.

Fibre Channel Technology and Definitions

Direct Connect

A Fibre Channel storage target can be connected in a point-to-point or direct fashion to a server (initiator) without needing to go through switches or an arbitrated loop.

Arbitrated Loop

Fibre Channel arbitrated loop (FC-AL) is a Fibre Channel design implementation where devices are connected in a one-way loop and bandwidth is shared by all devices. When a device needs to put data on the channel, it requests control of the loop by sending an arbitration signal to the elements of the loop. This signal is used to determine which device ports are granted ownership of the bandwidth and the associated ports are opened.

As FC switches became more affordable in the late 2000s, use of FC-AL between storage and servers declined. However, some FC target devices may use the FC-AL designs internally. Dual loops are often used inside of these storage devices to provide redundancy.

Switched Fabric

A network of Fibre Channel devices and switches is referred to as a fabric. Switched fabric is the general design implementation that most people today associate with Fibre Channel deployments. With this general design, SAN devices are found on the edge of the network, various configurations of switches are found in the core of the network, and fiber optic cabling is used to connect these components together.

The term topology describes how the switches are interconnected, such as core-edge, edge-core-edge or fully meshed. The goal of these topologies is to provide good performance for all of the edge devices, good resiliency of the network, while minimizing the number of interconnections between switches.

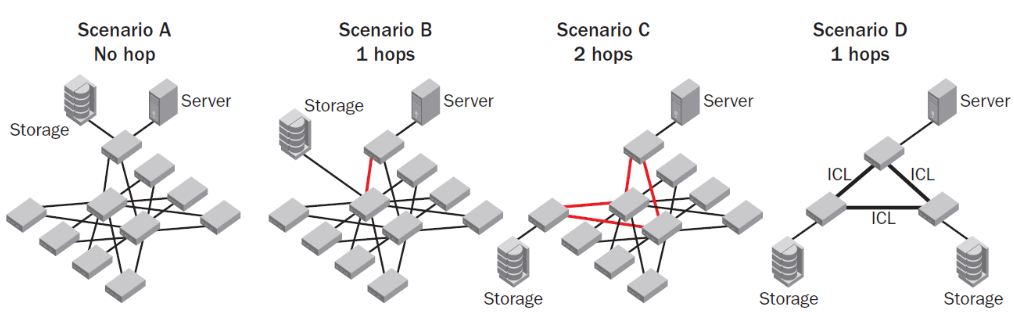

The placement of the devices and the associated traffic flows are the key differences between the four scenarios.

- **Scenario A** has localized traffic, which has no extra hops between switches, providing a slight performance advantage but does not provide scalability or resilience features.
- **Scenario B** is an edge-core switch configuration, separating the server and storage and introducing one switch hop. This topology provides increased scalability opportunities.
- **Scenario C** is an edge-core-edge configuration, with two hops between switches, and provides large scalability opportunities.
- **Scenario D** is a full-mesh topology and uses inter-switch connections so that there is no more than one hop between the edge devices such as servers and storage. These configurations can grow to thousands of device ports.

Inter-switch connections are sometimes called inter-switch links (ISLs) or inter-chassis links (ICLs).

Initiators

Initiators are typically host servers that need access to storage resources. This is in keeping with the standard SCSI protocol where initiators issue commands to storage devices and wait for replies. Fibre Channel initiators are usually implemented in a hardware host bus adapter, but are sometimes implemented in software. Fibre Channel initiators can access multiple Fibre Channel targets simultaneously.

Targets

Targets are devices that provide storage resources for Fibre Channel initiators. These follow the SCSI protocol for targets by responding to requests from initiators. Targets can be implemented as dedicated hardware devices such as disk arrays or as combination hardware and software in a server or appliance solution. Fibre Channel targets can respond to multiple Fibre Channel initiators simultaneously.

Fibre Channel Host Bus Adapters

Fibre Channel Host Bus Adapters (HBAs) connect a device such as a server or storage device to the Fibre Channel network or another FC device. Due to the architecture of Fibre Channel HBAs, most of the processing of the I/O requests is performed by the HBA itself, and is “offloaded” from the host CPU. As a result, Fibre Channel I/O traffic tends to be very host CPU efficient when compared to other interfaces that do not offload the I/O processing.

Fibre Channel Storage Arrays

Many storage vendors sell storage array products configured for the Fibre Channel protocol. FC storage arrays can be composed of hard disk drives (HDDs), solid state drives (SSDs) or a combination of the two. These arrays have target mode Fibre Channel HBAs present. An FC array can be connected to a host via direct connect, arbitrated loop, or switched fabric.

We are beginning to see products from storage vendors that have 16Gb Fibre Channel host ports. The performance of 16Gb Fibre Channel can be approximated by utilizing multiple 8Gb ports (or slower ports). However, using 16Gb Fibre Channel components requires fewer cables, optics and switch ports than it would to achieve equivalent performance with 8Gb Fibre Channel or 4Gb Fibre Channel.

Fibre Channel Switches and Backbones

The switches are designed to run stand-alone or as edge switches for part of a larger SAN fabric. The Backbones, also known as Directors, form the central core of an enterprise Fibre Channel fabric and provide reliable, scalable, high-performance switching infrastructure for mission-critical storage. There are also embedded Fibre Channel switches designed for the blade server market and Fibre Channel encryption solutions for protecting data-at-rest. Fibre Channel extension solutions provide ESCON, FICON and FCIP (Fibre Channel over IP). FCIP is often used to transmit Fibre Channel protocol over a WAN.

World Wide Node Name and World Wide Port Name

IEEE provides FC device vendors with unique identities called World Wide Names (WWN) in the form of a 64-bit address, not unlike a MAC address for Ethernet devices. The World Wide Node Name (WWPN) identifies the endpoint, such as an HBA, while the World Wide Port Name identifies individual ports on a device (for example, the single or dual ports of an HBA).

Switch Management Software

Software management tools are available that support full lifecycle management of all switches in the network to ensure network availability and optimal performance and available for use on most versions of Windows, Linux, VMware and Microsoft Hyper-V.

Fibre Channel Switch Zoning

Zoning is a service in Fibre Channel SANs that groups together servers and storage that need to communicate with each other. Elements of a zone can only communicate with other elements within the same zone. This prevents unauthorized access of storage. Servers and storage can be members of multiple zones.

Web interface management tools are available to zone switches, such as the example screenshots shown below.

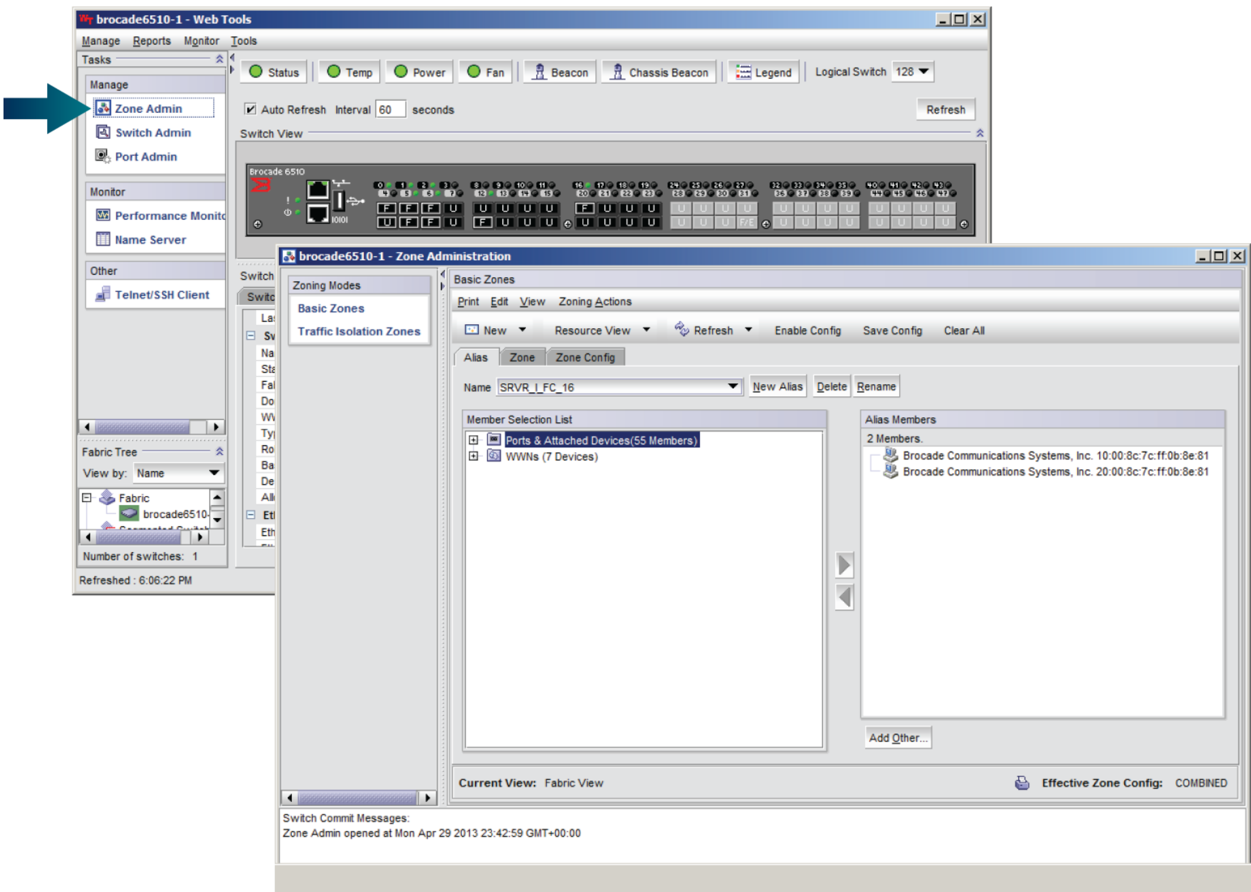

Aliasing allows the administrator to assign logical names to switch ports or WWNs for simpler management. Though not necessary for Fibre Channel zoning, aliasing can improve the ease of zoning administration.

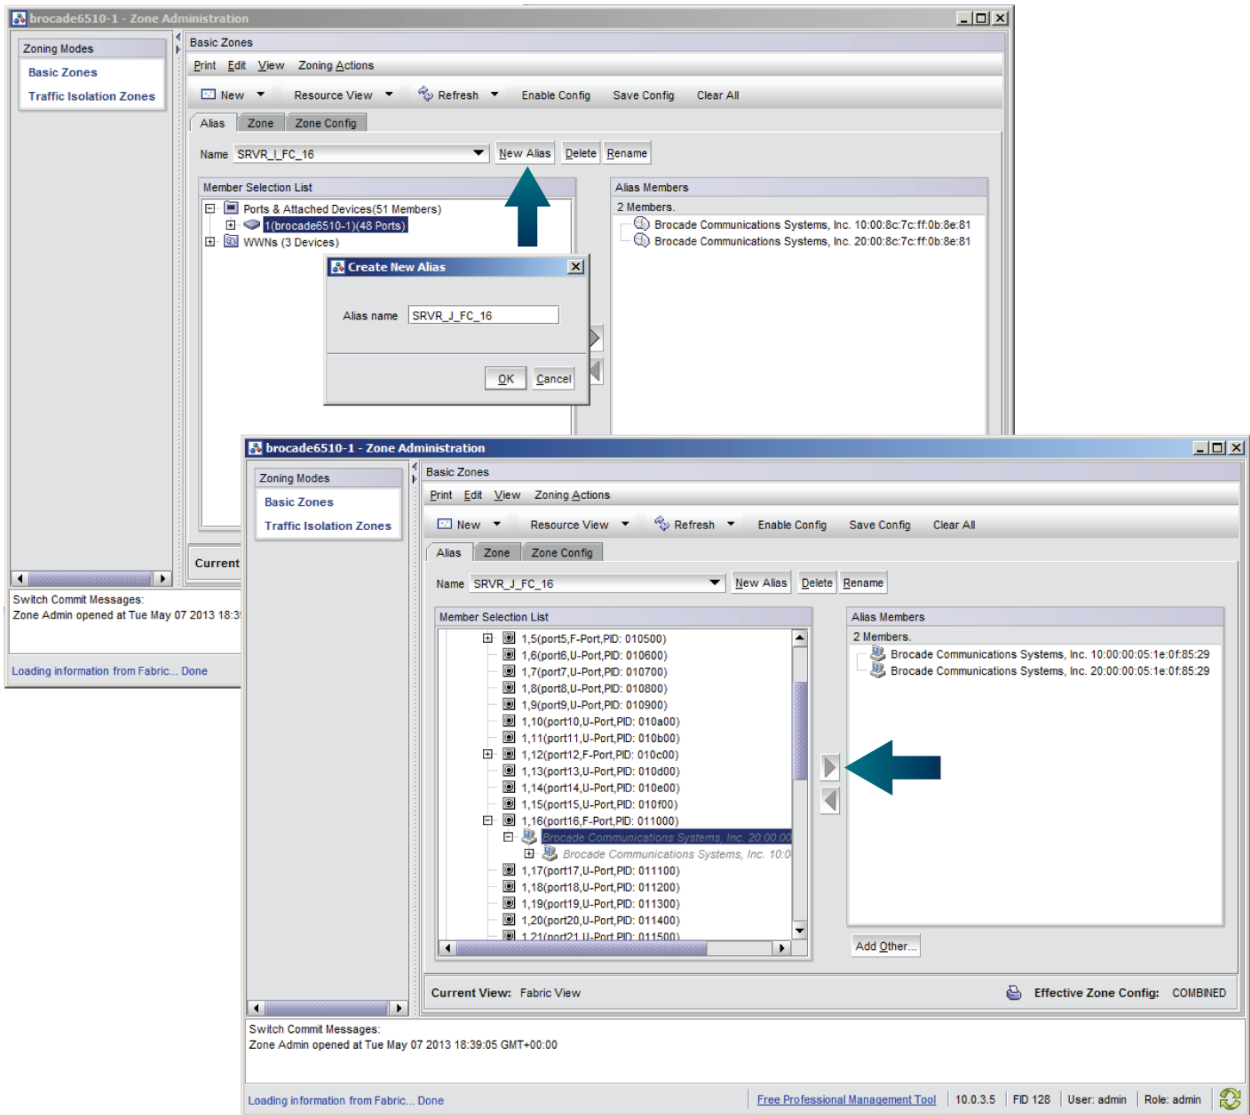

Selecting WWNs as alias members tends be more portable than using switch ports because the WWNN and WWPN can move from one switch port to another if the fiber optic cables are moved. However, if the Fibre Channel HBAs are moved from one server or storage target to another, it's possible that an alias or zone based on WWNs may have unexpected members, which can be a security concern.

The same ports, WWNs, and aliases can be members of more than one zone. A zone must include both storage clients (initiators) and storage targets to be effective.

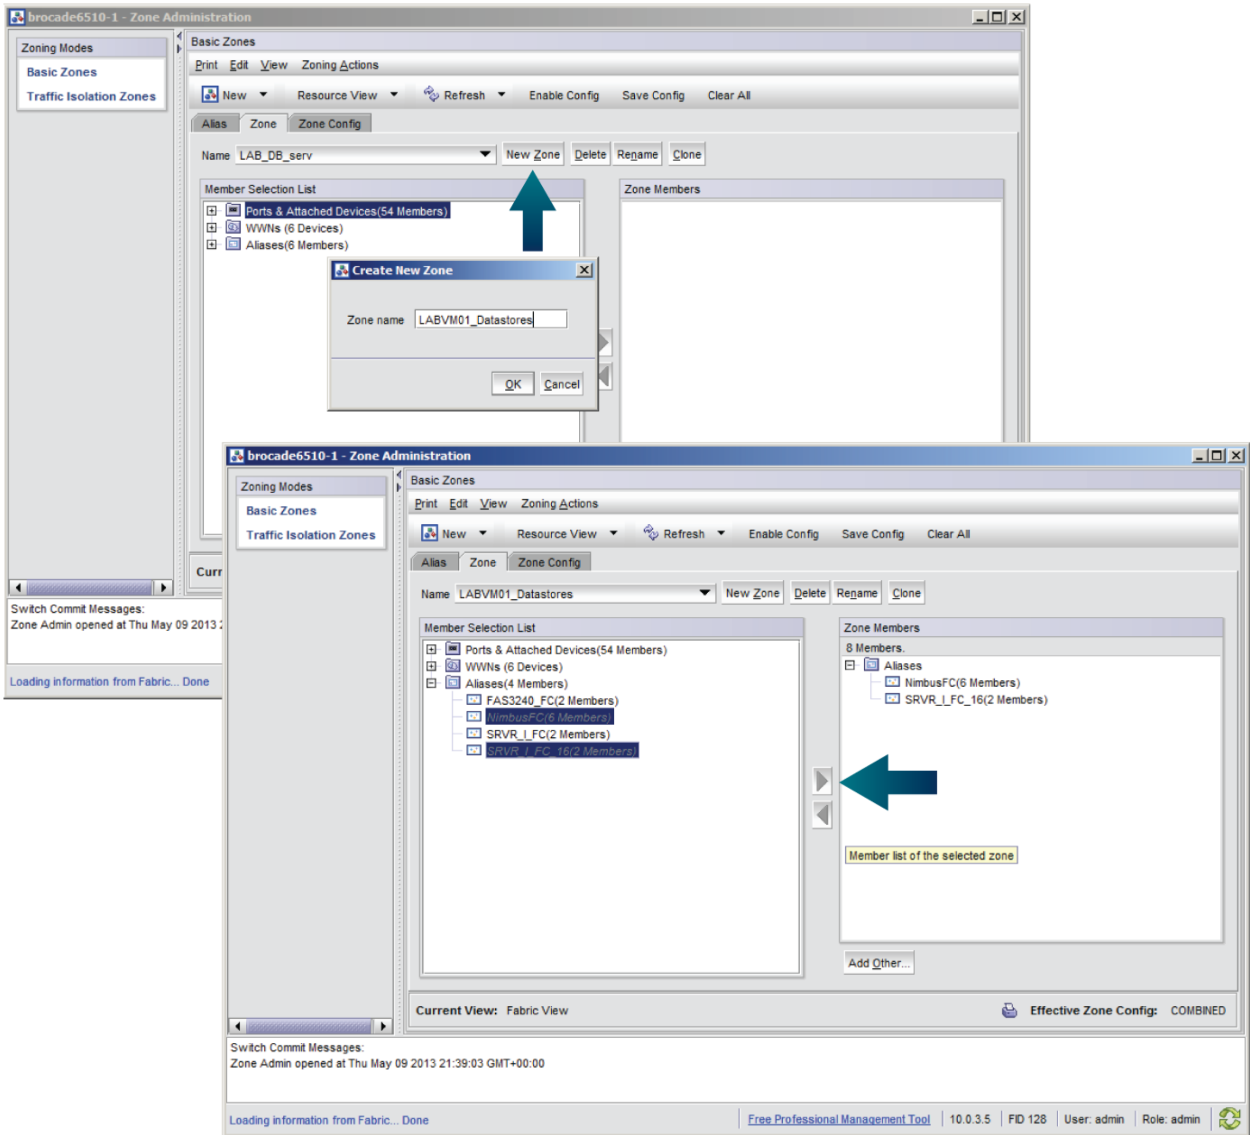

After a zone is created, it must be enabled before storage traffic will run across it. A zone configuration is a collection of zones that will be active at one time. After a zone configuration is created, a name for this zone configuration can be assigned.

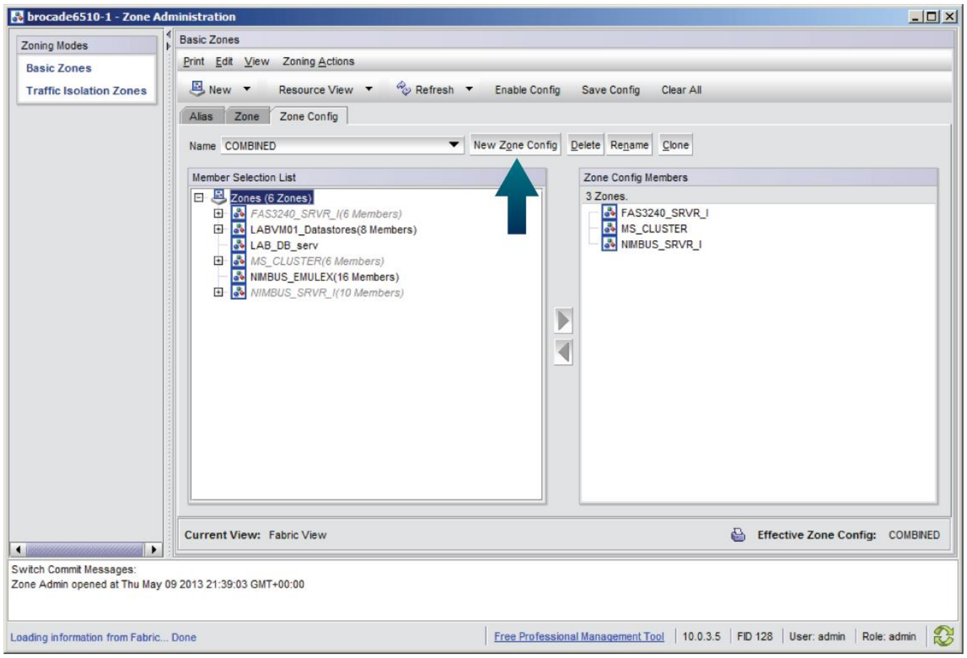

Zones are then added to the zone configuration, and then the entire configuration should be saved to preserve aliases, zones, and zone configurations if the switch is powered off or rebooted.

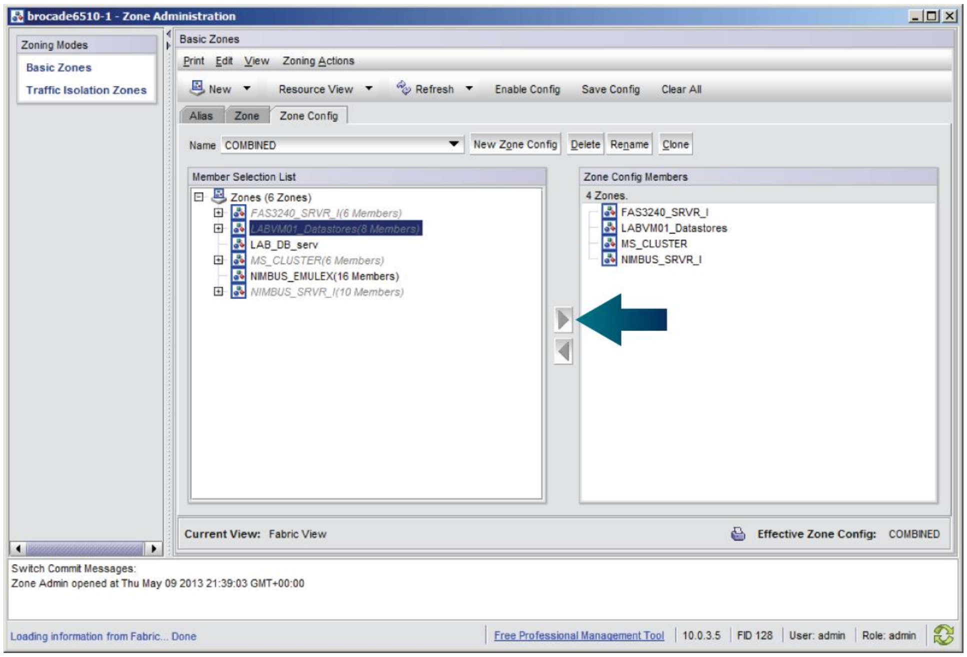

Deploying 16Gb Fibre Channel in Virtualized Environments

Fibre Channel is supported by all the common virtualization platforms such as VMware vSphere, Microsoft Hyper-V, and others. Fibre Channel LUNs are suitable to host both guest operating systems and as application space within guests. Provisioning storage can be accomplished with either virtual disks managed by the hypervisor or by passing the raw LUNs directly through to the guest OS.

A feature of 16Gb Fibre Channel that is particularly attractive to virtualization managers is the speed at which datastores or virtual machines can be migrated from one device to another.

Microsoft Hyper-V

Hyper-V requires no special configuration to connect to FC LUNs once the Windows Server operating system has the device drivers loaded. Microsoft Hyper-V supports two methods of configuring Fibre Channel LUNs for use by guest virtual machines—Virtual Hard Disk method and Pass-through method.

Virtual Hard Disk

In the VHD method, a Fibre Channel volume is allocated to the parent (Hyper-V) partition in the usual manner. This Fibre Channel volume is brought online, initialized, and formatted, and folders may be created if desired.

In the Hyper-V management role, we created a new hard disk using the "Actions" pane. There are three choices for the type of VHD: Fixed size, Dynamically expanding and Differencing.

The new VHD onto the FC volume that was just attached, and storage space allocated.

In the Hyper-V management role, select a virtual machine that is not running, and open "Settings." Add a hard drive under the IDE Controller, and specify the location of the VHD just created on the new FC volume.

After starting that virtual machine, the new volume appears in Disk Manager for that virtual machine.

Pass-through (Raw Device Mapping)

The pass-through method for allocating Fibre Channel LUNs to guest virtual machines requires that the LUN be allocated on the parent partition in the usual manner, similar to the VHD method. The difference is that the FC LUN is assigned to the guest by using the "settings" function for the guest VM in the Hyper-V management role, while the guest is not running.

The LUN is given to the guest by the "Add Hardware" function and selecting the "SCSI Controller". In the SCSI Controller section, select "Hard Drive", then "Add". Choose "Physical hard disk." Create a hard disk of appropriate size.

After starting that virtual machine, the new volume appears in Disk Manager for that virtual machine.

NPIV

NPIV, essentially virtualizes the FC HBA. With Microsoft Hyper-V, a virtual machine can be configured to use a Virtual Storage Area Network with N_Port ID Virtualization (NPIV) as long as it has a compatible FC HBA

Once a VM has booted in the hypervisor machine, the VM can use the Virtual SAN to connect to a Fibre Channel LUN. From Hyper-V, the Virtual SAN Manager is selected and a new Virtual Fibre Channel SAN created

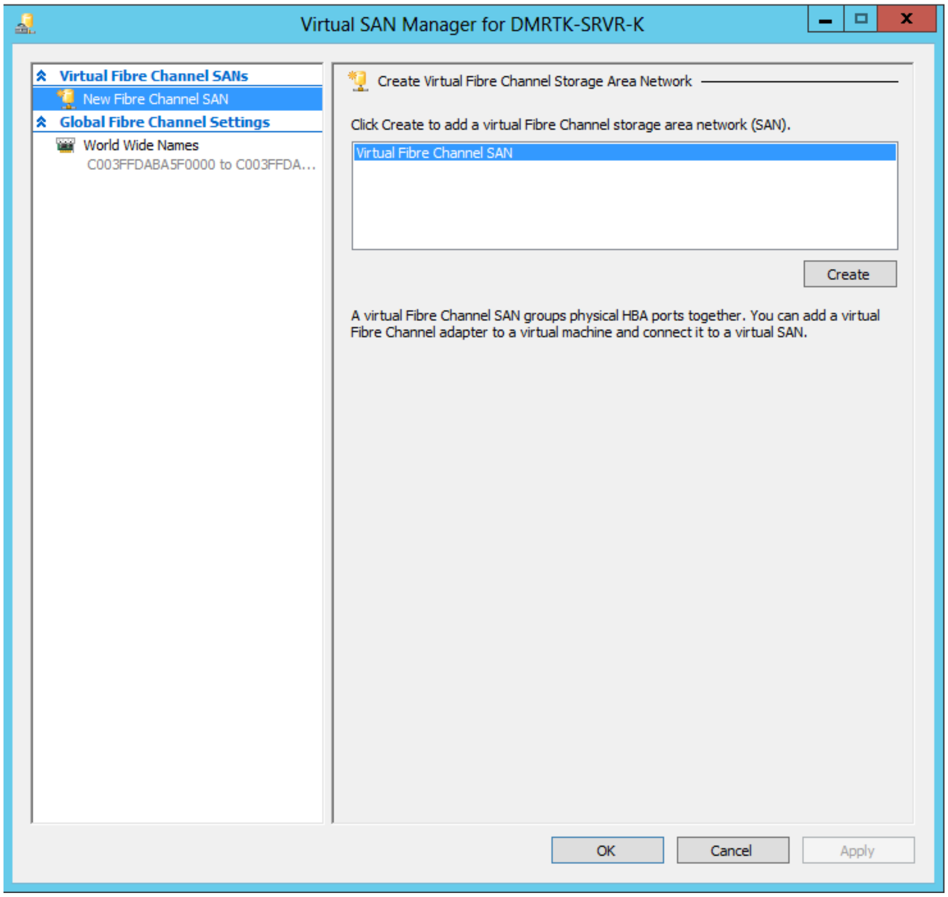
Virtual SAN Manager for DMRTK-SRVR-K
Virtual Fibre Channel SANs
New Fibre Channel SAN
Global Fibre Channel Settings
World Wide Names
C003FFDABA5F0000 to C003FFDA...
Create Virtual Fibre Channel Storage Area Network
Click Create to add a virtual Fibre Channel storage area network (SAN).
Virtual Fibre Channel SAN
Create
A virtual Fibre Channel SAN groups physical HBA ports together. You can add a virtual Fibre Channel adapter to a virtual machine and connect it to a virtual SAN.
OK
Cancel
Apply

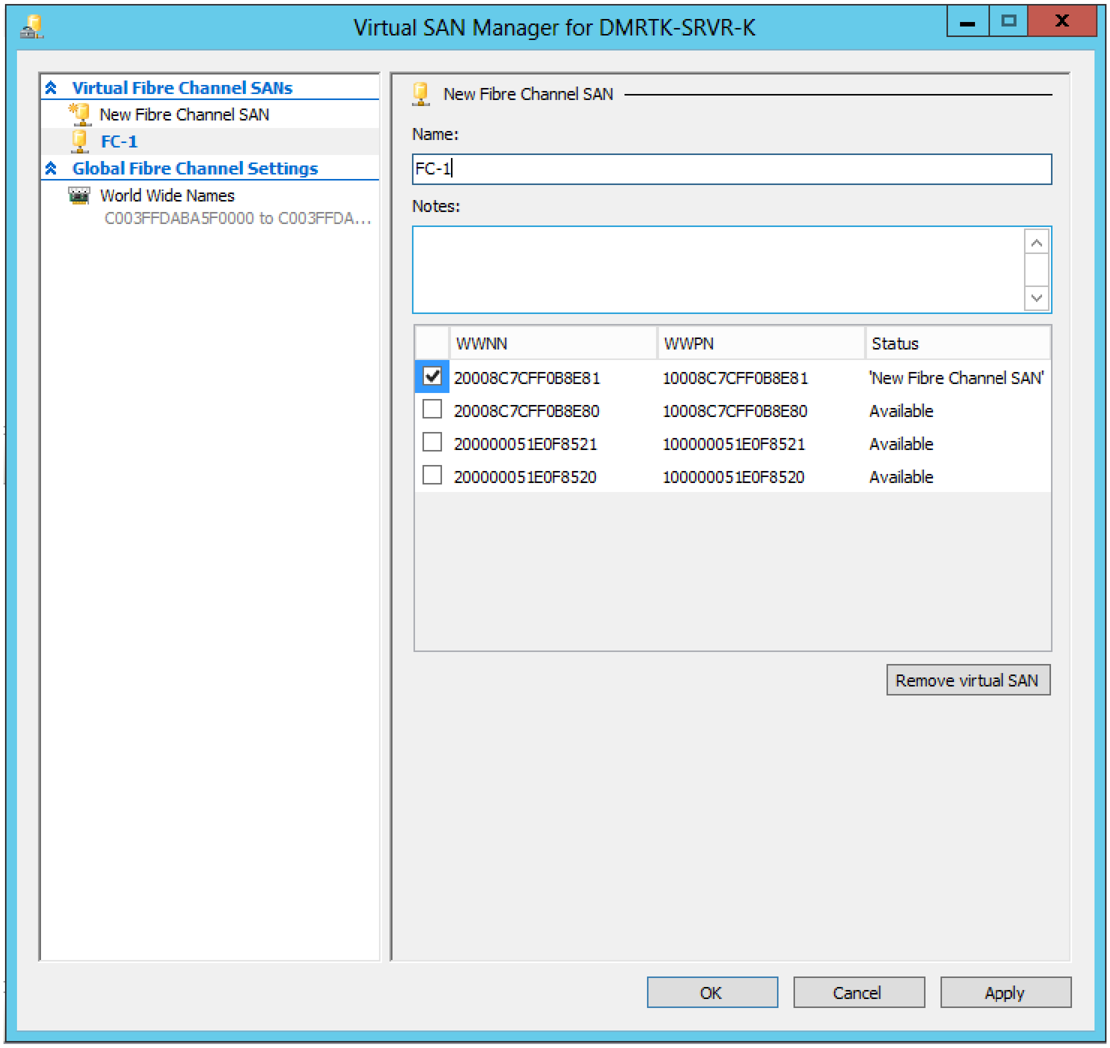

After the Virtual SAN is created, the settings for the VM are edited. Select Add Hardware. "Fibre Channel Adapter" should be listed as one of the hardware devices that can be added. A list of Virtual SANs will be displayed.

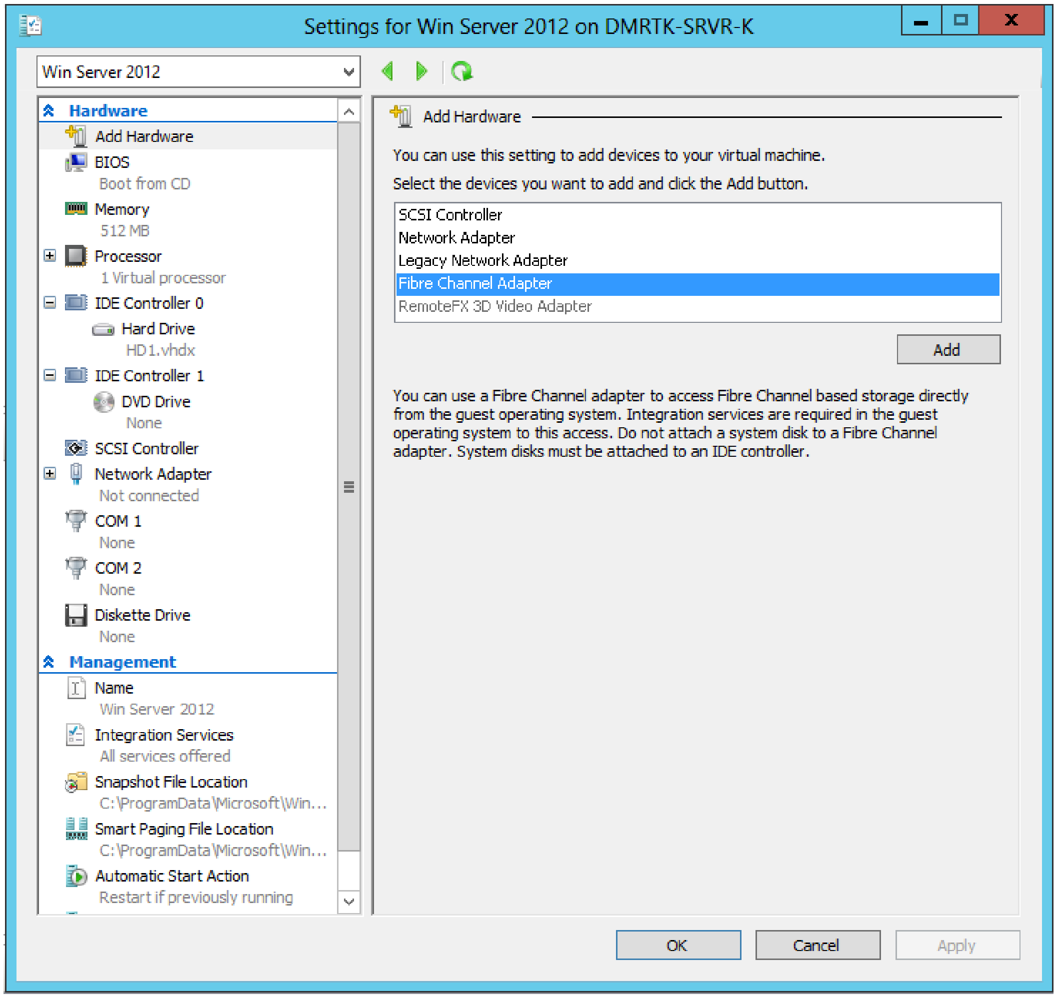
Settings for Win Server 2012 on DMRTK-SRVR-K
Win Server 2012
Hardware
Add Hardware
BIOS
Boot from CD
Memory
512 MB
Processor
1 Virtual processor
IDE Controller 0
Hard Drive
HD1.vhdx
IDE Controller 1
DVD Drive
None
SCSI Controller
Network Adapter
Not connected
COM 1
None
COM 2
None
Diskette Drive
None
Management
Name
Win Server 2012
Integration Services
All services offered
Snapshot File Location
C:\ProgramData\Microsoft\Win...
Smart Paging File Location
C:\ProgramData\Microsoft\Win...
Automatic Start Action
Restart if previously running
Add Hardware
You can use this setting to add devices to your virtual machine.
Select the devices you want to add and click the Add button.
SCSI Controller
Network Adapter
Legacy Network Adapter
Fibre Channel Adapter
RemoteFX 3D Video Adapter
Add
You can use a Fibre Channel adapter to access Fibre Channel based storage directly from the guest operating system. Integration services are required in the guest operating system to this access. Do not attach a system disk to a Fibre Channel adapter. System disks must be attached to an IDE controller.
OK
Cancel
Apply

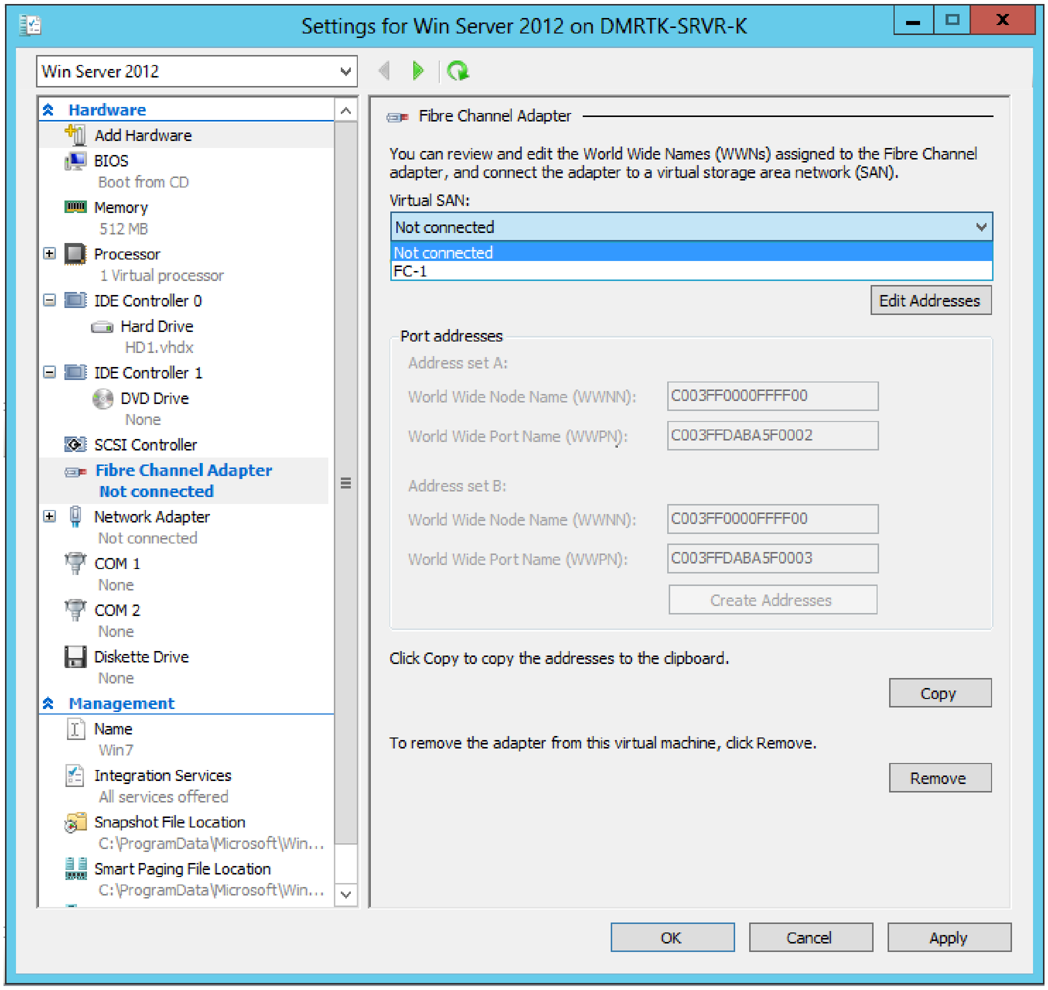
Settings for Win Server 2012 on DMRTK-SRVR-K
Win Server 2012
Hardware
Add Hardware
BIOS
Boot from CD
Memory
512 MB
Processor
1 Virtual processor
IDE Controller 0
Hard Drive
HD1.vhdx
IDE Controller 1
DVD Drive
None
SCSI Controller
Fibre Channel Adapter
Not connected
Network Adapter
Not connected
COM 1
None
COM 2
None
Diskette Drive
None
Management
Name
Win7
Integration Services
All services offered
Snapshot File Location
C:\ProgramData\Microsoft\Win...
Smart Paging File Location
C:\ProgramData\Microsoft\Win...
Fibre Channel Adapter
You can review and edit the World Wide Names (WWNs) assigned to the Fibre Channel adapter, and connect the adapter to a virtual storage area network (SAN).
Virtual SAN:
Not connected
Not connected
FC-1
Edit Addresses
Port addresses
Address set A:
World Wide Node Name (WWNN):
C003FF0000FFFF00
World Wide Port Name (WWPN):
C003FFDABA5F0002
Address set B:
World Wide Node Name (WWNN):
C003FF0000FFFF00
World Wide Port Name (WWPN):
C003FFDABA5F0003
Create Addresses
Click Copy to copy the addresses to the clipboard.
Copy
To remove the adapter from this virtual machine, click Remove.
Remove
OK
Cancel
Apply

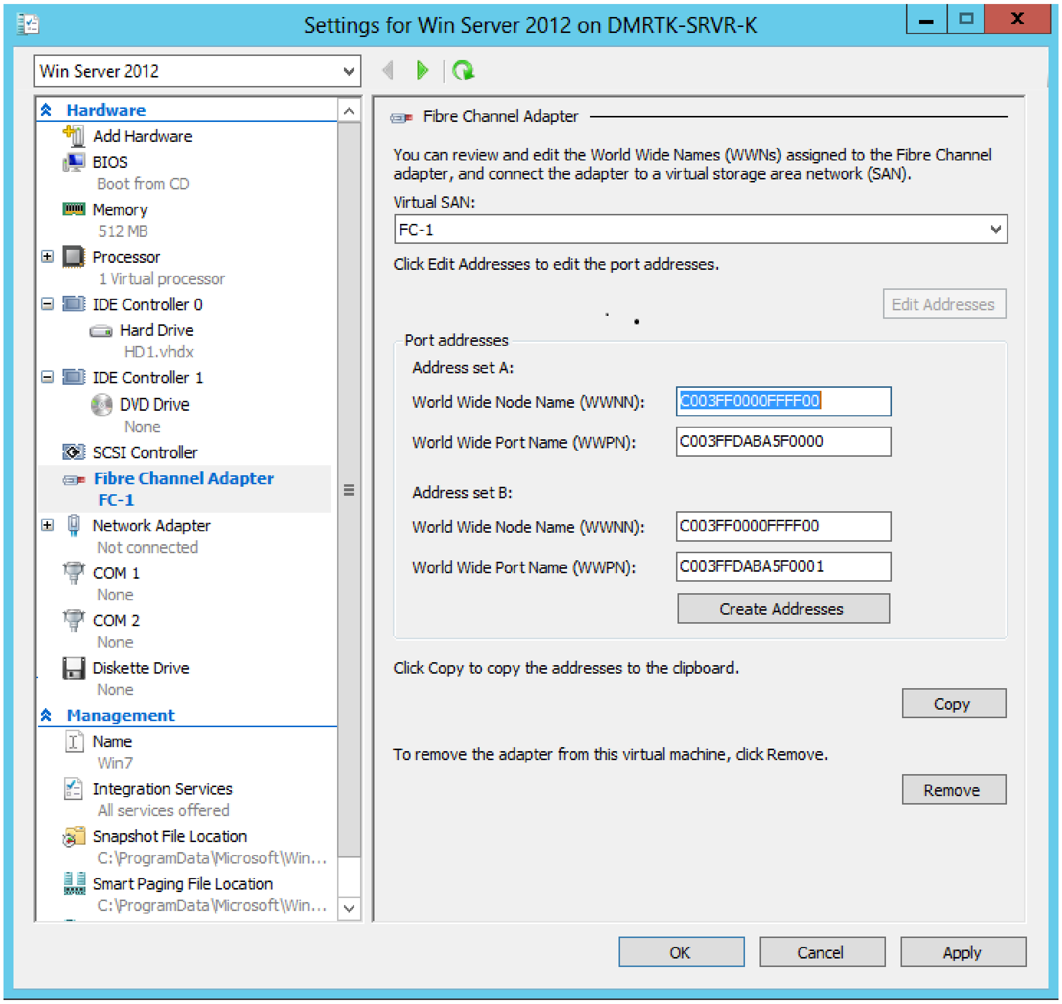

Windows can generate WWNs, or the user can enter their own based on their own needs. When the VM boots, the VM will obtain its LUN using the Virtual SAN in the same way a bare metal installation would use an HBA to access a physical SAN.

If another hypervisor host with a 16Gb FC HBA is available, the VMs can be migrated to it without detaching their storage. Migration can be accomplished while VMs are running, with Virtual Fibre Channel LUNs connected and being subjected to I/O

If two hypervisor hosts are clustered, a Fibre Channel LUN can be configured as a Cluster Shared Volume (CSV). This volume will be seen by both machines. Provided the VHD that stores the VM is on the CSV instead of local storage, the time required for the VM to migrate is drastically reduced as the corresponding VHD does not have to be transferred over the network to the new machine.

VMware

VMware vSphere drivers need to be installed on the vSphere host

```
/tmp # esxcli software vib install --depot=/tmp/BCD-bfa-3.2.0.0-00000-offline_bundle-
943825.zip

Installation Result
Message: The update completed successfully, but the system needs to be rebooted for
the changes to be effective.
Reboot Required: true
VIBs Installed: Brocade_bootbank_scsi-bfa_3.2.0.0-1OEM.500.0.0.472560
VIBs Removed:
VIBs Skipped:
```

A system reboot is required to complete the installation. The ESX/ESXi host is then ready to accept 16Gb Fibre Channel traffic over the HBA. After LUNs are created on the storage and the switch(es) are zoned appropriately, open the Configuration tab and navigate to the Storage adapters page. Rescan the storage adapters to recognize the LUNS.

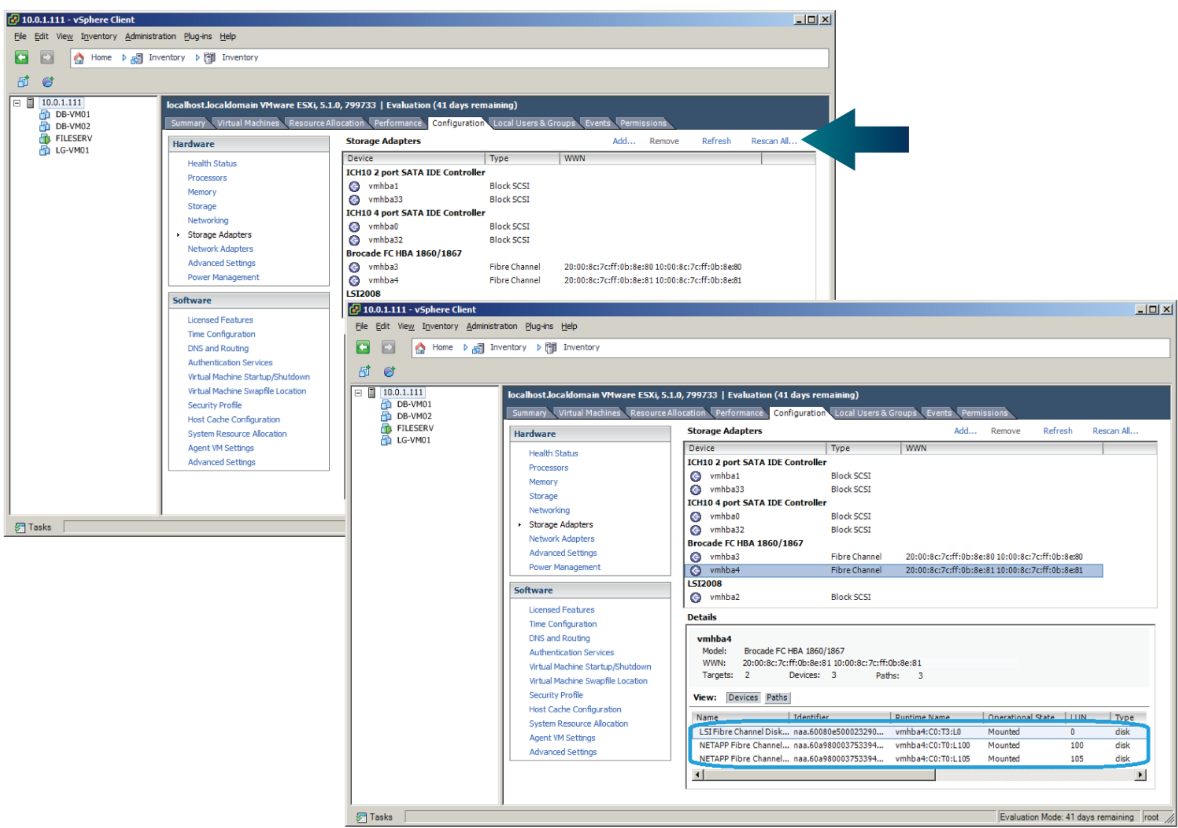

VMware ESX/ESXi supports two methods of configuring Fibre Channel LUNs for use by guest virtual machines. These methods are:

1. Virtual Hard Disk method

2. Pass-through method

Virtual Hard Disk

Using the virtual hard disk method, a Fibre Channel LUN is provisioned as a datastore via the vSphere client by accessing the Configuration tab for our ESX/ESXi host, Selecting "Storage" and then clicking on "Add Storage" on the right hand side of the screen.

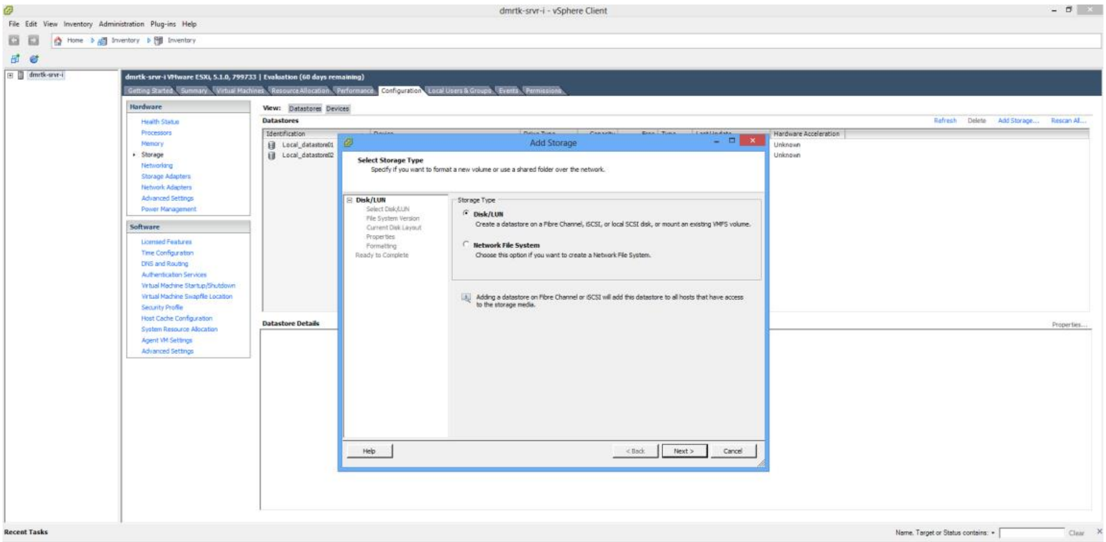

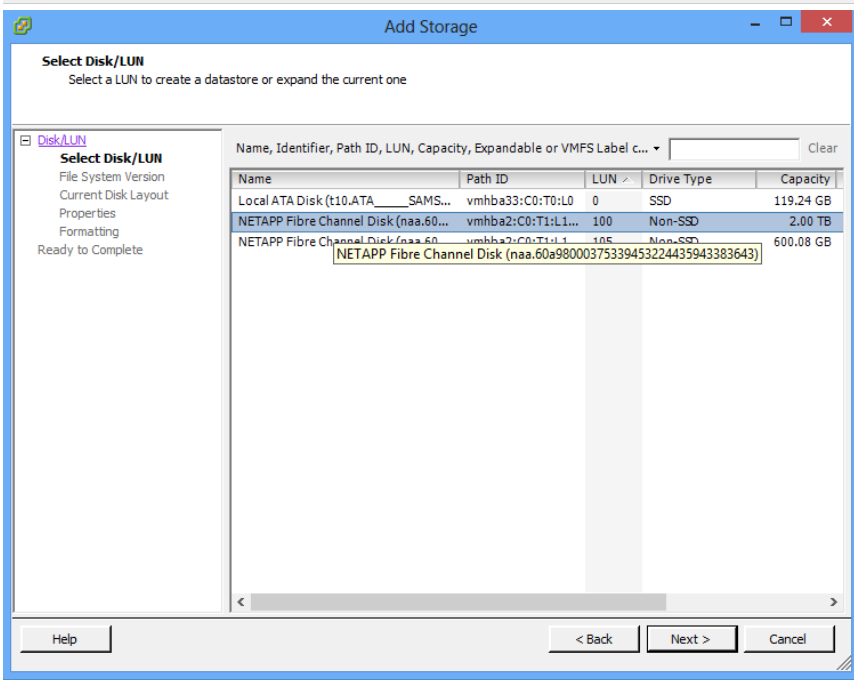

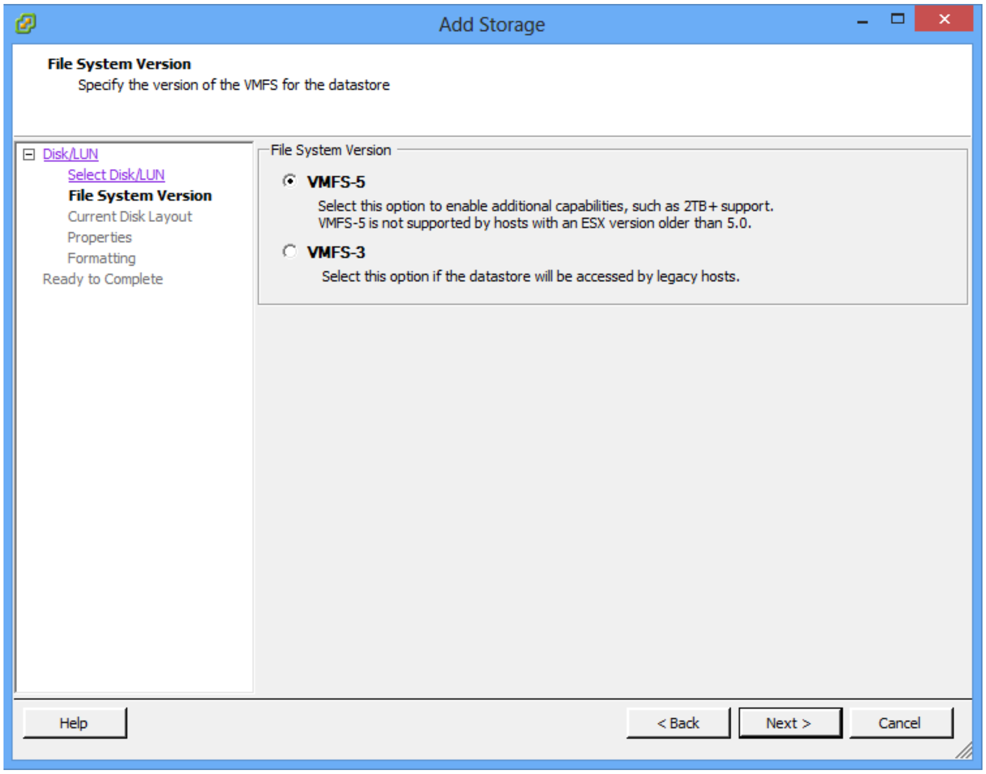
Add Storage
File System Version
Specify the version of the VMFS for the datastore
Disk/LUN
Select Disk/LUN
File System Version
Current Disk Layout
Properties
Formatting
Ready to Complete
File System Version
VMFS-5
Select this option to enable additional capabilities, such as 2TB+ support.
VMFS-5 is not supported by hosts with an ESX version older than 5.0.
VMFS-3
Select this option if the datastore will be accessed by legacy hosts.
Help
< Back
Next >
Cancel

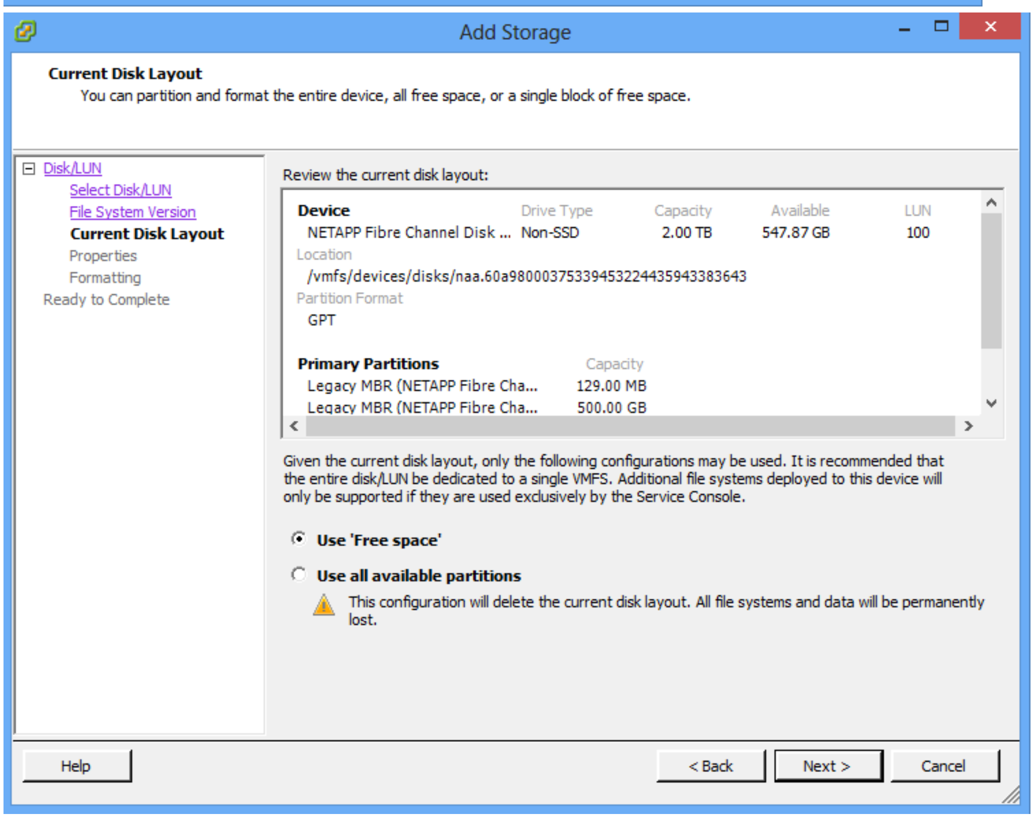
Add Storage
Current Disk Layout
You can partition and format the entire device, all free space, or a single block of free space.
Disk/LUN
Select Disk/LUN
File System Version
Current Disk Layout
Properties
Formatting
Ready to Complete
Review the current disk layout:
Device
Drive Type
Capacity
Available
LUN
NETAPP Fibre Channel Disk ...
Non-SSD
2.00 TB
547.87 GB
100
Location
/vmfs/devices/disks/naa.60a98000375339453224435943383643
Partition Format
GPT
Primary Partitions
Capacity
Legacy MBR (NETAPP Fibre Cha...
129.00 MB
Legacy MBR (NETAPP Fibre Cha...
500.00 GB
Given the current disk layout, only the following configurations may be used. It is recommended that the entire disk/LUN be dedicated to a single VMFS. Additional file systems deployed to this device will only be supported if they are used exclusively by the Service Console.
Use 'Free space'
Use all available partitions
This configuration will delete the current disk layout. All file systems and data will be permanently lost.
Help
< Back
Next >
Cancel

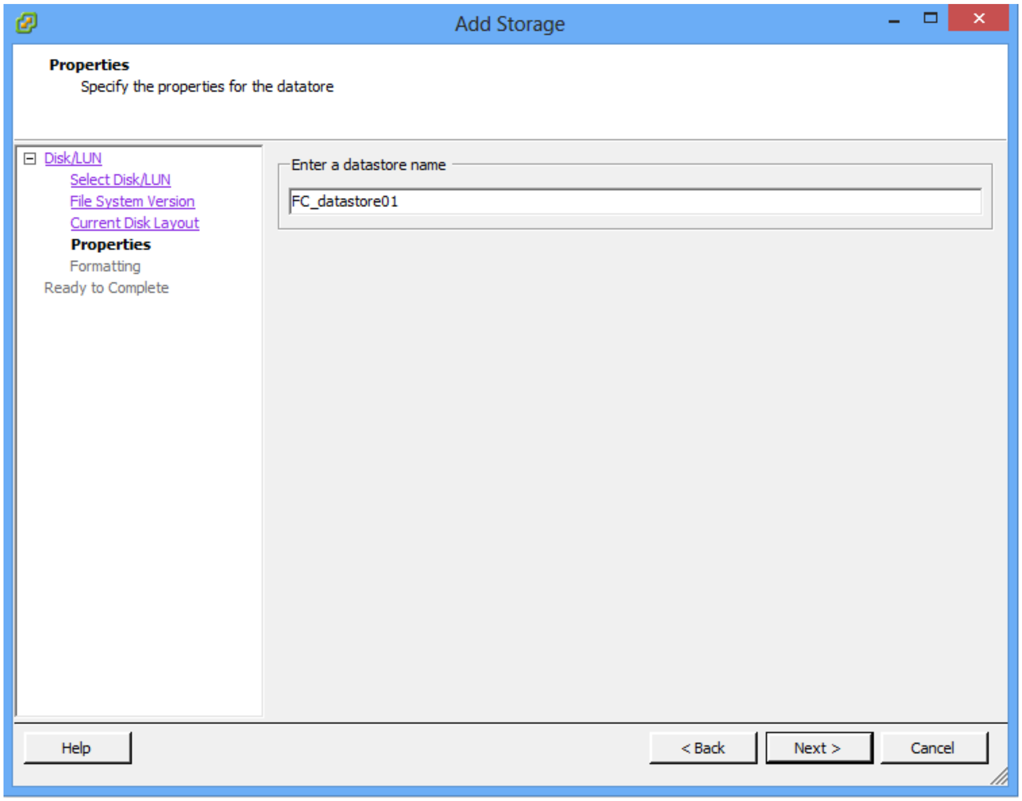
Add Storage
Properties
Specify the properties for the datatore
Disk/LUN
Select Disk/LUN
File System Version
Current Disk Layout
Properties
Formatting
Ready to Complete
Enter a datastore name
FC_datastore01
Help
< Back
Next >
Cancel

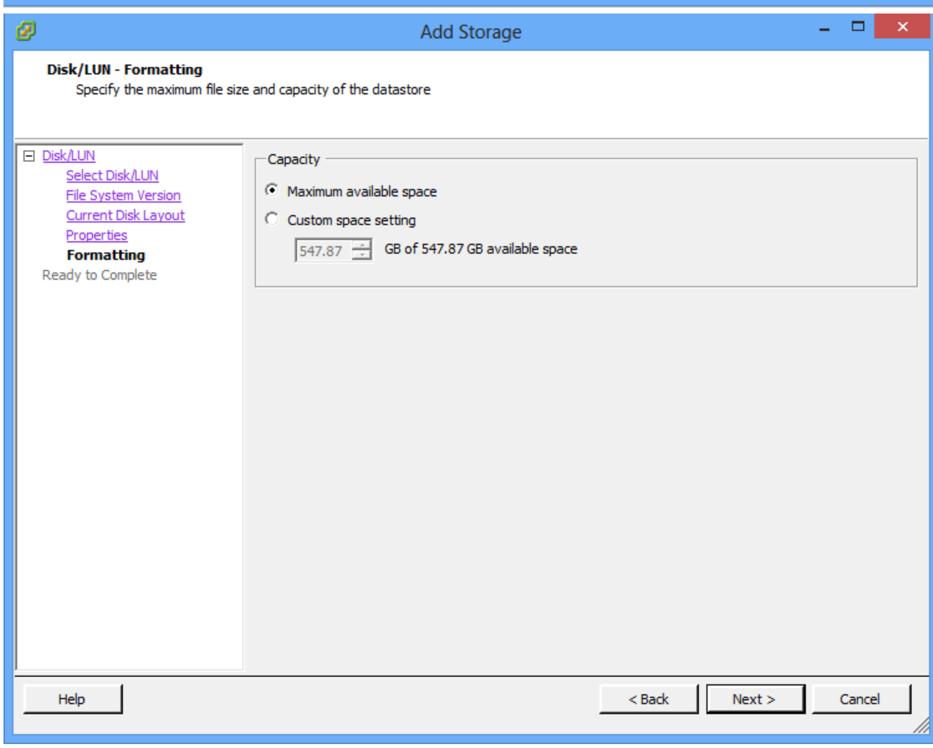
Add Storage
Disk/LUN - Formatting
Specify the maximum file size and capacity of the datastore
Disk/LUN
Select Disk/LUN
File System Version
Current Disk Layout
Properties
Formatting
Ready to Complete
Capacity
Maximum available space
Custom space setting
547.87
GB of 547.87 GB available space
Help
< Back
Next >
Cancel

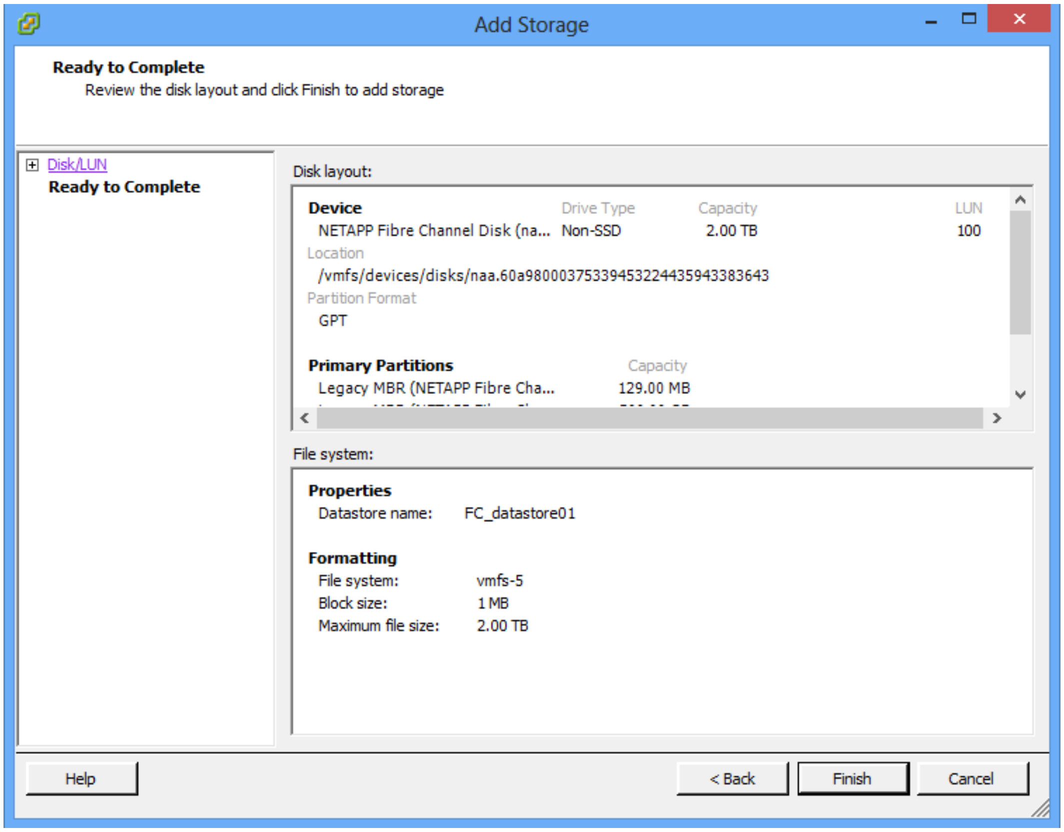

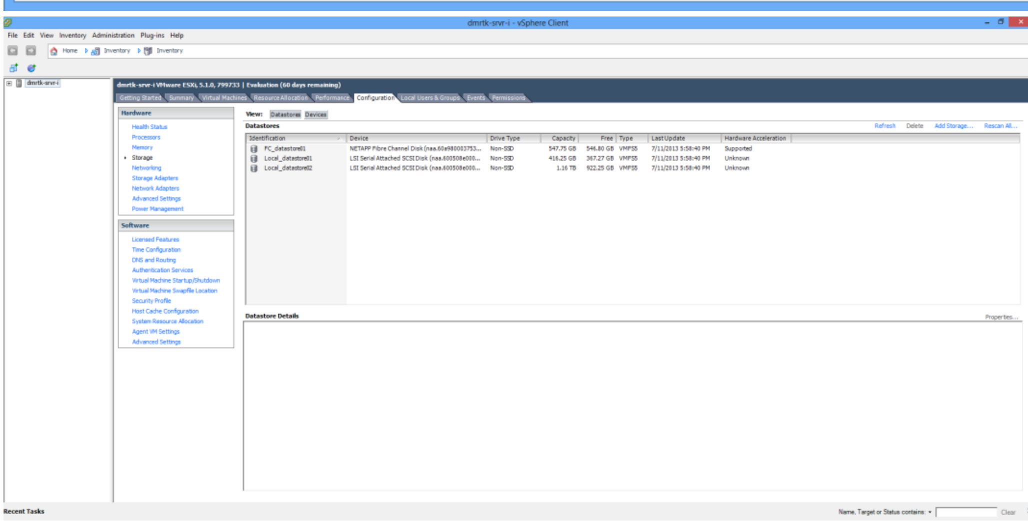

Then we created a new hard disk for the DB-VM01 VM by using the settings dialog box and clicking “Add” from the top left of the hardware tab.

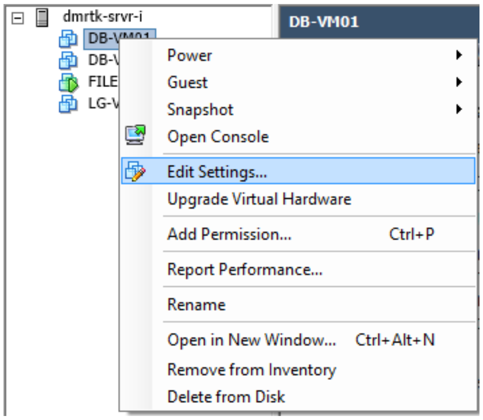
dmrtk-srvr-i
DB-VM01
Power
Guest
Snapshot
Open Console
Edit Settings...
Upgrade Virtual Hardware
Add Permission... Ctrl+P
Report Performance...
Rename
Open in New Window... Ctrl+Alt+N
Remove from Inventory
Delete from Disk

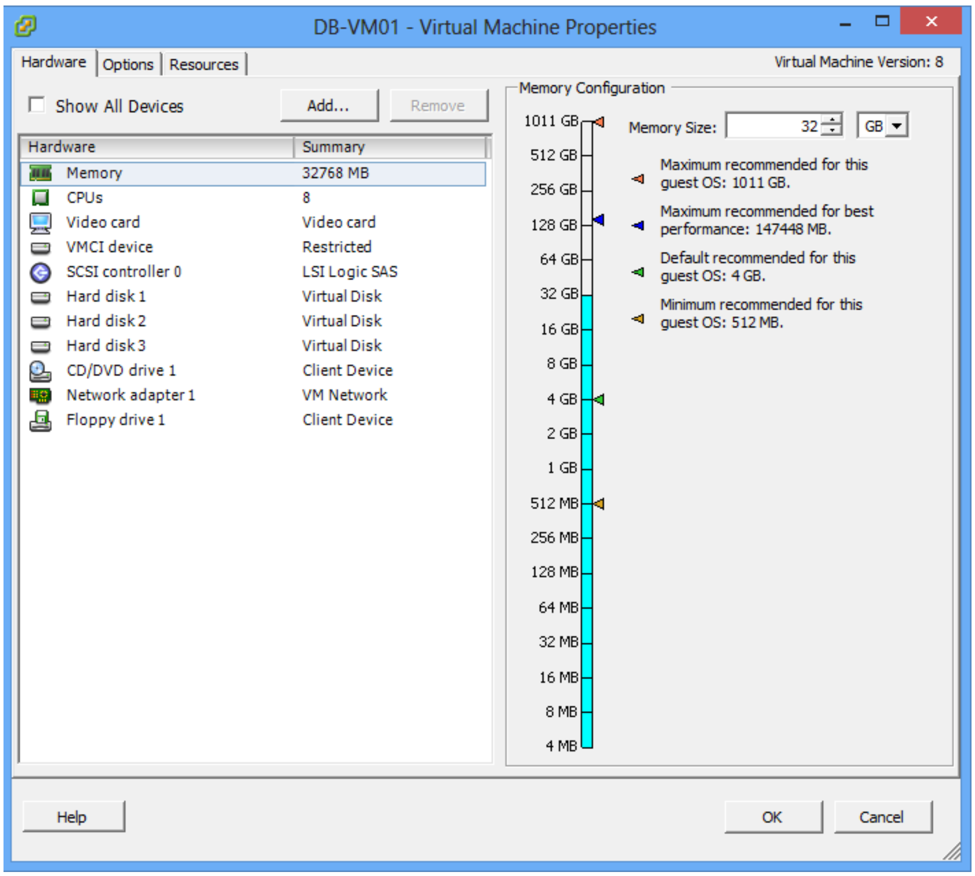
DB-VM01 - Virtual Machine Properties
Hardware
Options
Resources
Virtual Machine Version: 8
Show All Devices
Add...
Remove
Hardware
Summary
Memory 32768 MB
CPUs 8
Video card Video card
VMCI device Restricted
SCSI controller 0 LSI Logic SAS
Hard disk 1 Virtual Disk
Hard disk 2 Virtual Disk
Hard disk 3 Virtual Disk
CD/DVD drive 1 Client Device
Network adapter 1 VM Network
Floppy drive 1 Client Device
Memory Configuration
Memory Size: 32 GB
Maximum recommended for this guest OS: 1011 GB.
Maximum recommended for best performance: 147448 MB.
Default recommended for this guest OS: 4 GB.
Minimum recommended for this guest OS: 512 MB.
1011 GB
512 GB
256 GB
128 GB
64 GB
32 GB
16 GB
8 GB
4 GB
2 GB
1 GB
512 MB
256 MB
128 MB
64 MB
32 MB
16 MB
8 MB
4 MB
Help
OK
Cancel

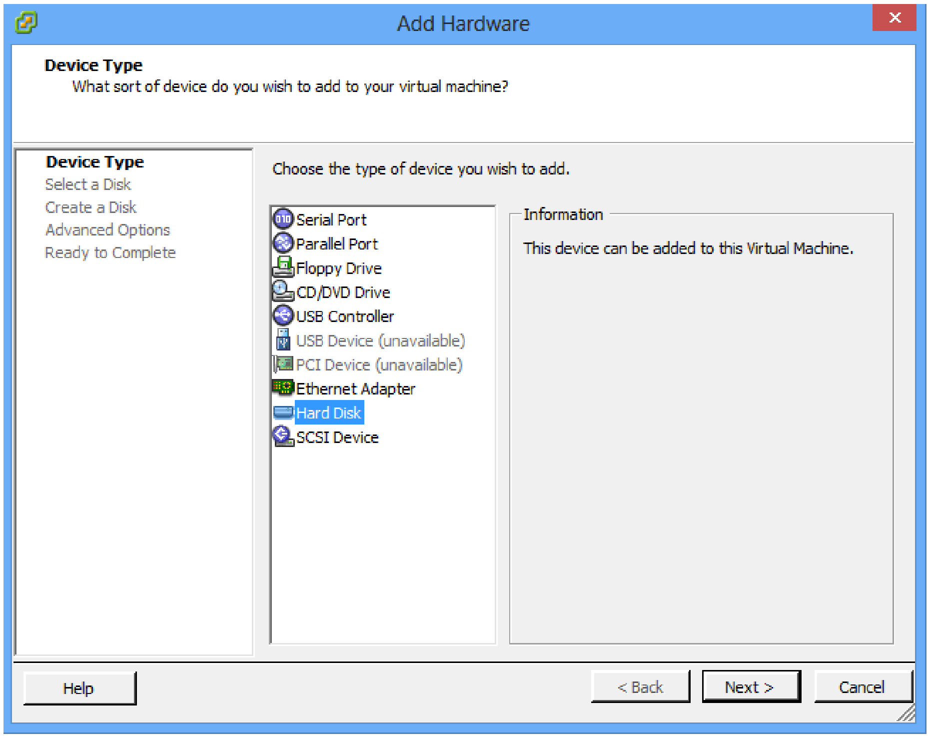

There are three choices for the type of VHD: Thick Provision Lazy Zeroed, Thick Provision Eager Zeroed, or Thin Provision.

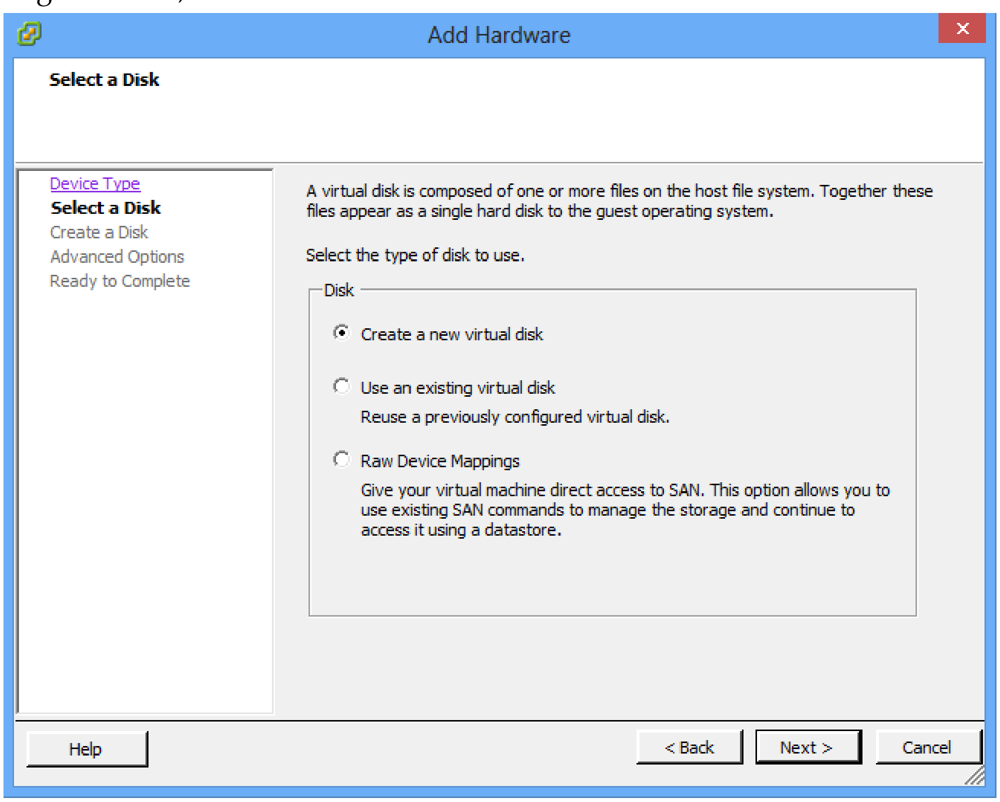

A datastore location can be selected. Using the Browse button brings up a list of datastores including our recently configured FC_datastore01.

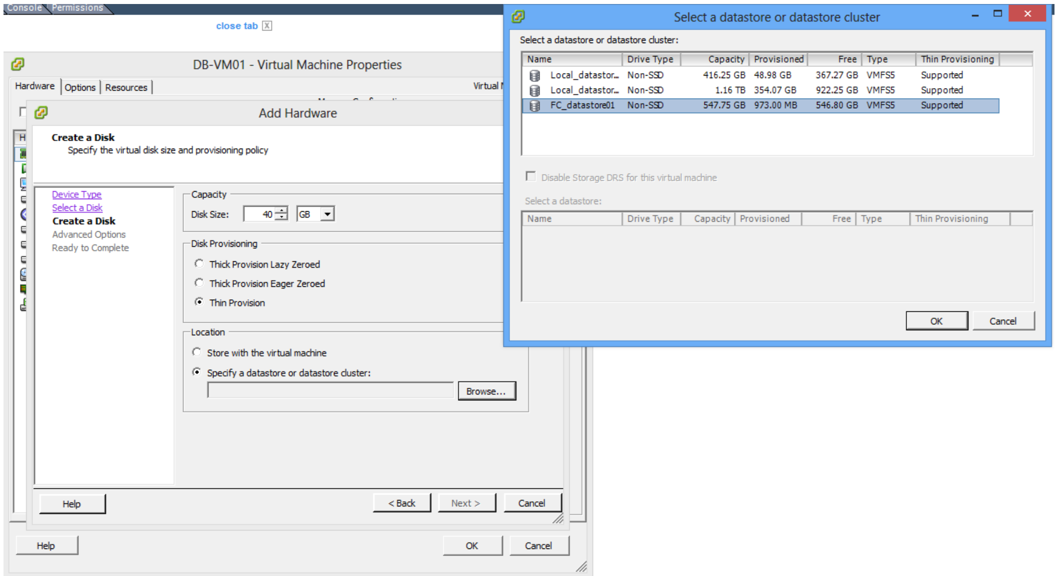

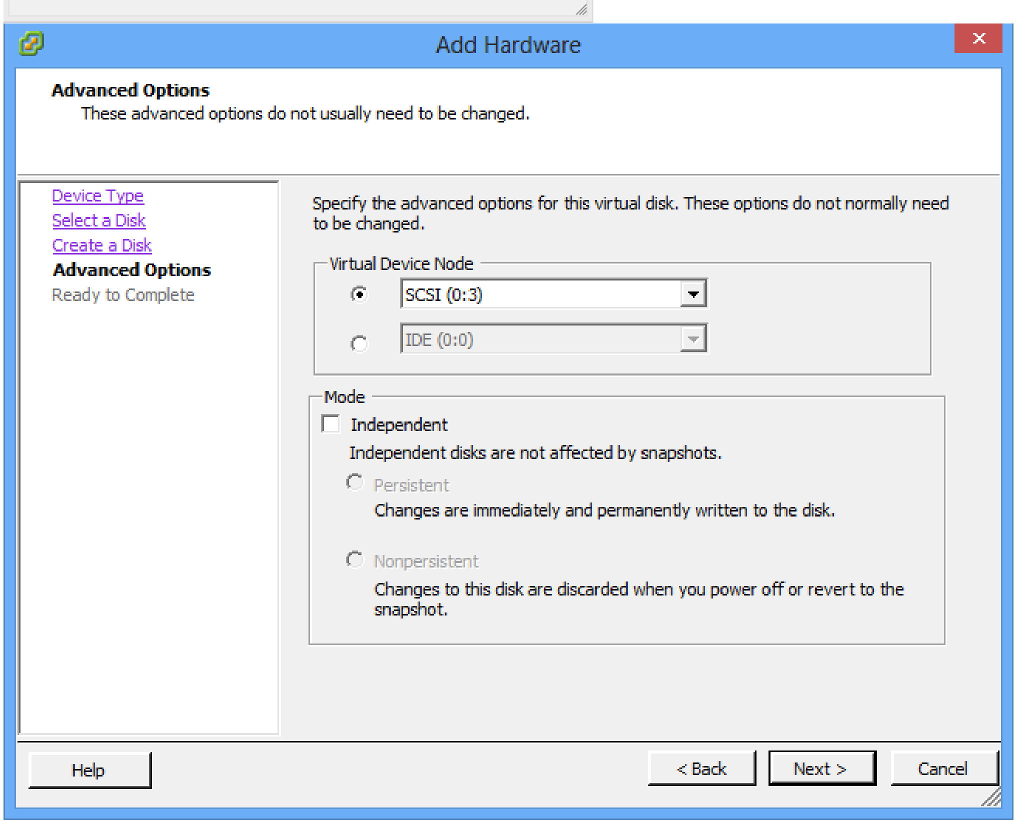

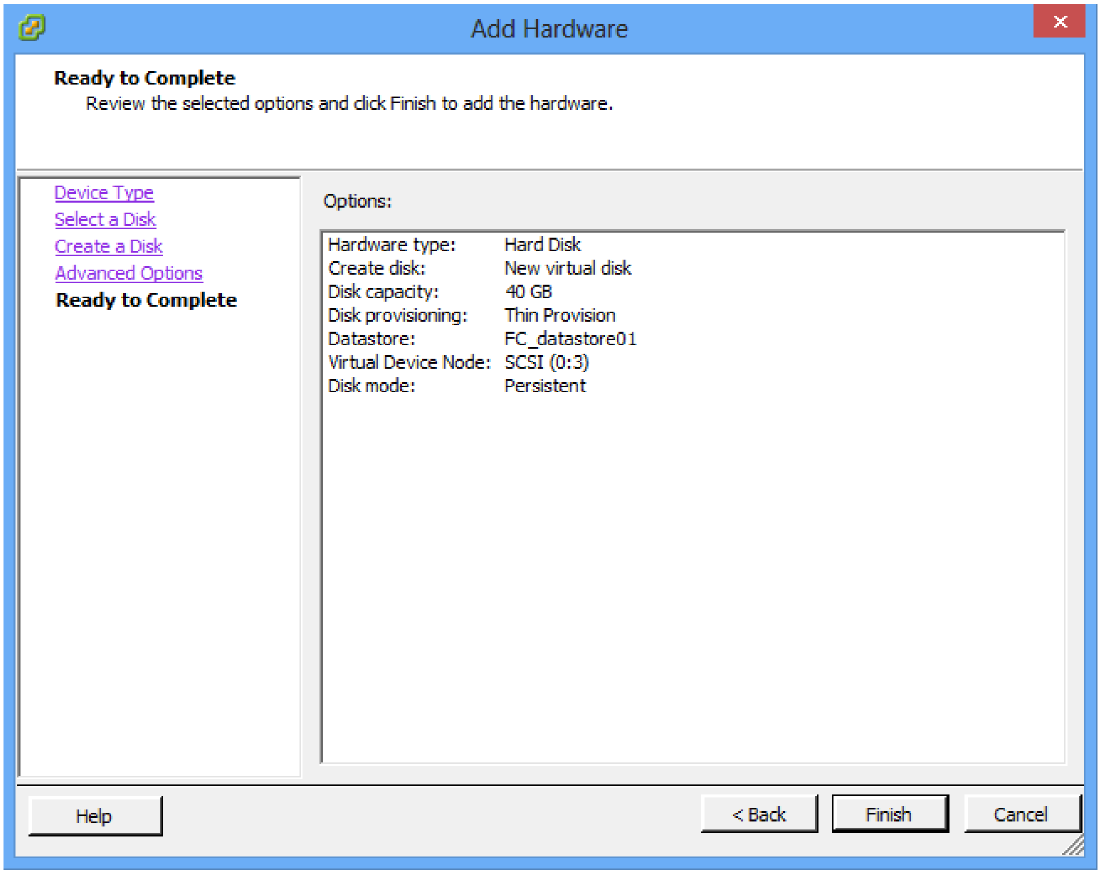
Add Hardware
Ready to Complete
Review the selected options and click Finish to add the hardware.
Device Type
Select a Disk
Create a Disk
Advanced Options
Ready to Complete
Options:
Hardware type: Hard Disk
Create disk: New virtual disk
Disk capacity: 40 GB
Disk provisioning: Thin Provision
Datastore: FC_datastore01
Virtual Device Node: SCSI (0:3)
Disk mode: Persistent
Help
< Back
Finish
Cancel

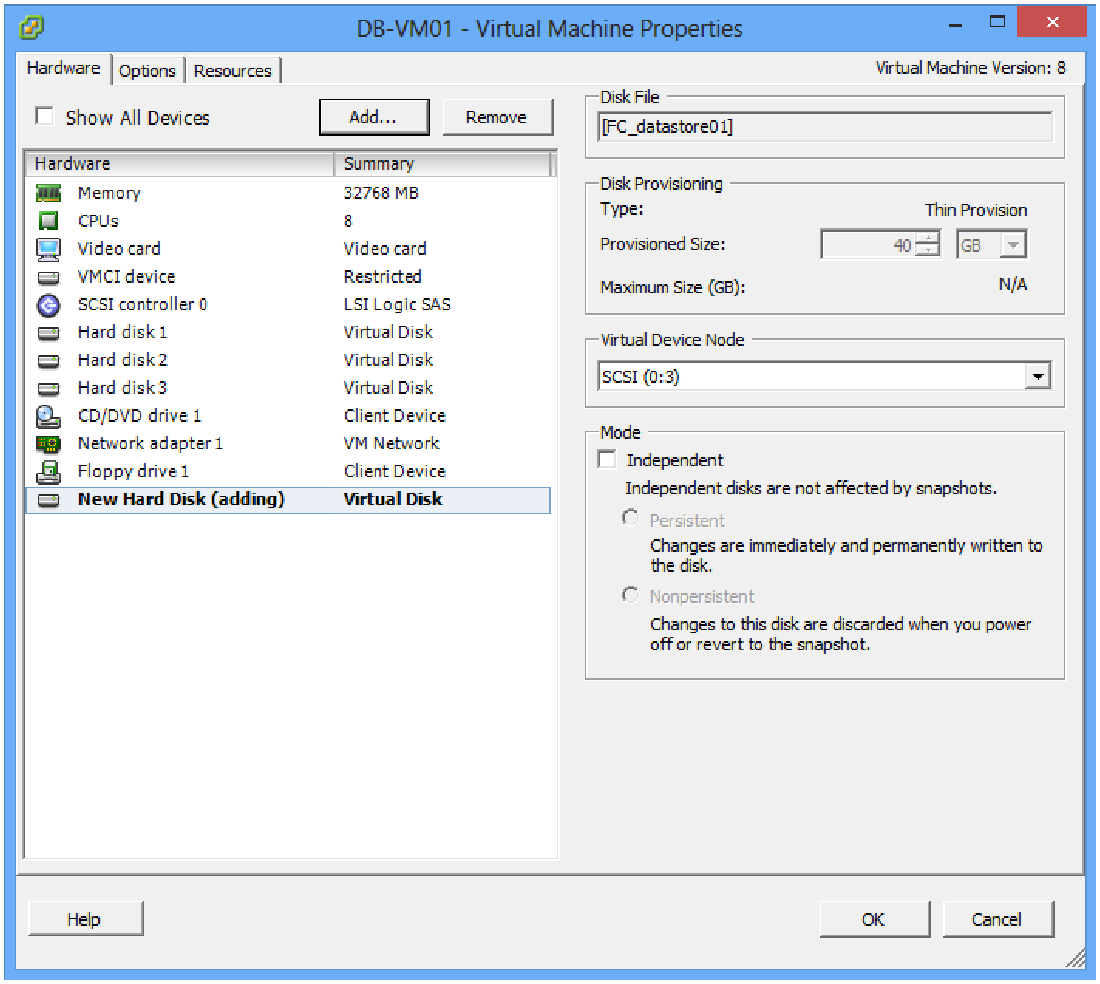
DB-VM01 - Virtual Machine Properties
Hardware
Options
Resources
Virtual Machine Version: 8
Show All Devices
Add...
Remove
Hardware
Summary
Memory
32768 MB
CPUs
8
Video card
Video card
VMCI device
Restricted
SCSI controller 0
LSI Logic SAS
Hard disk 1
Virtual Disk
Hard disk 2
Virtual Disk
Hard disk 3
Virtual Disk
CD/DVD drive 1
Client Device
Network adapter 1
VM Network
Floppy drive 1
Client Device
New Hard Disk (adding)
Virtual Disk
Disk File
[FC_datastore01]
Disk Provisioning
Type:
Thin Provision
Provisioned Size:
40
GB
Maximum Size (GB):
N/A
Virtual Device Node
SCSI (0:3)
Mode
Independent
Independent disks are not affected by snapshots.
Persistent
Changes are immediately and permanently written to the disk.
Nonpersistent
Changes to this disk are discarded when you power off or revert to the snapshot.
Help
OK
Cancel

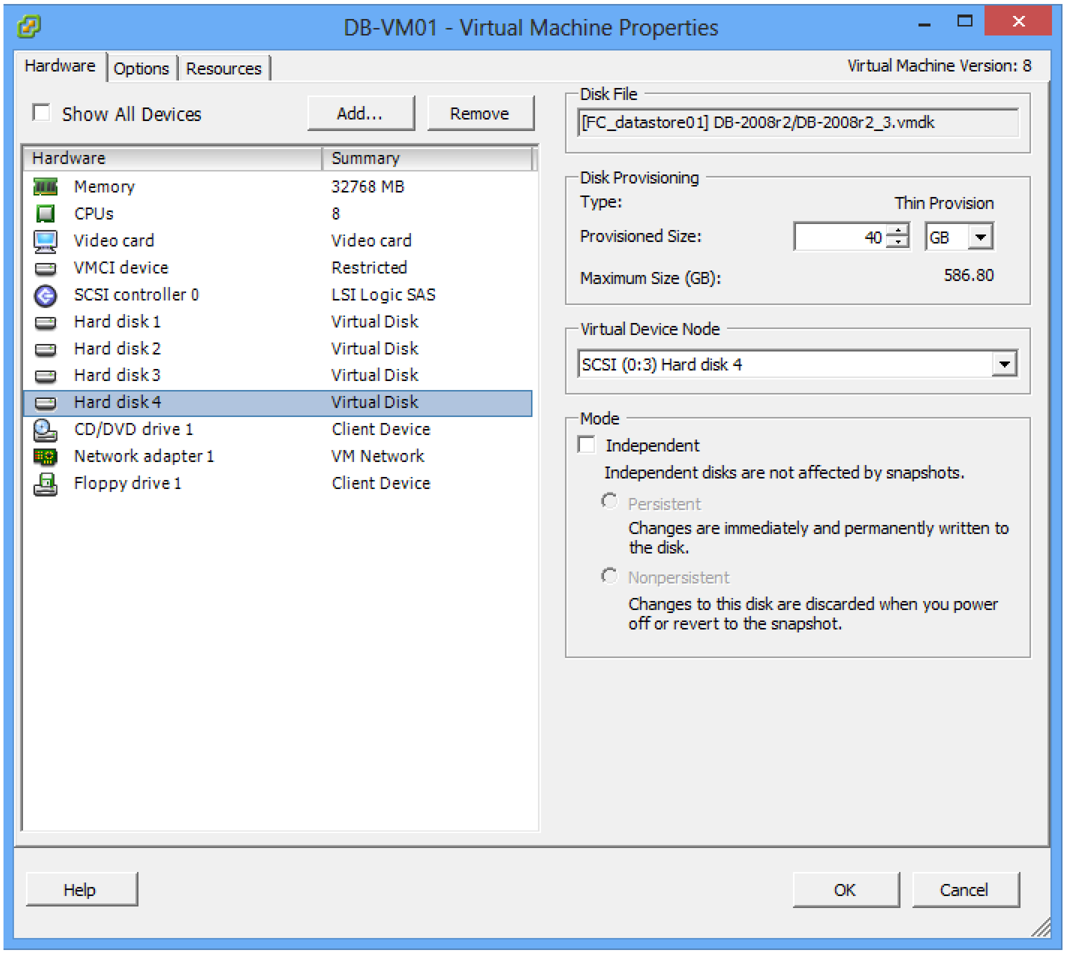

After starting that virtual machine, the new volume appears in Disk Manager for that virtual machine.

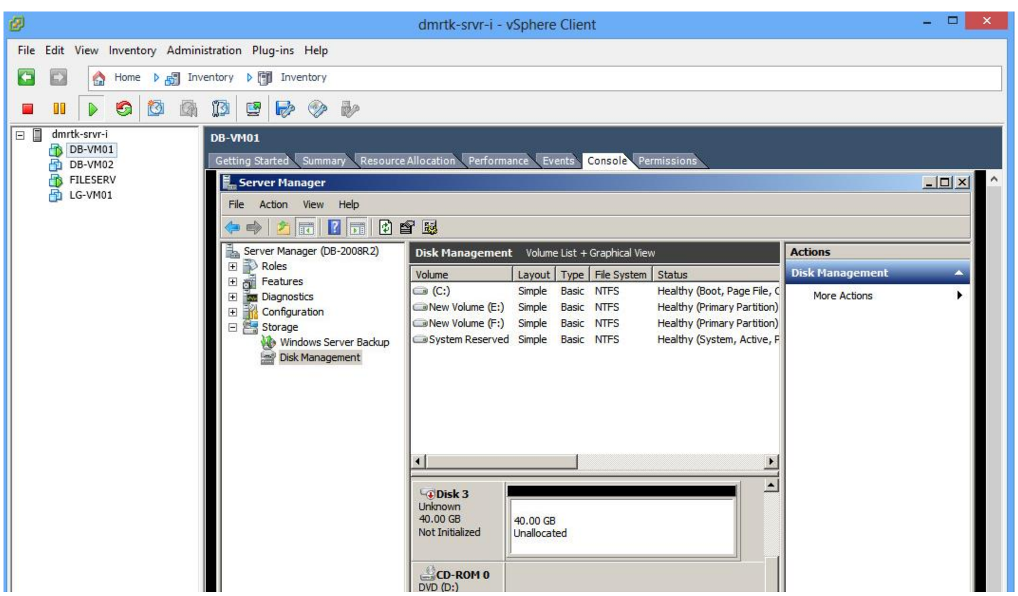

Pass-through (Raw Device Mapping)

A Fibre Channel LUN was made available to the hypervisor via the vSphere console in the usual manner. The LUN is then passed directly to the virtual machine. This process is similar to the virtual hard disk method above, except that this LUN will be allocated using “raw device mappings” and is mapped with physical device mappings that are stored with the virtual machine. Selected screenshots from this process are shown below.

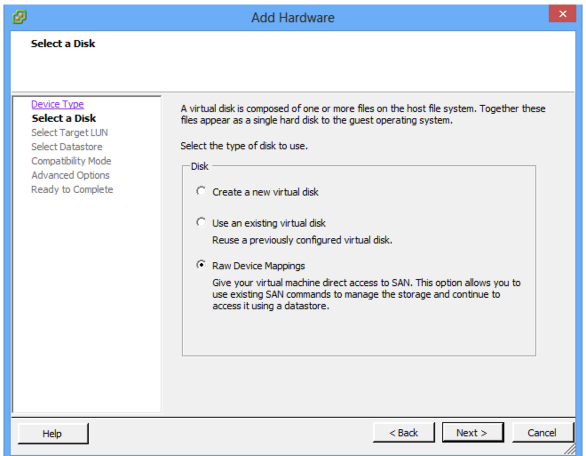

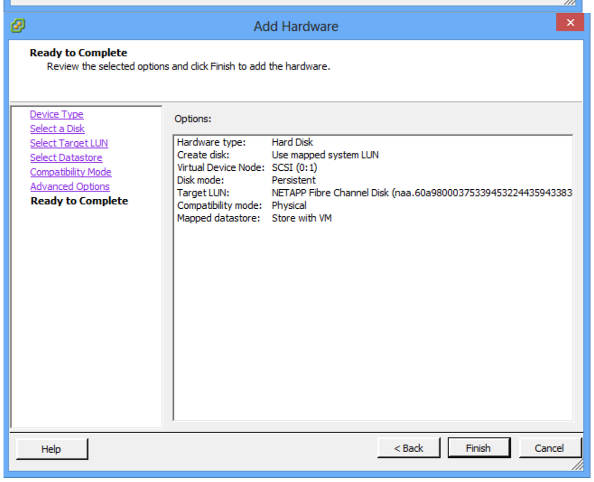

The LUN is now available to the virtual machine and is described as a “Mapped Raw LUN”.

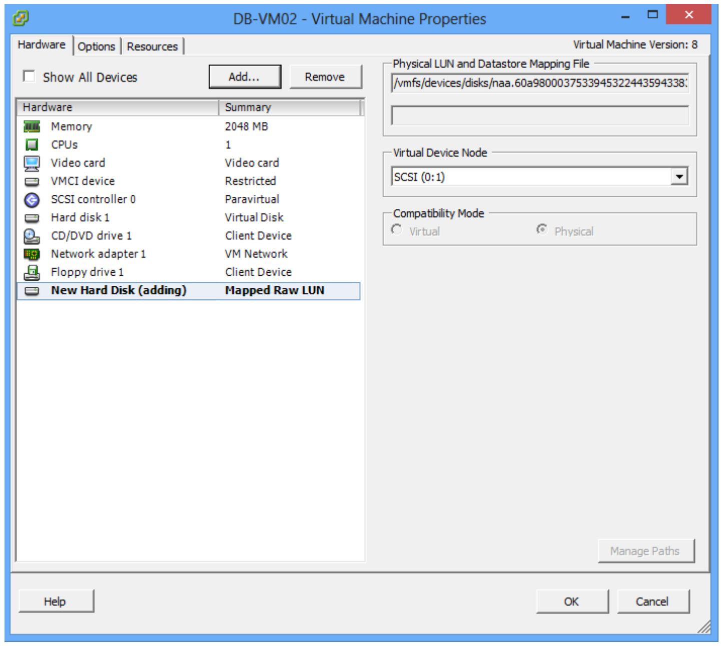

NPIV

NPIV features can be used in VMWare, enabling a virtual machine to use a virtual Fibre Channel port as long as the ESX/ESXi machine has a compatible FC HBA.

The host ESX/ESXi machine should be able to access the LUN that will eventually be assigned to your VM. The VM that will be using the virtual FC must first be given a RDM (Raw Device Mapped) Hard Drive using the instructions in the previous section. After this, access the settings for your VM by right clicking on the VM and accessing “Edit Settings.”

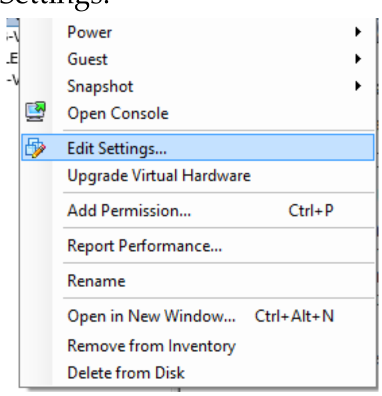

The Virtual Machine Properties screen will open with the Hardware tab open. Click on the “Options” tab, and “Fibre Channel NPIV” will be listed as one of the settings that can be modified.

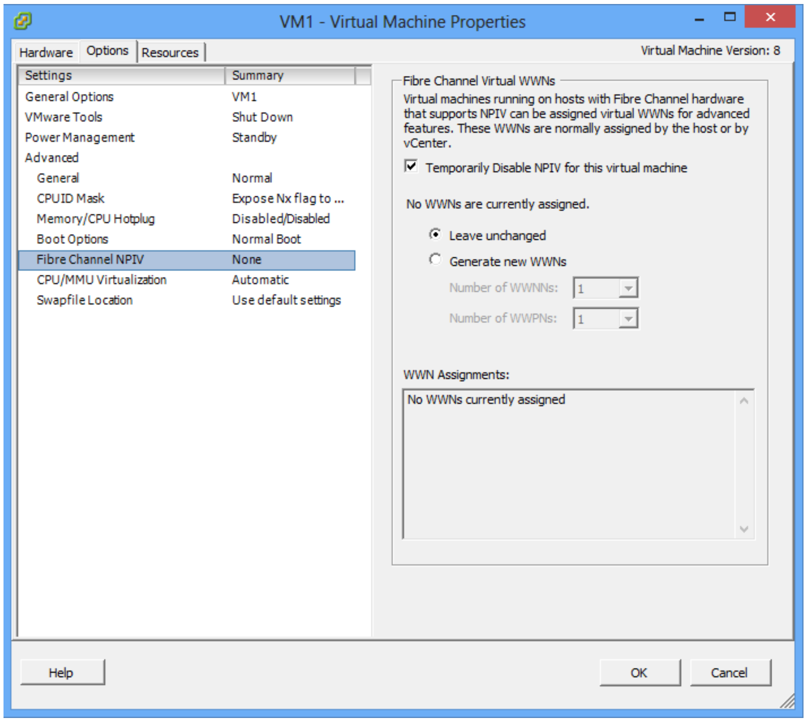

Clicking on the Fibre Channel NPIV will show a menu that tells us that NPIV is Temporarily Disabled. Click on the checkbox to enable it. Then click on the radio button next to Generate New WWNs so that the machine gets WWNs to use. Click ok to generate the numbers and the Virtual Machine Properties box will close.

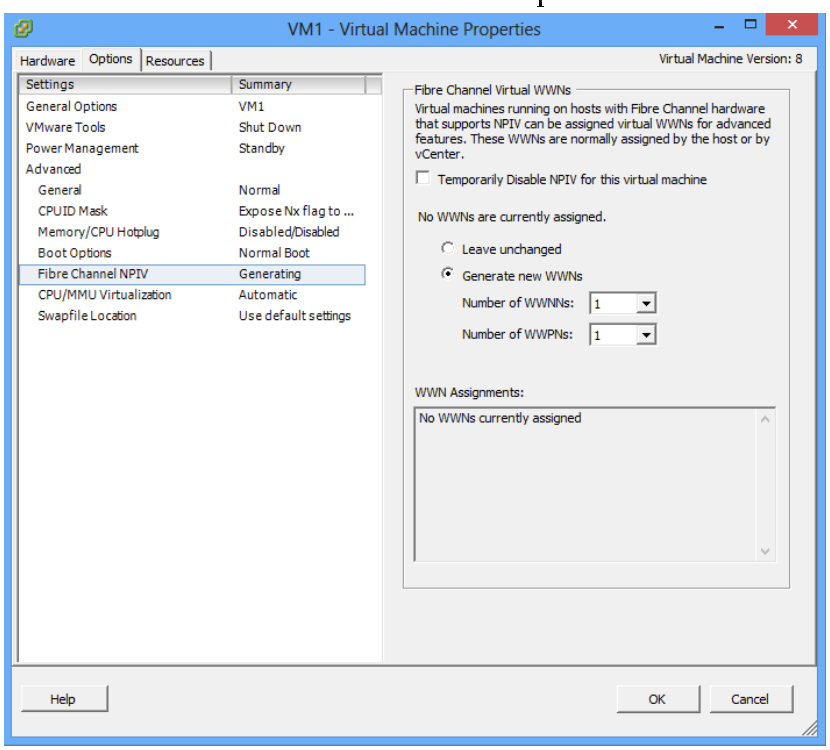

Re-access the same menu again to find out what WWNs were generated for Target configuration.

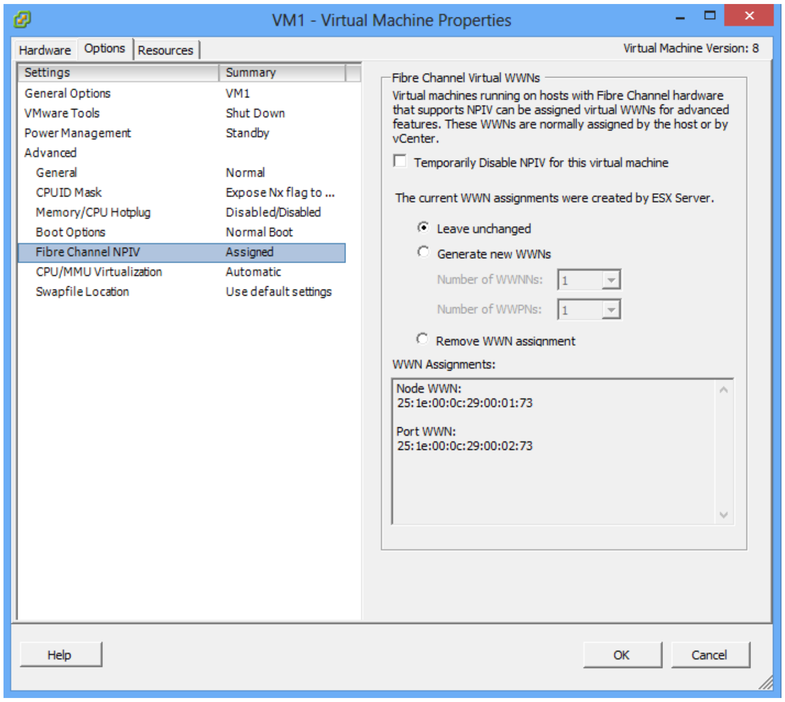

Once the target is configured, the VM can be booted and will obtain its LUN using the virtual port in the same way a bare metal installation would use an HBA to access a physical SAN.

Measuring 16Gb Fibre Channel Performance

Storage system performance can be measured by a variety of tools that are either included in the operating system or easily installable. Workload performance data in this guide was generated using readily available measuring and monitoring tools such as Windows Perfmon, Linux iostat and switch monitoring utilities. Additional open source and third party tools are also available, but beyond the scope of this guide.

Bandwidth, IOPS, and Latency

Bandwidth, IOPS and latency are the standard industry measurements for storage performance. Bandwidth is often expressed in units of data transferred (read or written) over a period of time, such as MB/s.

IOPS, inputs and outputs per second, is the total number of read and write requests serviced by the storage or interface during each second.

Latency is the time taken to service the I/O requests. Latency is commonly expressed in milliseconds, however, as solid state storage continues to improve, it is becoming more common to see sub-millisecond latency measurements. In our lab testing, we have found that Fibre Channel solutions generally have low latency (lower is better).

Windows Perfmon

Windows Performance Monitor (Perfmon) is a very granular measurement tool that is provided with the Windows operating system. The user can define performance counters for the measurements desired and either execute measurement s in real-time or batch measurements to coincide with workloads. There is a staggering array of performance counters available for metering system performance in Perfmon.

Linux sysstat

Linux sysstat is an add-on Linux package that can be installed into most Linux distributions. It contains a collection of individual command-line utilities such as “iostat” that provide various performance statistics for CPU, I/O devices, filesystems and more.

Workload Testing

A variety of synthetic I/O benchmarks exist to push I/O through the storage infrastructure. Open source tools like IOmeter and vdbench can be configured to drive multiple I/O patterns to simulate recognizable workload profiles or just force maximum bandwidth and IOPS. Several industry associations and companies also produce industry recognized benchmark kits that can be obtained in varying price ranges from free to several thousands of dollars.

Real-world applications may more closely align with actual use cases and are the workloads used for this guide. The following graph is from a workload application which produces database "data warehousing" I/O in a virtualized environment. The purpose of this workload is not to show that we can achieve line rate across 16Gb Fibre Channel, which can be easily done with any synthetic benchmark, but to instead demonstrate the limitation of the 8Gb Fibre Channel in a real-world I/O intensive environment.

Here we see the same database workload executed across 8 virtual machines. We can see that the dual channel, 8Gb Fibre Channel is frequently saturated by this workload as the bandwidth often plateaus around 1600 MBps.

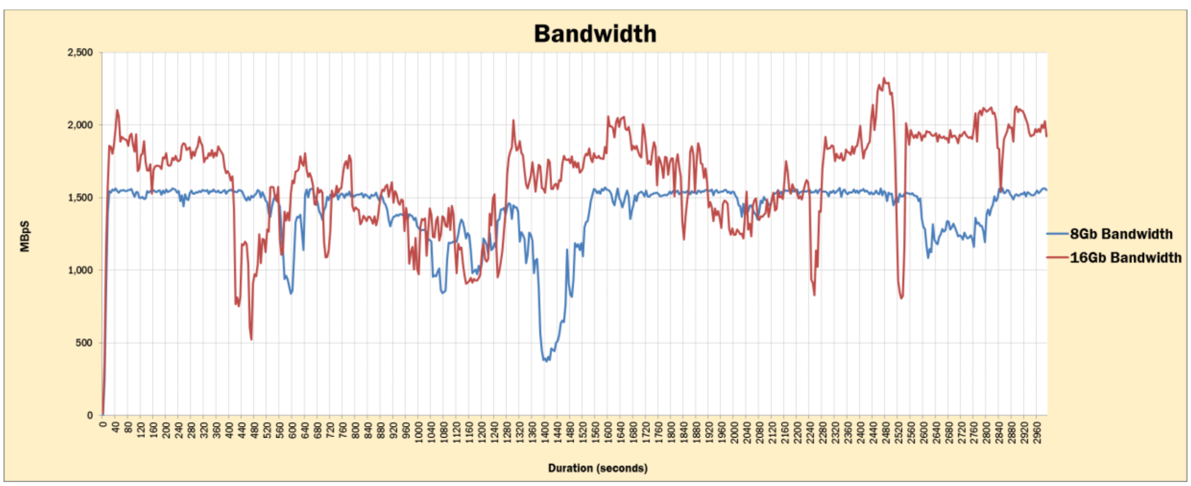

Simply switching the interface to dual channel, 16Gb Fibre Channel demonstrates this environment is capable of more work with the same compute resources. The bottleneck generated by the storage interface is gone and there is capacity for additional I/O if more compute resources are added to the environment, such as additional VMs running the workload, or more CPU and memory to the existing VMs.

The higher bandwidth and IOPS attainable by 16Gb FC compared to 8Gb FC has the potential to identify new performance bottlenecks. These could include a back-end storage system that may be incapable of generating 16Gbps performance. On the other hand, the higher bandwidth may allow the system to achieve its maximum processing capacity by fully utilizing the CPU or memory through the reduction of I/O waits from the storage system.

16Gb Fibre Channel Best Practices

Like all technology, proper deployment of 16Gb Fibre Channel is essential to achieving the best performance and reliability that can be delivered.

Multipath I/O

Single path I/O may be appropriate for non-production, non-critical functions, or environments in which redundancy is provided by some mechanism other than multipath I/O (MPIO). However, most production environments will want to take advantage of the redundancy and performance benefits provided by having more than one I/O channel.

Multiple paths can be configured at several points in the environment. However, if at any point in the system architecture a single path is introduced (including the use of a single 16Gb FC switch, even if multiple switch ports are zoned), then there is a single point of failure and overall redundancy cannot be guaranteed.

The simplest multipath configuration is to bypass the SAN altogether. For fully redundant multipath, two HBAs must be present in both the server and the storage target, regardless of whether they are single port or dual port adapters. A 16Gb FC port from each HBA on the server connects to corresponding ports on the two target HBAs. This deploys the FC target as direct attached storage, but does provide redundancy in case of FC port or HBA failures in both the server and the storage target. So long as all paths remain operational, multipathing can be configured to take advantage of both paths in an active-active fashion, essentially doubling the potential bandwidth.

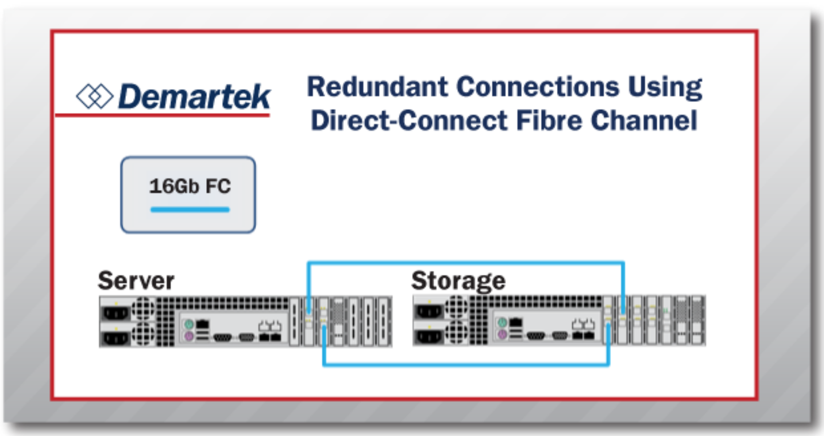

To enhance the robustness of this configuration, dual port HBAs can be used to further increase both redundancy and bandwidth. In theory, as many HBAs as are supported by the server and storage can be deployed in this fashion, though there will be a point of diminishing returns if the server, storage and workloads are not able to drive the aggregated FC bandwidth.

More likely, the intent will be to deploy FC storage in a 16Gb FC SAN, rather than directly attaching it to individual servers. The server and storage side configuration remains unchanged, however, two or more 16Gb FC switches are placed between them. Each separate HBA, on both the server and the storage, needs a port connected to each switch. If multiple ports are in use, it is best never to connect ports on the same HBA to the same switch in order to eliminate single points of failure.

An operating system will initially see multiple "copies" of each LUN or virtual disk provisioned to the server, one for each path. Each operating system has its own mechanism for enabling and configuring multipath management, usually a multipath driver, which resolves the multiple copies situation and which is outside the scope of this Deployment Guide. Detailed documentation is available from your operating system vendor.

Note that in the diagram on the right, each HBA is connected to each switch, and that there is no single point of failure from a SAN fabric perspective.

This deployment model allows for the addition of more servers and storage targets to the same SAN fabric by connecting them to available ports on the switches in a similar fashion.

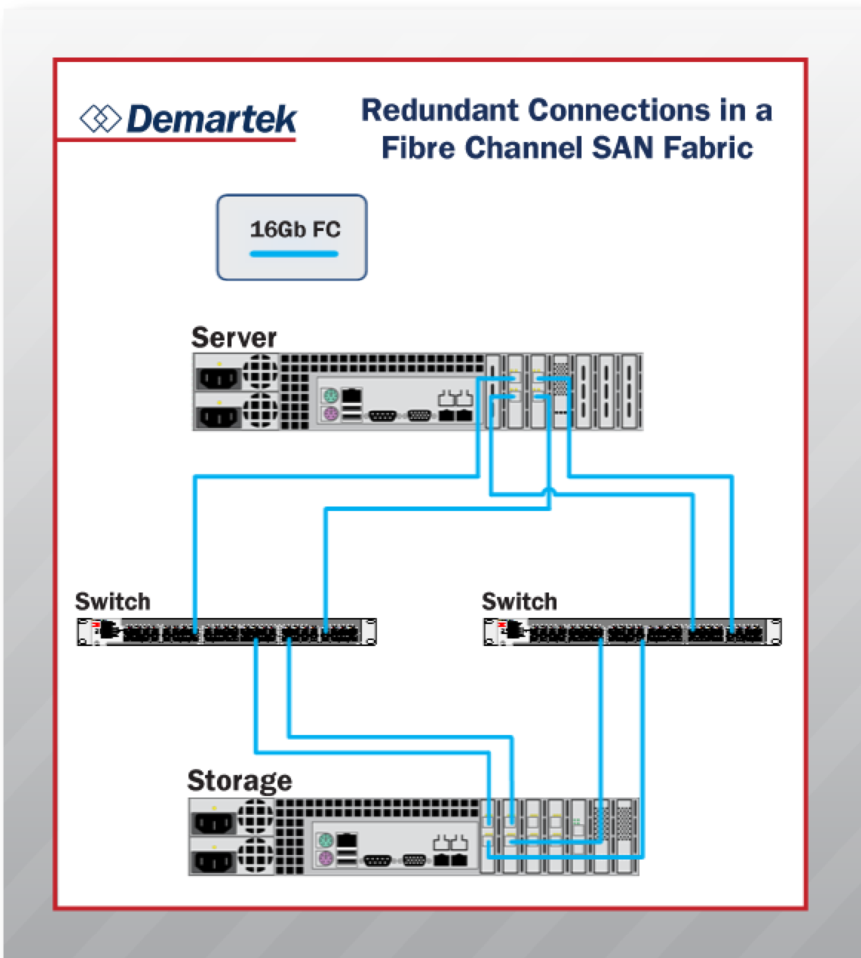

Increasing VM Density

Deploying 16Gb Fibre Channel in virtualized environments provides additional opportunities and considerations. VM density is a phrase that refers to the number of VMs supported by a hypervisor host or cluster. Density is governed by the amount and performance of compute power (CPU and memory), network resources (Ethernet and SAN), and backend storage available to the VM environment. When any component is utilized to its maximum capacity, the number of VMs that can be deployed is capped at that point.

In the case of a storage I/O bottleneck, such as was demonstrated earlier in this guide, the hypervisor host is not able to fully utilize the compute power available to it. Therefore, removing that constraint by deploying 16Gb FC allows the environment to use more of its compute resources. Theoretically, more VMs could be deployed to take advantage of the additional bandwidth increasing the VM density.

However, removing one bottleneck can highlight another. By removing storage I/O constraints, the processors and memory are able to do more work, potentially enough to fully utilize those resources without the addition of more VMs.

Legal and Trademarks

The latest version of this report is available at http://www.demartek.com/Demartek_FibreChannel_Deployment_Guide.html on the Demartek website.