

CITRIX XENSERVER AND NETAPP STORAGE BEST PRACTICES

Citrix Systems, Inc. and NetApp, Inc.

May 2008 | Version 1.0



TABLE OF CONTENTS

The Challenge of Today’s Enterprise.....	3
Citrix XenServer for Enterprise-Ready Virtualization.....	3
NetApp Storage Solutions for Enterprise-Ready Virtualization.....	3
Overview of XenServer storage	4
Storage repositories (SRs).....	4
Virtual Disk Images (VDIs).....	4
Managing Storage	4
XenServer Shared Storage Options	5
Configuration and Setup	9
NetApp Filer Configuration Basics.....	9
XenServer Shared Storage Configuration Basics.....	13
Backup and Recovery	30
Snapshot for Virtual Disk on NetApp Data ONTAP and NFS Storage Repositories	30
Backing up Storage Repositories on NetApp Filer	32
Recovery of Snapshot Virtual Disks.....	32
Summary	34
Appendix A	35
Script to perform Snapshot	35
Appendix B	36
Security Considerations.....	36
Acknowledgements.....	36

THE CHALLENGE OF TODAY'S ENTERPRISE

IT departments face the constant demand to respond more rapidly to changing business priorities, application demands, and user dynamics – all without compromising security or manageability or increasing server count. They must deliver robust data management, business uptime, and complete backup and recovery capabilities. In order to address these challenges, enterprises need to:

- Standardize on a reduced number of hardware configurations – drastically reducing the time required to deploy upgrades and patches
- Provide effective, near-term high-availability for cost-effective protection against physical server and storage failures in an application-agnostic manner
- Adjust allocation of server and storage resources for different application workloads on the fly
- Consider scalability for the largest enterprise applications
- Deploy a tightly unified server and storage virtualization solution that is reliable, not overly complex and leverages all available capabilities

This document presents best practices for using NetApp® storage systems as networked attached storage solutions for Citrix® XenServer™.

CITRIX XENSERVER FOR ENTERPRISE-READY VIRTUALIZATION

Citrix XenServer is a native 64-bit virtualization platform, with the scalability required by Microsoft Exchange Server, Microsoft SQL Server, Citrix Presentation Server, and other business-critical applications. The highest host and guest CPU and memory limits available, coupled with fine-grained controls for processor, network, and disk, enable it to deliver optimal quality of service. With Citrix XenServer virtualization, businesses can increase server and storage utilization, reducing costs of equipment, power, cooling, and real estate. By combining servers and storage into resource pools that can be apportioned to the applications with the highest business need, IT operations can be aligned to changing demands and business priorities. With XenMotion™, running virtual machines can be migrated to new servers with no service interruption, allowing essential workloads to get needed resources and enable zero-downtime maintenance. Citrix XenServer products install directly on bare metal servers, requiring no dedicated host operating system. Open command-line (CLI) and programming (API) interfaces make it easy for companies to integrate Citrix XenServer virtualization with existing processes and management tools, rather than requiring rip-and-replace reimplementations. Key benefits and features include:

- Enterprise ready performance and scalability
- Simple deployment and installation
- Flexible shared infrastructure
- On-demand deployment of Windows and Linux virtual machines
- Powerful storage management
- Efficient, secure virtual networking
- Live migration
- XenCenter™ multi-server management, included with product
- Deliver server workloads on demand via streaming

NETAPP STORAGE SOLUTIONS FOR ENTERPRISE-READY VIRTUALIZATION

Unified storage solutions from NetApp complement the manageability, utilization and cost-saving benefits of Citrix XenServer. NetApp solutions enable powerful thin provisioning, simplified data management and scalable and consistent I/O performance for all IT applications across NAS, Fibre Channel and iSCSI SAN in a single pool of storage. Key benefits and features include:

- Supports SAN, IP-SAN, or NAS
- Scale non-disruptively to 100's of TB
- Easily installed, configured, managed, and maintained
- Rapid backup and recovery with zero penalty snapshots
- Simple, cost effective replication for Disaster Recovery
- Instant space efficient data clones for provisioning and testing
- Dynamically expand and contract storage volumes as needed
- Data deduplication to reduce capacity requirements

NetApp storage solutions offers these powerful data management and data protection capabilities allowing Citrix XenServer users the ability to lower cost while meeting their capacity, utilization, and performance requirements.

OVERVIEW OF XENSERVER STORAGE

STORAGE REPOSITORIES (SRs)

The XenServer host accesses containers named Storage Repositories (SRs) in which Virtual Disk Images (VDIs) are stored. A VDI is a disk abstraction which, when attached to a host, appears as a physical disk drive to the Virtual Machine.

The interface to storage hardware provided on the XenServer host allows VDIs to be supported on a large number of different SR substrate types. VDIs may be files on a local disk, on an NFS share, Logical Volumes within a LUN or a raw LUN itself directly attached to the VM. The SR and VDI abstractions on the host provide for advanced storage features such as sparse provisioning, image snapshots, and fast cloning to be leveraged on storage targets where available.

Each XenServer host can access multiple SRs in parallel of any type. These SRs can be shared between a pool of hosts, or a dedicated repository to a single. Shared storage must be accessible to all hosts in a resource pool, and is utilized for both non-live and live migration of VMs via XenMotion. When hosting shared Storage Repositories on a Netapp filer, there are 4 options; the optimized Netapp storage adapter, an NFS fileshare, an iSCSI LUN or a Fibre Channel LUN.

VIRTUAL DISK IMAGES (VDIs)

There are three fundamental VDI types (Files, LUNs and Logical Volumes) that can be accessed with a NetApp filer as the backend over 4 different SR driver types:

- **NetApp managed LUNs:** Managed NetApp LUNs are accessible via the NetApp Data ONTAP SR type, and are hosted on a NetApp filer running a version of Data ONTAP 7.0 or greater. LUNs are allocated on demand via the XenServer management interface and mapped dynamically to the host via the XenServer host management framework (using the open iSCSI software initiator) while a VM is active. All the thin provisioning and fast clone capabilities of the filer are exposed via the Netapp Data ONTAP adapter.
- **VHD files:** The VHD format can be used to store VDIs on an NFS share exported from a Netapp filer in a sparse format. Being sparse, the image file grows proportionally to the number of writes to the disk by the the Virtual Machine (VM), so large portions of the disk which are typically unused do not consume unnecessary space. VHD on NFS SRs can be shared to all hosts in a pool.
- **Logical Volume Manager (LVM):** Shared storage can be provided using a Logical Volume Manager layered over either an iSCSI or a Fibre Channel LUN hosted on a NetApp filer. The filer must be configured to provide access to the LUN by multiple iSCSI initiators or Fibre Channel HBAs.

The section entitled 'XenServer Shared Storage Options' discusses each option in more detail.

MANAGING STORAGE

In working with the XenServer host CLI, there are four object classes that are used to describe, configure, and manage storage:

- **Storage Repositories (SRs)** are storage targets containing homogeneous virtual disks (VDIs). SR commands provide operations for creating, destroying, resizing, cloning, connecting and discovering the individual Virtual Disk Images (VDIs) that they contain. A storage repository is a persistent, on-disk data structure. So the act of "creating" a new SR is similar to that of formatting a disk -- for single LUN-based SR types, i.e. LVM over iSCSI or Fibre Channel, the creation of a new SR involves erasing any existing data on the specified LUN. The exception is the Netapp Data ONTAP adapter which allocates it's own FlexVol resources on the filer and NFS SRs, which create a new directory on the filer leaving existing NFS SRs as they were. SRs are long-lived, and may in some cases be shared among XenServer Hosts, or moved between them. CLI operations to manage storage repositories are described in the section called "SR commands".
- **Physical Block Devices (PBDs)** represent the interface between a physical server and an attached SR. PBDs are connector objects that allow a given SR to be mapped to a XenServer Host. PBDs store the device configuration fields that are used to connect to and interact with a given storage target. For example, NFS device configuration includes the IP address of the NFS server and the associated path that the XenServer Host mounts. PBD objects manage the run-time attachment of a given SR to a given XenServer Host. CLI operations relating to PBDs are described in the section called "PBD commands".
- **Virtual Disk Images (VDIs)** are an on-disk representation of a virtual disk provided to a VM. VDIs are the fundamental unit of virtualized storage in XenServer. Similar to SRs, VDIs are persistent, on-disk objects that exist independently of XenServer Hosts. CLI operations to manage VDIs are presented in the section called "VDI commands".
- **Virtual Block Devices (VBDs)** are a connector object (similar to the PBD described above) that allows mappings between VDIs and Virtual Machines (VMs). In addition to providing a mechanism to attach (or plug) a VDI into a VM, VBDs allow the fine-tuning of parameters regarding QoS (quality of service), statistics, and the bootability of a given VDI. CLI operations relating to VBDs are described in the section called "VBD commands".

XENSERVER SHARED STORAGE OPTIONS

When using a NetApp filer as your networked, backend storage array, it is recommended with XenServer v4.1 onwards to use the NetApp Data ONTAP SR type. The customized architecture of the NetApp Data ONTAP adapter allows XenServer to utilize the capabilities of the NetApp filer to provide data efficiency, performance and ensure compatibility with existing ONTAP filer management tools. This allows for:

- Fast provisioning of VDIs
- Fast cloning of VDIs
- Fast Snapshot® of VDIs

To use non-customized storage options with a Netapp filer, the following types can also be used:

- Network Attached Storage using NFS
- iSCSI
- Fibre Channel

The following sections give an overview of the above storage types and the benefits associated with them. All shared storage options enable VM agility using XenMotion -- VMs can be started on any XenServer host in a resource pool and migrated between them.

Shared Storage using NetApp Data ONTAP

The XenServer adapter for NetApp Data ONTAP uses the Zephyr API (ZAPI) interface to the filer to create a group of FlexVol®s which correspond to a XenServer SR. VDIs are created as virtual LUNs on the filer, and attached to XenServer hosts using an iSCSI data path. There is a direct mapping between a VDI and a raw LUN without requiring any additional volume metadata. Thus, at a logical level, the NetApp SR is a managed volume and the VDIs are the LUNs within the volume.

Since the adapter exposes LUNs to the XenServer host as VDIs, there is a one-to-one mapping between the Virtual Machine disks in the XenServer environment and the Netapp disk data abstraction. This enables much simpler integration with existing physical host backup and array provisioning tools that already operate at a LUN granularity.

For the I/O data path, the NetApp Data ONTAP adapter directly controls the host built-in software initiator and its assigned server iSCSI Qualified Name or IQN to map the data path on demand without requiring user intervention.

Storage can be thinly or fully provisioned depending on the administrative preference for the Netapp array. When thin provisioning is utilized, data de-duping can also be switched on to reclaim common blocks between VM images, thereby conserving even more space. All these configuration options are exposed via the XenServer storage configuration interface and managed by the Netapp Data ONTAP adapter directly.

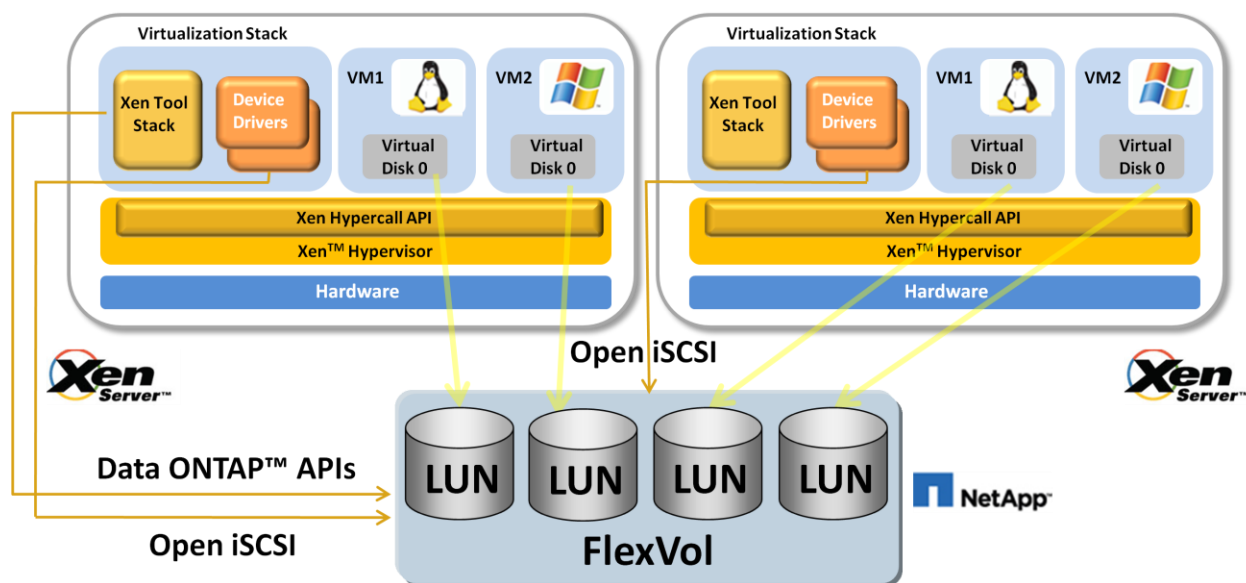


Figure 1 Shared storage using NetApp Data ONTAP

Shared NAS using NFS

XenServer supports shared access to Virtual Disk storage exported as NFS v3 over TCP/IP based on the Microsoft VHD format.

VDIs stored on NFS are *sparse*. The image file grows as the Virtual Machine (VM) writes data into the disk, so unused portions of the virtual disk do not consume space on the array. This is a considerable benefit since VM image files take up only as much space on the NFS filer as is required. If a 100-GB VDI is allocated for a new VM and an OS is installed, the VDI file will only reflect the size of the OS data that has been written to the disk.

VHD files may also be *chained*, allowing two VDIs to share common data. In cases where a NFS-based VM is cloned, the resulting VMs will share the common on-disk data at the time of cloning. Each will proceed to make its own changes in an isolated *copy-on-write* version of the VDI. This feature allows NFS-based VMs to be quickly cloned from templates, facilitating very fast provisioning and deployment of new VMs.

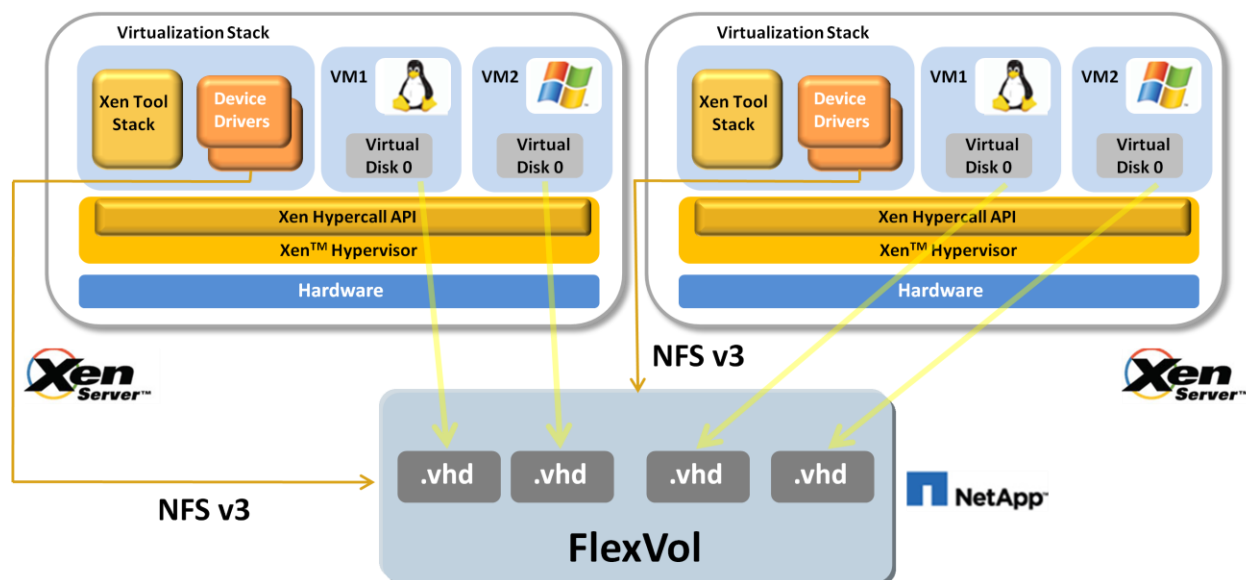


Figure 2 Shared NFS storage

Shared iSCSI Storage

XenServer provides support for shared SRs on iSCSI LUNs. iSCSI is supported using the open-iSCSI software iSCSI initiator or using a supported iSCSI *Host Bus Adapter* (HBA).

Shared iSCSI support is implemented based on a Logical Volume Manager (LVM). LVM-based storage is high-performance and allows virtual disks to be dynamically resized. Virtual disks are fully allocated as an isolated volume on the underlying physical disk and so there is a minimum of storage virtualization overhead imposed. As such, this is a good option for high-performance storage, but lacks the flexibility of file-based storage options. Additionally, iSCSI VDIs do *not* provide support for sparse provisioning or fast cloning from the XenServer API level.

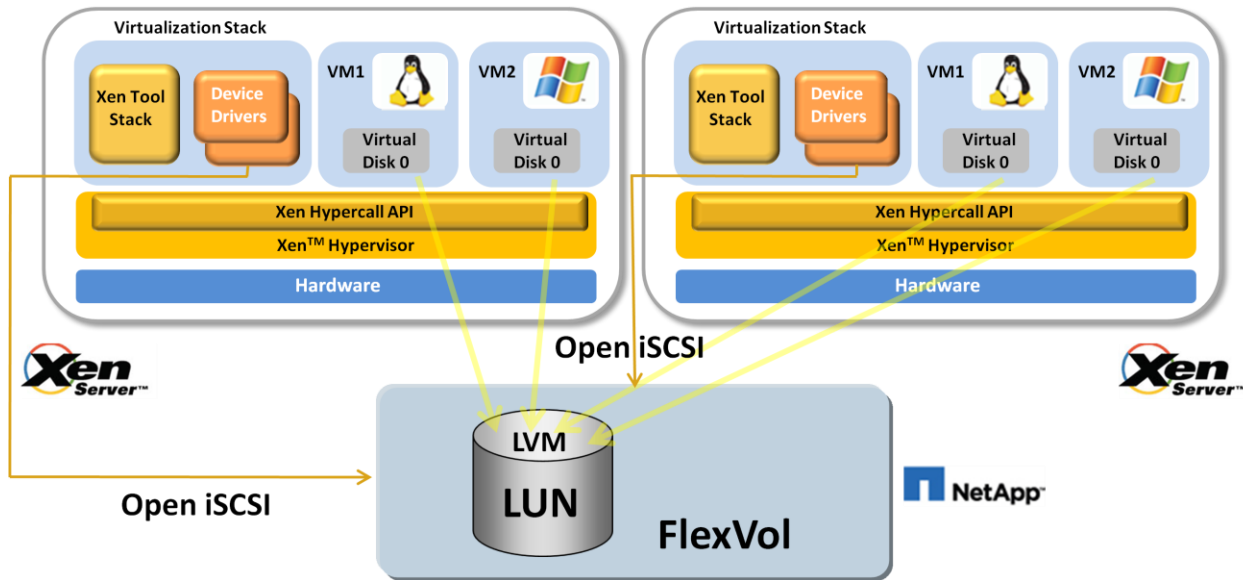


Figure 3 Shared iSCSI storage using Open iSCSI initiator

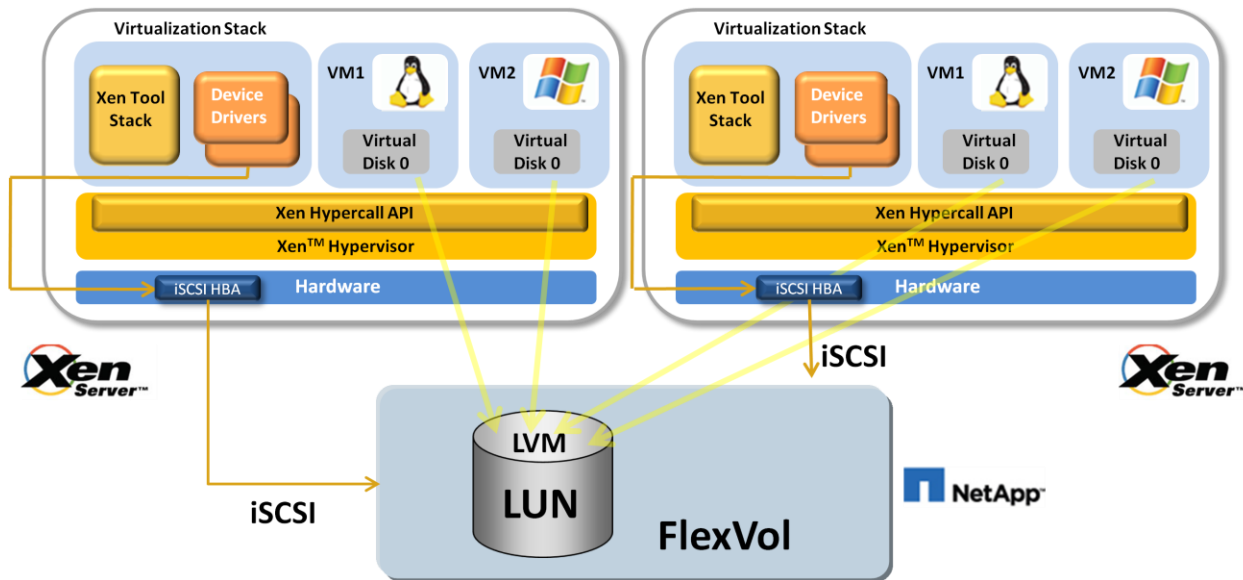


Figure 4 Shared iSCSI storage using iSCSI HBA

Shared Fibre Channel Storage

XenServer hosts supports Fibre Channel SANs using the Emulex or QLogic *host bus adapter* (HBA). Logical unit numbers (LUNs) are mapped to the XenServer host as disk devices.

Like iSCSI storage, Fibre Channel storage support is implemented based on the same Logical Volume Manager with the same benefits as iSCSI storage, just utilizing a different data IO path.

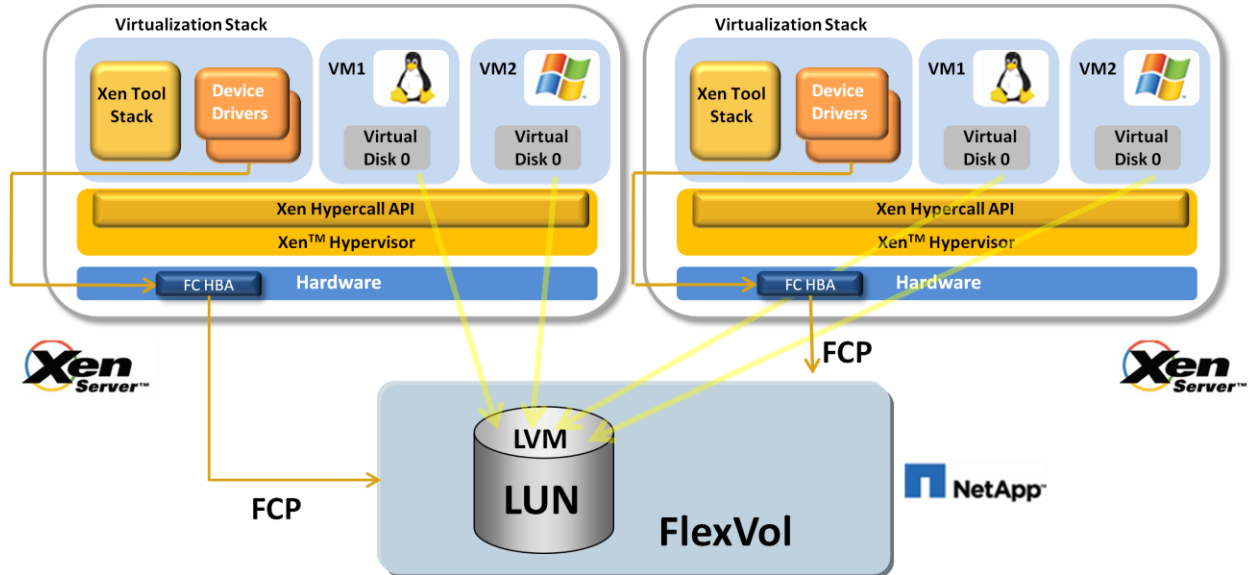


Figure 5 Shared FC storage

CONFIGURATION AND SETUP

NETAPP FILER CONFIGURATION BASICS

This section covers the best practices for configuring the NetApp storage box.

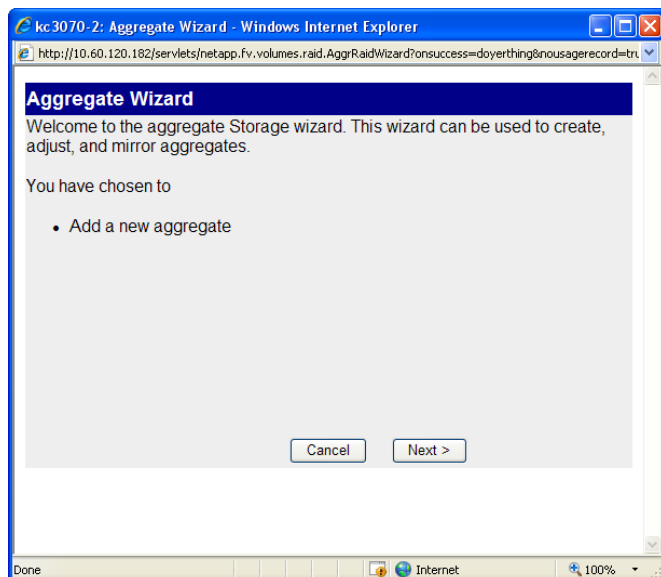
Configuring an Aggregate

It is recommended that a separate aggregate be used for XenServer hosts in a data center.

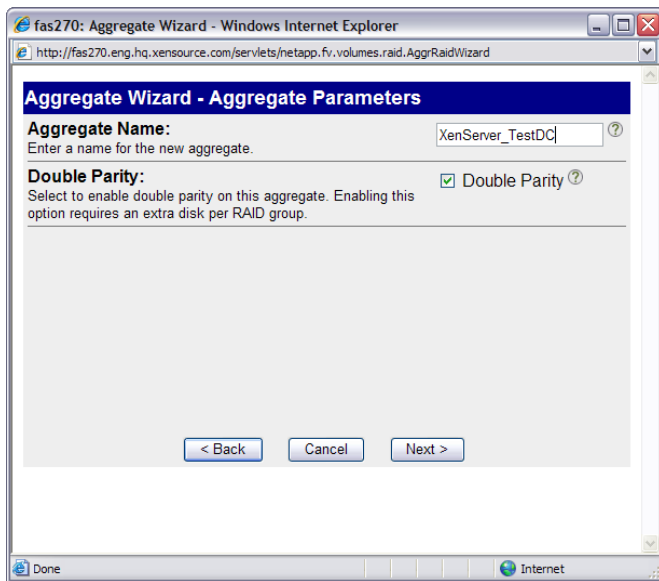
If FAS De-Dupe is going to be enabled for the FlexVol® from XenCenter, then note that the aggregate size should match the maximum supported FAS De-Dupe limit for the filer. The current limits are listed in the table below:

NetApp Filer Model	Aggregate Limit
FAS6070 / N7800	16 TB
FAS6030 / N7600	10 TB
FAS3070 / N5600	6 TB
NearStore R200	4 TB
FAS3040 / N5300	3 TB
FAS3050 / N5500	2 TB
FAS3020 / N5200	1 TB
FAS2050	1 TB
FAS2020	.5 TB

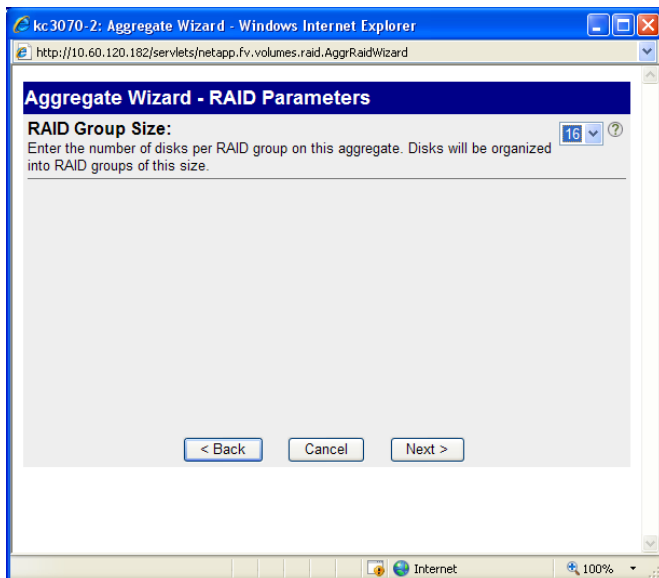
1. Open the NetApp FilerView®, and click *Aggregates->Add* to add a new aggregate on the filer.



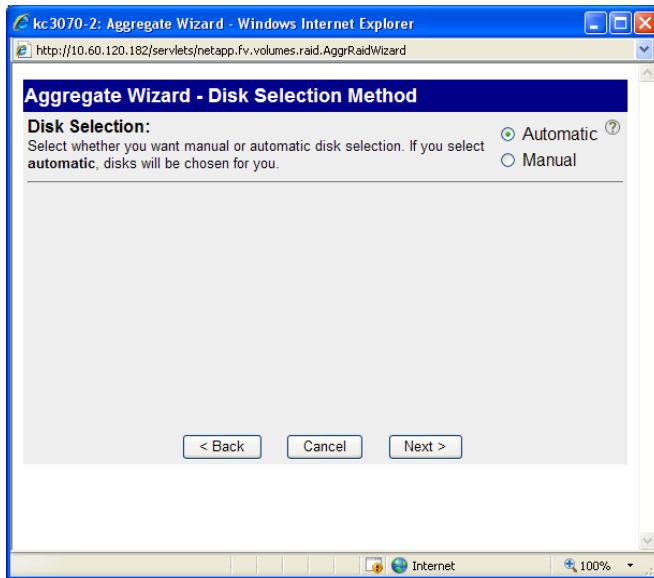
- Choose an aggregate name that reflects the data center that will be using it, along with "XenServer" in the name. This will also make storage configuration easier when accessed via the XenCenter management interface. Choose the *Double Parity* option if there is an extra disk per RAID group available. This is the recommended RAID-level for the aggregate.




- Select the number of disks to assign for the *Raid Group Size*. By default the filer will default to 16.



4. Select the disk that will be used for the aggregate. *Automatic* is selected by default for this section.



5. Choose the disk size to be used in the aggregate. By default *Any Size* will be selected.



6. Choose the number of disks to be added to the aggregate.



7. Assign at least three disks in order to provision an aggregate. Click *Next* and then *Commit* to create the aggregate.

Thin Provisioning vs. Thick Provisioning

Thin provisioning allows the admin to present more storage space to the VMs connecting to the SR than is actually available on the SR. It is important to note that there are no space guarantees, and allocation of a LUN does not claim any data blocks in the FlexVol until the VM writes data into it.

The following scenarios benefit from thin provisioning

- a. Quick provisioning
- b. Not realistic to decide how much storage is required up front
- c. Allow storage capacity additions without downtime

In comparison, thick provisioning ensures that whenever a LUN is allocated on the filer, sufficient space is reserved to guarantee that it will never run out of space and consequently experience failed writes to disk. Due to the nature of the ONTAP FlexVol space provisioning algorithms, the best practice guidelines for the filer require that at least twice the LUN space is reserved to account for background Snapshot data collection and to ensure that writes to disk are never blocked. In addition to the double disk space guarantee, ONTAP also requires some additional space reservation for management of unique blocks across snapshots. The guideline on this amount is 20% above the reserved space. Therefore, the space guarantees afforded by "thick provisioning" will reserve up to 2.4 times the requested virtual disk space.

Using FAS De-Dupe

FAS De-Dupe reclaims redundant disk space by dividing newly-stored data objects into small blocks, each block containing a digital signature, which is compared to all other signatures in the data volume. If an exact block match exists, the duplicate block is discarded and the disk space reclaimed. The default De-Dupe parameters typically run the algorithm every 24 hours.

Below are some results on the testing FAS De-Dupe¹

- Supported by all storage data access types; iSCSI, FCP & NFS
- 50% – 70% typical storage reduction
- >90% with Virtual Desktop Infrastructure
- Saves up to 95% for full backups; 25% to 55% for most data sets.

¹ The testing was done in NetApp's Kilo Client Center as part of a Virtual Desktop Infrastructure Proof of Concept. The backend filer used was a FAS 3050 clustered, and XenServer running on IBM Blade Center H.

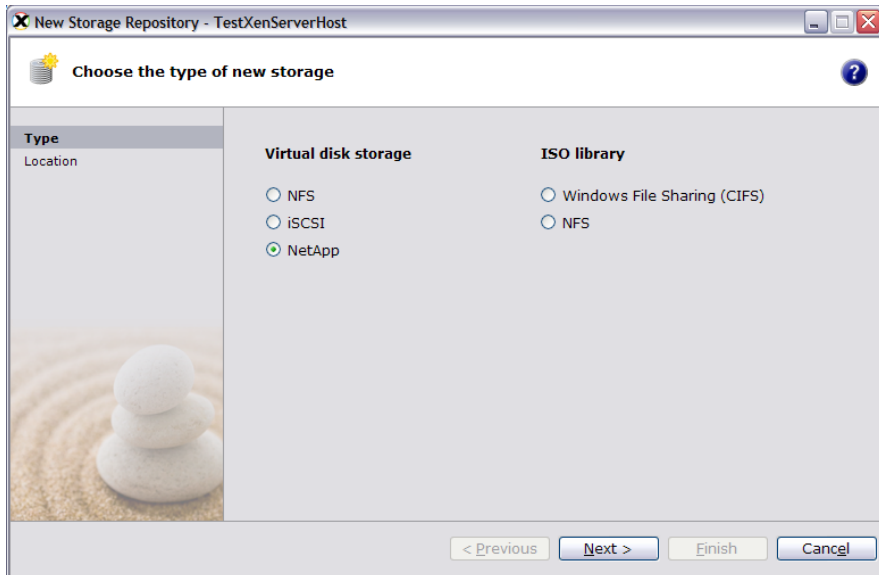
XENSERVER SHARED STORAGE CONFIGURATION BASICS

This section covers the best practices for configuring the various available storage options (Data ONTAP, NFS, iSCSI, FP) with a NetApp storage box. The recommended storage configuration is to utilise the NetApp Data ONTAP SR type since it leverages NetApp Data ONTAP capabilities for better performance and storage management.

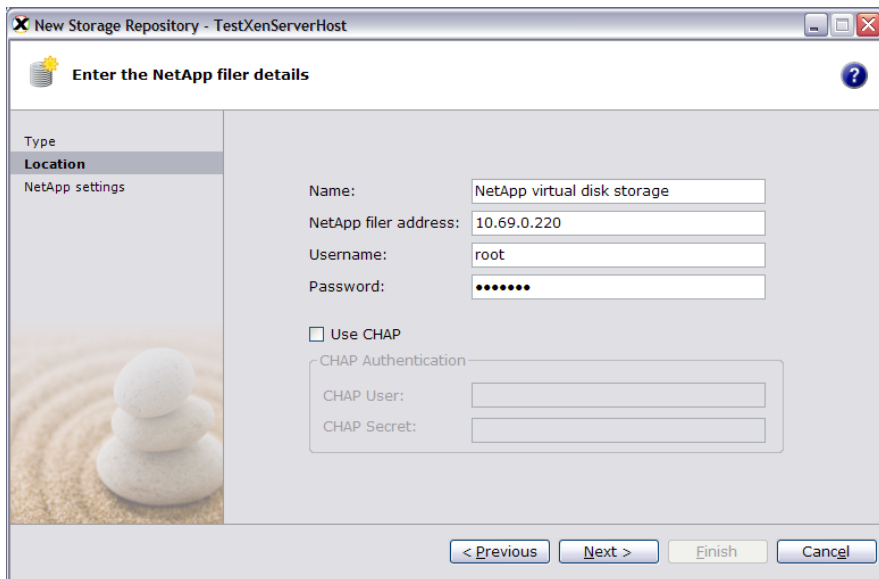
Shared Storage using NetApp Data ONTAP

This type of storage repository uses NetApp Data ONTAP for its control path, with in-band configuration and management of the data path via the host software initiator. The only backend NetApp storage configuration required is to create an aggregate that will house the FlexVols used by this storage repository type. See Appendix B for information for security considerations to allow XenServer administrators to have root access to the filer to be able to provision FlexVols and generate Snapshots on the filer.

1. For using NetApp Data ONTAP to create the storage repository, in XenCenter, choose *New Storage*. Select *NetApp*.



2. Provide the name of the filer (or its IP address), and authentication credentials. If CHAP is required, then select *Use CHAP* and provide the username and password. CHAP is only required if there is a security requirement between the NetApp storage and the XenServer.



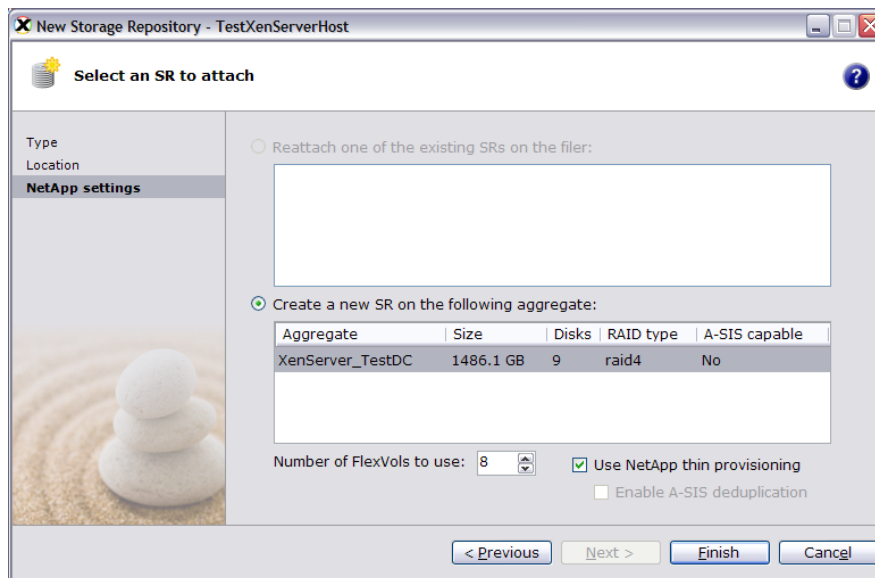
Note, the filer manages CHAP authentication credentials per host IQN. This must be managed and configured directly on the filer if required.

3. The NetApp Data ONTAP adapter will poll the filer for existing aggregates. Choose the aggregate to create the storage repository in. If thin provisioning is required, then check the box entitled *Use NetApp thin provisioning*. With thin provisioning, it is also possible to make enable FAS de-dupe on the backend.

Thin provisioning is a very useful space conserving mechanism for Virtual Machine disk storage, since many Virtual Machines are likely to significantly under-utilize all the virtual disk space allocated. Furthermore, FAS de-dupe can be very effective where many VMs contain the same, or similar, Operating System, and there is likely to be a significant amount of duplication of data across disks. Selecting *Use NetApp thin provisioning* with *Enable A-SIS deduplication* can significantly reduce the amount of space required on disk. Note, however that there are no space guarantees when operating in thin provisioning mode, so it is quite feasible to over-provision the amount of allocated space. If an over-provisioned aggregate runs out of space, a LUN will be forcibly set offline when any data writes are received which may cause a Virtual Machine to crash. Management of space usage can be greatly improved by utilizing the NetApp Data ONTAP alert mechanisms. Alternatively, Virtual Machine disk space can be guaranteed by disabling thin provisioning.

The *Number of FlexVols to use* is by default 8. This can be changed to a number in the range 2 - 32. The number of FlexVols to be used is based on:

- a. Increase the number of FlexVols if the data center will have:
 - i. A lot of VMs that will be Snapshot often
 - ii. A lot of VMs with multiple VDIs
- b. Decrease the number of FlexVols if the data center will have:
 - i. SnapMirror® replication to a Disaster Recovery Site




4. Click *Finish* for the storage repository to be created using Data ONTAP for the control path.

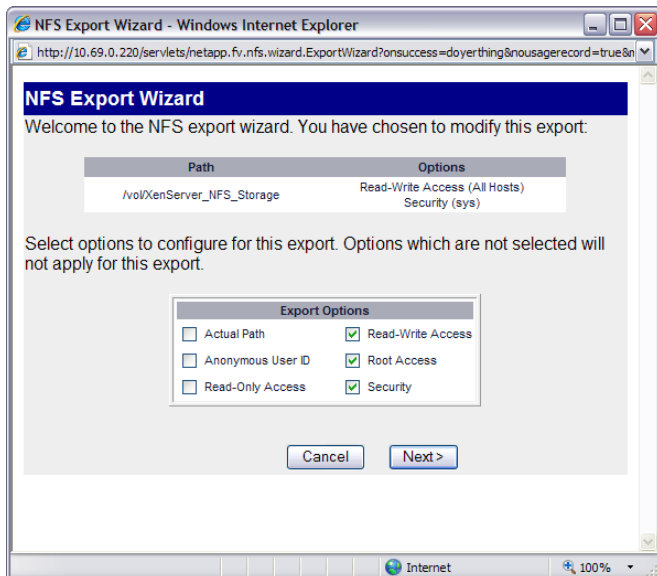
Configuration Shared NAS using NFS

To use the NetApp storage box as a shared NAS storage option using NFS, it is recommended that a separate volume be created for VDI storage. To do so:

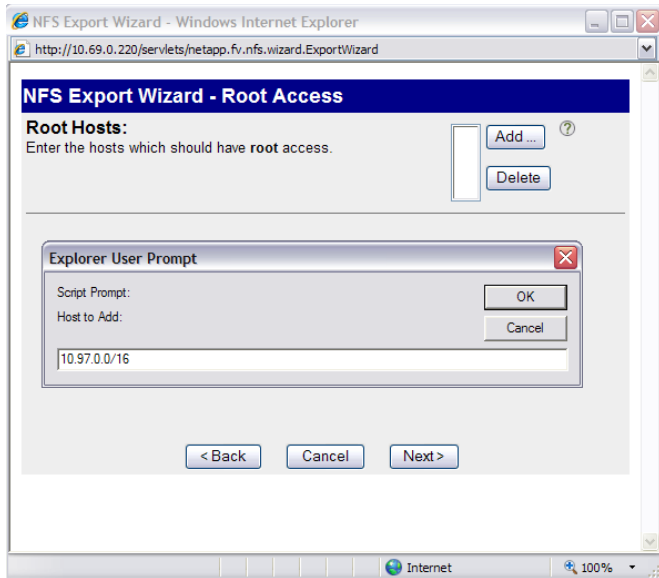
1. Open the FilerView for the storage box by pointing a browser to http://filer/na_admin
2. Click on *Volumes*, and then *Add* to open the *Volume Wizard*
3. Click *Next* and select the volume type as *Flexible*.



4. It is recommended to give a name that the NetApp storage server's automatic support system can identify as specific to XenServer storage, for example, a name such as "XenServer_NFS_Storage" would be appropriate.
5. To create an NFS export on the above volume, select *NFS->Manage Exports*. Click on the above created volume to modify the options associated with the export. For the XenServer host to be able to create SRs on the exported NFS target, the host's IP address or subnet mask needs to be granted *Root Access*. Select *Root Access* from the export wizard.




6. For the *Root Hosts* page, click *Add...* If all hosts on a particular subnet should be given access to the NFS storage repository, then enter the subnet. Else, enter the individual host names (or IP addresses) separated by a comma. In the case of subnet, the format is *dotted_ip/num_bits*, where *dotted_ip* is the IP subnet and *num_bits* is the number of leading 1 bits on the subnet mask. For the example in figure below, the XenServer host is in the 10.97.0.0 subnet – so the entry is 10.97.0.0/16.



In the following wizard panes, set the *Security* to *Unix Style*, and *Commit* the changes in the end.

7. In XenCenter, connect to the XenServer host, and choose *New Storage*. Select *NFS*.



8. Give the NFS share a name, and set the path to the NFS export point from the filer and click *Finish*.

New Storage Repository - XenServerTest

Enter a name and path for the new NFS storage

Type

Location

Name:

Share Name:

Example: server:/path

Create a new SR

Reattach an existing SR:

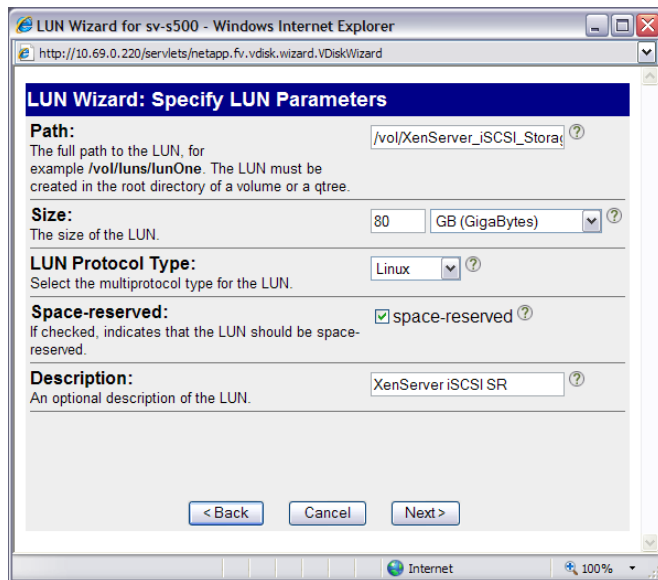
Configuring iSCSI Storage

To set up the NetApp storage box to be used for an iSCSI SR from XenServer, follow steps 1-4 from the previous section to create a FlexVol. In step2, choose the *Space Guarantee* as *volume*. In step 4, give the FlexVol a name that can easily be identified by the filer's autosupport program as a volume being used by a XenServer host, for example `XenServer_iSCSI-Storage` used in illustrations below. Keep the *Snapshot Reserve* to 20% if Snapshot will be done from the backend – XenServer does not support VDI snapshots for VDIs in an iSCSI SR. Otherwise, bring it down to 0% if there are concerns for space usage.

Once the FlexVol has been created, initiator groups (igroups) need to be created. These igroups are used to provide a mechanism of host access to a LUN via LUN masking. It is recommended that a single igroup be used for a data center. To simplify the management of igroups, we suggest that the igroup name contain “XenServer”, the protocol type (*iSCSI* in this case), and the data center name. This also helps the XenServer administrator to choose the correct igroup LUNs for the pool of hosts in the data center. The igroup itself contains iSCSI qualified names or IQNs that match the IQN of each XenServer host in the pool.

The LUN wizard from the NetApp FilerView can be used to create LUN that will be used by the iSCSI SR.

1. Select *Wizards*, and then *LUN Wizard* from the FilerView
2. Enter the path for the LUN for the above created volume as `/vol/XenServer_iSCSI-Storage/{LUN Name}`. Set the *Protocol Type* to *Linux*. It is recommended to use *Space-reserved* LUNs so that blocks in the LUN can be updated when they are written to. Click *Next*>



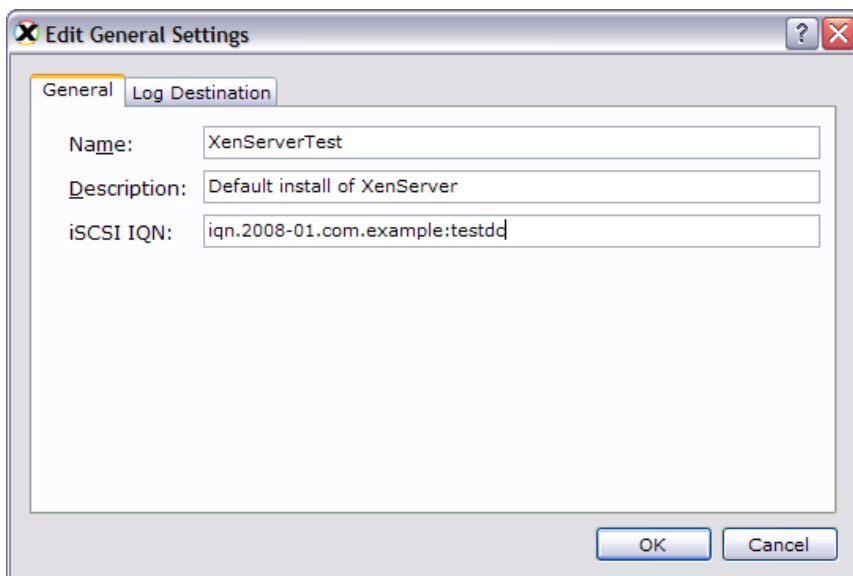

3. To enable LUN masking, the LUN must be mapped to the igroup. Click *Add Group* to either create a new igroup or assign an already existing igroup. Select *Create new initiator group* and continue.




4. Give a name that uniquely identifies the igroup to be used by XenServer hosts in the data center. Select the initiator group type as *iSCSI*, and the operating system as *Linux*.



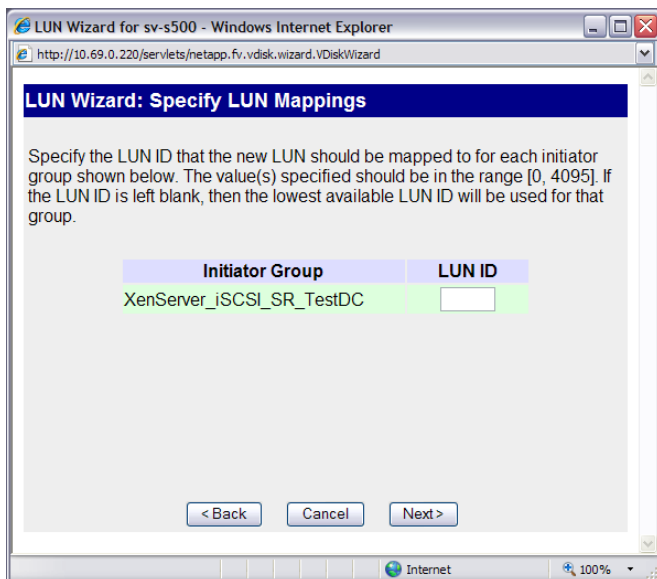
5. The next steps need IQNs from the XenServer host. The IQN of a XenServer host can be seen from the *General* tab for the host in XenCenter (in the figure below, the IQN for XenServerTest host is "iqn.2008-01.com.example:44b80d2a"). To change the IQN to a more recognizable name, click on the *Edit* button at the bottom right of the *General* tab, and modify the *iSCSI IQN* field (see figure below where IQN is changed to iqn.2008-01.com.example:testdc).



6. Once you have set the appropriate IQN on the XenServer host, it needs to be added to the igroup in the LUN wizard from the NetApp FilerView. Click *Add Initiator*, and enter the IQN from the XenServer host here.

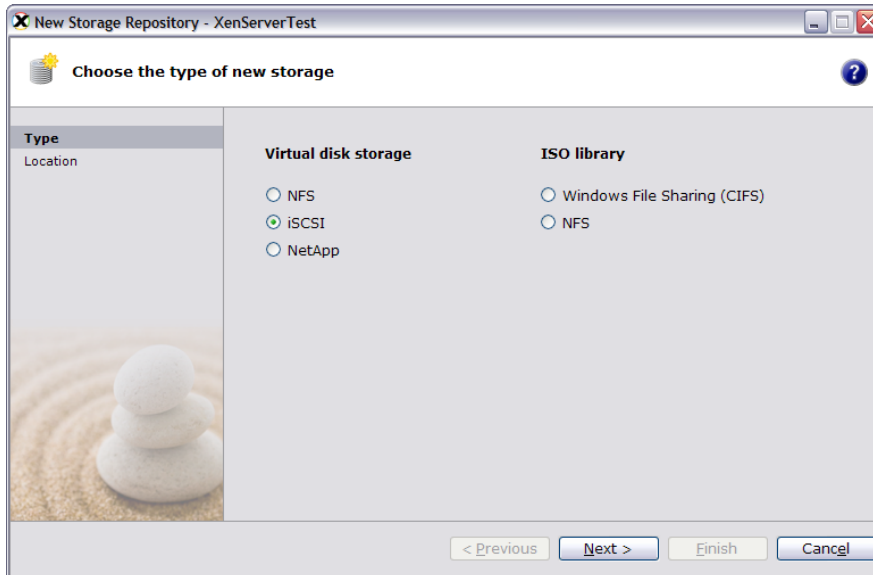


7. Continuing on, the LUN Wizard will ask for a LUN ID that the newly created LUN should be mapped to in the igroup. This can be left blank so that lowest available LUN ID used.

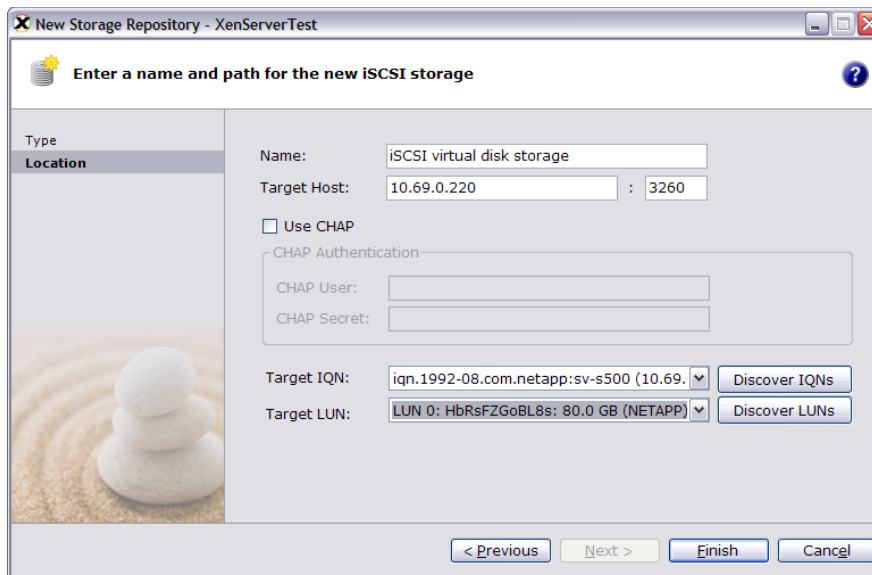


8. Commit the changes from the wizard, and switch to the XenCenter GUI to start creating an iSCSI based storage repository on the newly created LUN.


9. For using Open iSCSI software initiator to create the storage repository, in XenCenter, choose *New Storage*. Select *iSCSI*.



- a. Enter the *Target Host* as the hostname or IP address of the NetApp filer in which the LUN was set up in steps 1-8. If *CHAP Authentication* was set up for iSCSI security on the filer, then enter the *CHAP User* and *CHAP Secret*. It is recommended to have the same CHAP username/password for initiators in the same igroup (as is the case with a pool of XenServer hosts connecting to the same igroup). Most of the time, customers typically do not enable security for iSCSI unless there is a security requirement. If you are required to enable security for the iSCSI connection, then we recommend that you utilize the CHAP option. Click on *Discover IQNs* to get the list of *Target IQNs* on the NetApp filer. Select the relevant *Target IQN* and click on *Discover LUNs* to get a list of LUNs associated with the mapped igroup and LUNs in it. From the initial steps, the LUN that was created should show up in this list. Select the LUN, and click *Finish*.



- b. The new LUN will be overlaid with LVM, and XenCenter will ask the LUN to be formatted as such. Click *Yes* on the pop-up for the LUN to be formatted with LVM.



10. For using an iSCSI HBA to create the iSCSI SR, the CLI from the control domain needs to be used. Depending on what HBA is being used, the initiator IQN for the HBA needs to be configured. Given the type of HBA used, the documentation for that HBA should be consulted to configure the IQN.
- Once the IQN has been configured for the HBA, use the NetApp FilerView to create a new LUN as in steps 1-4. Instead of using the XenServer's IQN, specify the IQN of port 0 of the HBA. Because XenServer host does not support dual multi-path, this guide uses port 0 only. Two HBA CLI's are included in the XenServer host to configure the HBA:

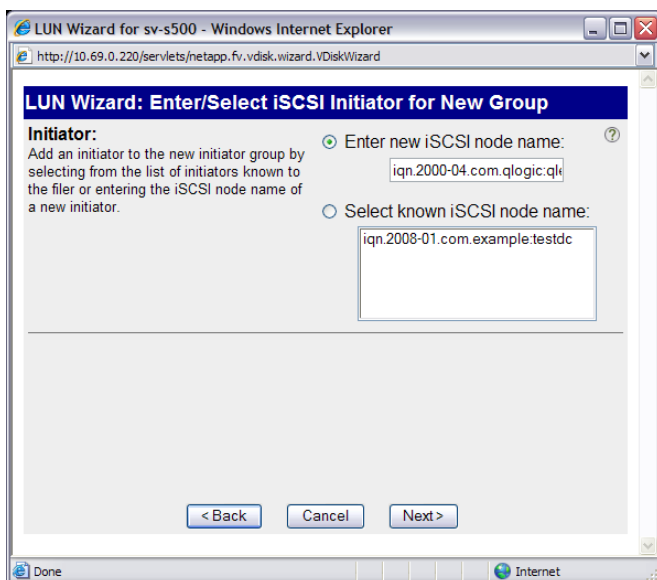
```
Emulex: /usr/sbin/hbaanyware
QLogic iSCSI: /opt/QLogic_Corporation/SANsurferiCLI/iscli
```

For the purposes of an example, this guide illustrates how the QLogic iSCSI HBA CLI *iscli* can be used. The IQN of port 0 of the HBA can be read from the output of *iscli* command.

```
C:\WINDOWS\system32\cmd.exe - ssh root@10.70.0.41
main Interactive Menu

-----
Program Version: 1.1.00.13 Driver Version: 5.01.00.08 IC: 2
FW Version: 3.0.1.24 Type: Copper
Current HBA/Port Information: HBA Alias:
HBA: 0 Port: 0 HBA Port Index: 1 HBA Model: QLE4062C
IP Address: 0.0.0.0 Link: Up
IPv6 Protocol is currently disabled.
Port iSCSI Name: iqn.2000-04.com.qlogic:qle4062c.gs40743a48539.1
Port iSCSI Alias:
-----

1. Display Program Version Information
2. Host Level Info & Operations
3. HBA Level Info & Operations
4. Port Level Info & Operations
5. List All QLogic iSCSI HBA Ports detected
6. Help
7. Select HBA Port
8. Refresh
9. Exit
enter selection: 5
1. HBA: 0 Port: 0 HBA Port Index: 1 HBA Model: QLE4062C
   HBA Serial Number: <GS40743A48539> FW Version: 3.0.1.24 Type: Copper
   IP Address: 0.0.0.0
   Alias:
   iSCSI Name: iqn.2000-04.com.qlogic:qle4062c.gs40743a48539.1
2. HBA: 0 Port: 1 HBA Port Index: 2 HBA Model: QLE4062C
   HBA Serial Number: <GS40743A48539> FW Version: 3.0.1.24 Type: Copper
   IP Address: 0.0.0.0
   Alias:
   iSCSI Name: iqn.2000-04.com.qlogic:qle4062c.gs40743a48539.2
Press the Enter key to continue.
```



- After creating the LUN, set the IP address for the HBA. In the control domain, use the *iscli* CLI to do so.

```
/opt/QLogic_Corporation/SANsurferiCLI/iscli -ipdhcp 0
```

```
C:\WINDOWS\system32\cmd.exe - ssh root@10.70.0.41
[root@TestXenServerHost ~]# /opt/QLogic_Corporation/SANsurferiCLI/iscli -ipdhcp 0
inst 0 Waiting for port to be ready
[root@TestXenServerHost ~]#
```

- c. Now add a persistent target to the HBA. The target iSCSI IQN can be retrieved from the NetApp FilerView by clicking *LUNs->Initiator Groups->iSCSI->Manage Names*.



```
/opt/QLogic_Corporation/SANsurferiCLI/iscli -pa 0 <iSCSI filer IP Address> -INAME
<target iSCSI IQN>
```

```
C:\WINDOWS\system32\cmd.exe - ssh root@10.70.0.41
[root@TestXenServerHost ~]# /opt/QLogic_Corporation/SANsurferiCLI/iscli -pa 0 10.69.0.220 -INAME iqn.1992-08.com.netapp:sv-s500
[root@TestXenServerHost ~]#
```

This will assign port 0 the specified iSCSI IP address and specific target IQN

- d. Use the *xe sr-probe* command to force a scan of HBAs installed on the system and detect the new LUN zoned to the host. It will return the list of properties for each LUN found. One of these properties will be *<path>* which is the global device path of the HBA LUN. Specify the host-uuid of the system from where the *xe sr-probe* command is run.

```
xe sr-probe type=lvmohba host-uuid=<UUID of host>
```

```

C:\WINDOWS\system32\cmd.exe - ssh root@10.70.0.41
[root@TestXenServerHost ~]# xe sr-probe type=lvmohba host-uuid=0e8ff72c-d00d-47df-860b-0aac1d05b675
Error code: SR_BACKEND_FAILURE_90
Error parameters: , The request is missing the device parameter. <?xml version="1.0" ?>
<Devlist>
  <BlockDevice>
    <path>
      /dev/disk/by-id/scsi-360a98000686f64616b4a483466363972
    </path>
    <vendor>
      NETAPP
    </vendor>
    <serial>
      hodakJH4f69r
    </serial>
    <size>
      5368709120
    </size>
    <adapter>
      2
    </adapter>
    <channel>
      0
    </channel>
    <id>
      0
    </id>
    <lun>
      0
    </lun>
    <hba>
      qla2xxx
    </hba>
  </BlockDevice>
</Devlist>

```

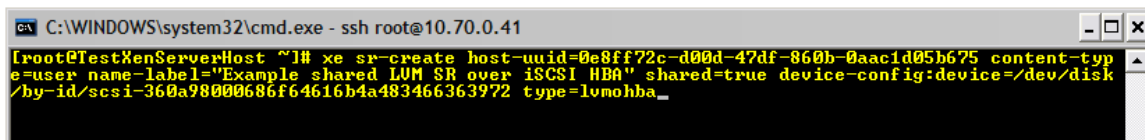
To validate that the device path is for the newly created LUN on the filer, match the serial number from the `<serial>` field of the `xe sr-probe` output with the serial number of the LUN in FilerView. To determine the LUN serial number from the NetApp FilerView, click *LUNs->Manage* and click on the newly created LUN. Then note the *Serial Number* field.

The screenshot shows the 'Modify LUN' page in the NetApp FilerView web interface. The 'Serial Number' field is highlighted with a red box and contains the value 'HbRsFZH4JfS'. The interface includes a navigation menu on the left, a breadcrumb trail 'LUNs -> Manage -> Modify', and various configuration fields for the LUN.

The device referred by `/dev/disk/by-id/<new-device>` can be used to create the SR as in the next step.

- e. In the control domain, use the `xe sr-create` command to create the storage repository


```
xe sr-create host-uuid=<valid_UUID> content-type=user name-label="Example shared LVM  
over HBA SR" shared=true device-config:device=/dev/disk/by-id/<new-device>  
type=lvmohba
```

A terminal window titled 'C:\WINDOWS\system32\cmd.exe - ssh root@10.70.0.41' showing a command being executed. The command is: xe sr-create host-uuid=0e8ff72c-d00d-47df-860b-0aac1d05b675 content-type=user name-label="Example shared LVM SR over iSCSI HBA" shared=true device-config:device=/dev/disk/by-id/scsi-360a98000686f64616b4a483466363972 type=lvmohba_. The output is not visible.

```
C:\WINDOWS\system32\cmd.exe - ssh root@10.70.0.41  
[root@TestXenServerHost ~]# xe sr-create host-uuid=0e8ff72c-d00d-47df-860b-0aac1d05b675 content-type=user name-label="Example shared LVM SR over iSCSI HBA" shared=true device-config:device=/dev/disk/by-id/scsi-360a98000686f64616b4a483466363972 type=lvmohba_
```

The LUN is now used as an iSCSI SR, with XenServer overlaying it with LVM to manage VDIs as logical volumes.


Configuring Fibre Channel Storage

To set up the NetApp storage box to be used for a Fibre Channel (FC) SR from XenServer, follow steps 1-4 from the previous section to create a FlexVol. In step 2, choose the *Space Guarantee* as *volume*. In step 4, give the FlexVol a name that can easily be identified by the filer's autosupport program as a volume being used by a XenServer host. Keep the *Snapshot Reserve* to 20% if snapshot will be done from the backend – XenServer does not support VDI snapshots for VDIs in an FC SR. Otherwise, bring it down to 0% if there are concerns for space usage.

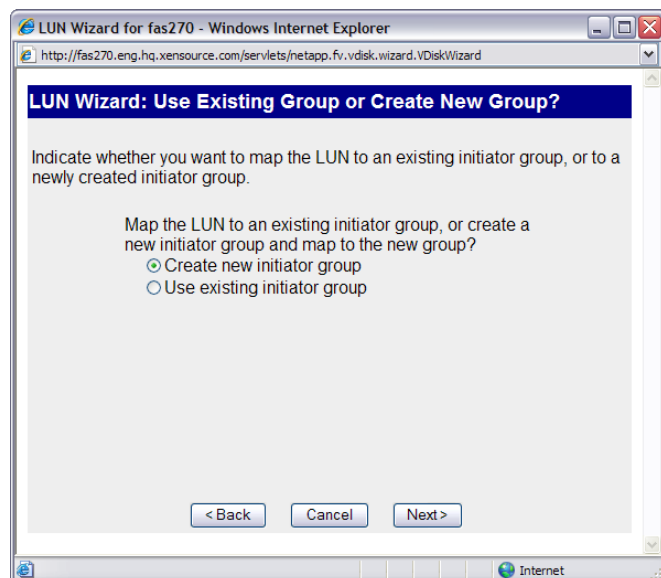
As with the iSCSI SR FlexVols, initiator groups (igroups) need to be created. To simplify the management of igroups, it is suggested that the igroup name contain “XenServer”, the protocol type (*FCP* in this case), and the data center name. This also helps the XenServer administrator to choose the correct igroup LUNs for the pool of hosts in the data center. The igroup itself contains World Wide Port Names (WWPNs) that match the WWPN for the HBA ports on the XenServer host.

The LUN wizard from the NetApp FilerView can be used to create LUN that will be used by the SR using FCP.

1. Select *Wizards*, and then *LUN Wizard* from the FilerView
2. Enter the path for the LUN for the above created volume as `/vol/XenServer_FCP_Storage/{LUN Name}`. Set the *Protocol Type* to Linux. It is recommended to use *Space-reserved* LUNs so that blocks in the LUN can be updated when they are written to. Click *Next*>



3. To enable LUN masking, an igroup needs to be assigned to the LUN being created. Click *Add Group* to either create a new igroup or assign an already existing igroup. Select *Create new initiator group* and continue.



4. Give a name that uniquely identifies the igroup to be used by XenServer hosts in the data center. Select the initiator group type as *FCP*, and the operating system as *Linux*.



5. The next steps need WWPN for the HBA from the XenServer host. Since different FC vendors have specific and different configuration requirements, it is recommended that the documentation for the specific HBA be consulted for configuration settings.

This guide will assume a QLogic 2342 HBA, and as such use the `/opt/QLogic_Corporation/SANsurferCLI/scli` to get configuration information. Run `/opt/QLogic_Corporation/SANsurferCLI/scli` in the control domain, and enter 5 from the main menu. The figure below shows the WWPN for Port 1 highlighted.

```

C:\WINDOWS\system32\cmd.exe - ssh root@10.70.0.41
[root@TestXenServerHost ~]# /opt/QLogic_Corporation/SANsurferCLI/scli
Searching for QLogic FC HBA(s) and attached device(s) ...

      SANsurfer FC HBA CLI
      v1.7.0 Build 14

      Main Menu
      1: Display System Information
      2: Display HBA Settings
      3: Display HBA Information
      4: Display Device List
      5: Display LUN List
      6: Configure HBA Settings
      7: Target Persistent Binding
      8: Selective LUNs
      9: Boot Device Settings
     10: Driver Settings
     11: HBA Utilities
     12: Flash Beacon
     13: Diagnostics
     14: Statistics
     15: Help
     16: Quit

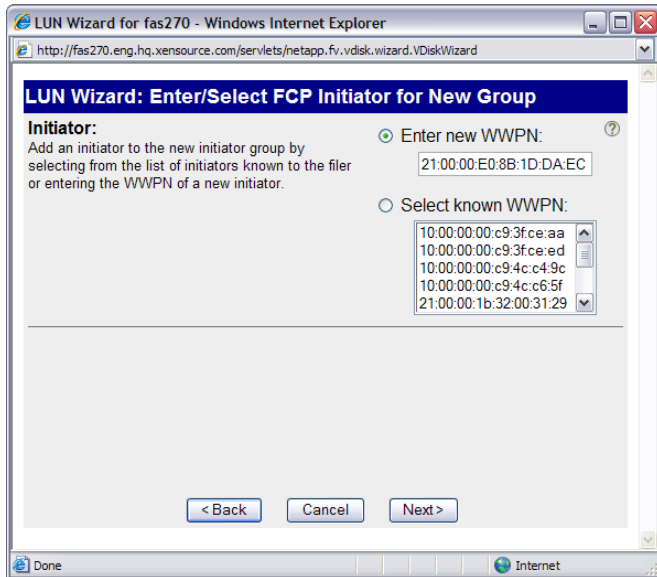
      Enter Selection: 5

      SANsurfer FC HBA CLI
      v1.7.0 Build 14

      LUN List Menu
      HBA Model QLA2342:
      1: Port 1: WWPN: 21-00-00-E0-8B-1D-DA-EC Online
      2: Port 2: WWPN: 21-01-00-E0-8B-3D-DA-EC Loop down
      3: All HBAs

      Note: 0 to return to Main Menu
      Enter Selection:
  
```

6. Once you have set the appropriate WWPN from the XenServer host, it needs to be added to the igroup in the LUN wizard from the NetApp FilerView. Click *Add Initiator*, and enter the WWPN for the FC HBA on the XenServer host here. Note that the WWPN in the wizard needs to be in `xx:xx:xx:xx:xx:xx:xx:xx` format where 'x' is a hexadecimal digit.



7. Continuing on, the LUN Wizard will ask for a LUN ID that the newly created LUN should be mapped to in the igroup. This can be left blank so that lowest available LUN ID used.
8. Commit the changes from the wizard
9. The newly created LUN now needs to be zoned in to the XenServer host and will appear as a SCSI device. For this, use the `xe sr-probe` command similar to the usage as when creating an iSCSI HBA SR.
- a. Use the `xe sr-probe` command to force a scan of HBAs installed on the system and detect the new LUN zoned to the host. It will return the list of properties for each LUN found. One of these properties will be `<path>` which is the global device path of the HBA LUN. Specify the host-uuid of the system from where the `xe sr-probe` command is run.

```
xe sr-probe type=<SR type> host-uuid=<UUID of host>
```

```
C:\WINDOWS\system32\cmd.exe - ssh root@10.70.0.41
[root@TestXenServerHost ~]# xe sr-probe type=lvmohba host-uuid=0e8ff72c-d00d-47df-860b-0aac1d05b675
Error code: SR_BACKEND_FAILURE_90
Error parameters: , The request is missing the device parameter, <?xml version="1.0" ?>
<Devlist>
  <BlockDevice>
    <path>
      /dev/disk/by-id/scsi-360a98000686f64616b4a483466363972
    </path>
    <vendor>
      NETAPP
    </vendor>
    <serial>
      hodakJH4f69r
    </serial>
    <size>
      5368709120
    </size>
    <adapter>
      2
    </adapter>
    <channel>
      0
    </channel>
    <id>
      0
    </id>
    <lun>
      0
    </lun>
    <hba>
      qla2xxx
    </hba>
  </BlockDevice>
</Devlist>
```

To validate that the device path is for the newly created zoned in LUN on the filer, match the serial number from the `<serial>` field of the `xe sr-probe` output with the serial number of the LUN in FilerView. To determine the LUN serial

number from the NetApp FilerView, click *LUNs->Manage* and click on the newly created LUN. Then look at the *Serial Number* field.

The screenshot shows the 'Modify LUN' page in the NetApp FilerView interface. The page is titled 'Modify LUN' and includes a navigation breadcrumb 'LUNs → Manage → Modify'. There are three buttons at the top: '[Manage LUNs]', '[Map LUN]', and '[Delete]'. The main configuration area includes the following fields:

- Path:** /vol/XenServer_FCP_Stora
- Status:** online
- LUN Protocol Type:** Linux
- Description:** (empty field)
- Size:** 5
- Units:** GB (GigaBytes)
- Space Reserved:** Space Reserved
- Serial Number:** hodakJH4f69r (highlighted with a red box)
- LUN Share:** none

An 'Apply' button is located at the bottom of the form.

The device referred by `/dev/disk/by-id/<new-device>` can be used to create the SR as in the next step.

- b. In the control domain, use the `xe sr-create` command to create the storage repository

```
xe sr-create host-uuid=<valid_UUID> content-type=user name-label="Example shared LVM over FC HBA SR" shared=true device-config:device=/dev/disk/by-id/<new-device> type=lvmohba
```

```
C:\WINDOWS\system32\cmd.exe - ssh root@10.70.0.41
[root@TestXenServerHost ~]# xe sr-create host-uuid=0e8ff72c-d00d-47df-860b-0aac1d05b675 content-type=user name-label="Example shared LVM over FC HBA SR" shared=true device-config:device=/dev/disk/by-id/scsi-360a98000686f64616b4a483466363972 type=lvmohba
8b8fd1da-2bca-6982-5bfe-ce839b62f040
[root@TestXenServerHost ~]#
```

The LUN is now used as an SR over FCP, with XenServer overlaying it with LVM to manage VDIs as logical volumes.

BACKUP AND RECOVERY

There are three elements for backup and recovery:

- Snapshot the virtual disks time to time,
- Backup the snapshots, and
- Recover data from snapshots when needed.

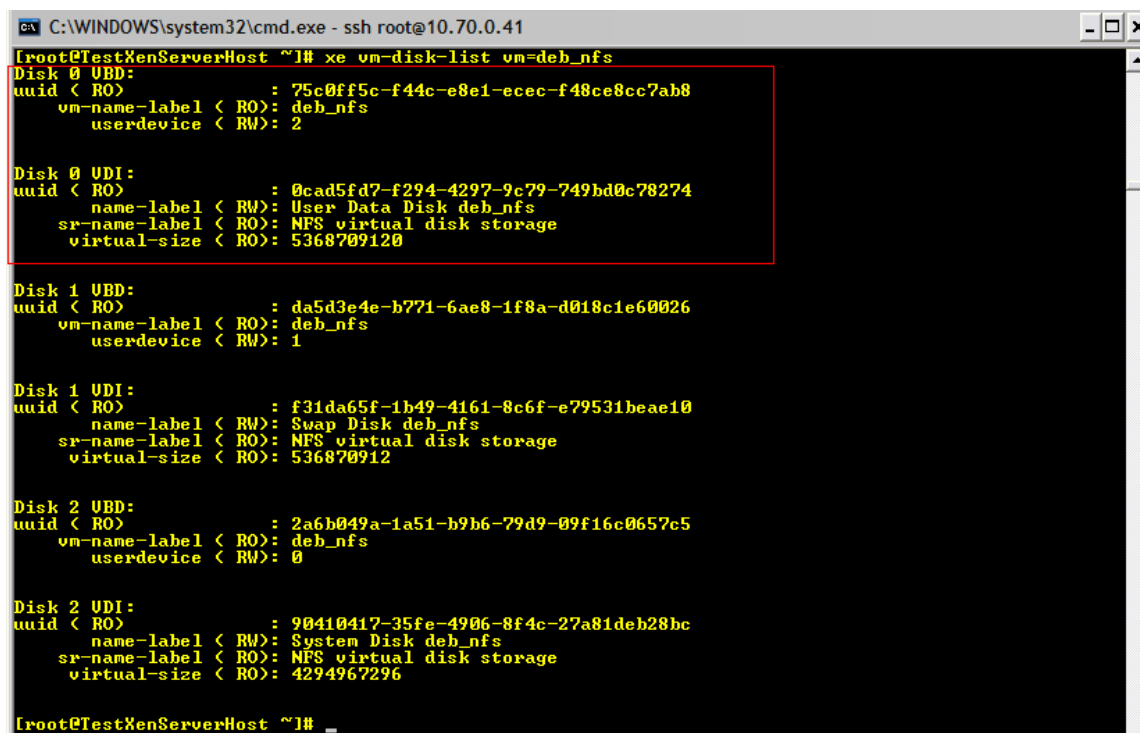
SNAPSHOT FOR VIRTUAL DISK ON NETAPP DATA ONTAP AND NFS STORAGE REPOSITORIES

Creating VDI snapshots for the NetApp Data ONTAP SR utilizes the NetApp filer Snapshot technology directly by invoking the Snapshot at the filer level. This results in minimal resource usage on the XenServer host in terms of CPU and memory during the snapshot process.

The snapshot process for VDIs on NFS SRs, however, does not invoke any NetApp filer capability. It uses the VHD capability to allow chaining for the original and snapshot VDI to share common data. The original VDI proceeds to make its own changes in an isolated copy-on-write version, with the snapshot VDI being Read Only.

1. In the control domain of the XenServer host, get the list of virtual disks for the virtual machine that you want to snapshot. `xe vm-disk-list` will return the attached Virtual Block Devices (VBDs) to the virtual machine, and list the VDI associated with the VBD. To learn more about VBD and VDI relationship, please see the XenServer Software Development Kit Guide for an overview of the XenServer API model.

```
xe vm-disk-list vm=<Name of Virtual Machine>
```



```
C:\WINDOWS\system32\cmd.exe - ssh root@10.70.0.41
[root@TestXenServerHost ~]# xe vm-disk-list vm=deb_nfs
Disk 0 UBD:
uuid < RO> : 75c0ff5c-f44c-e8e1-ecce-f48ce8cc7ab8
vm-name-label < RO>: deb_nfs
userdevice < RW>: 2

Disk 0 VDI:
uuid < RO> : 0cad5fd7-f294-4297-9c79-749bd0c78274
name-label < RW>: User Data Disk deb_nfs
sr-name-label < RO>: NFS virtual disk storage
virtual-size < RO>: 5368709120

Disk 1 UBD:
uuid < RO> : da5d3e4e-b771-6ae8-1f8a-d018c1e60026
vm-name-label < RO>: deb_nfs
userdevice < RW>: 1

Disk 1 VDI:
uuid < RO> : f31da65f-1b49-4161-8c6f-e79531beae10
name-label < RW>: Swap Disk deb_nfs
sr-name-label < RO>: NFS virtual disk storage
virtual-size < RO>: 536870912

Disk 2 UBD:
uuid < RO> : 2a6b049a-1a51-b9b6-79d9-09f16c0657c5
vm-name-label < RO>: deb_nfs
userdevice < RW>: 0

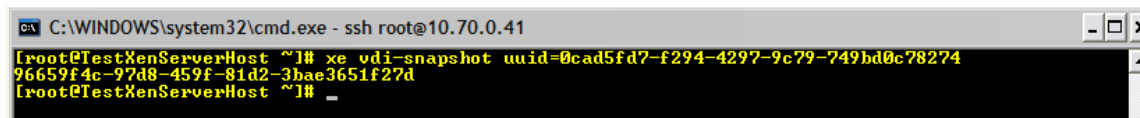
Disk 2 VDI:
uuid < RO> : 90410417-35fe-4906-8f4c-27a81deb28bc
name-label < RW>: System Disk deb_nfs
sr-name-label < RO>: NFS virtual disk storage
virtual-size < RO>: 4294967296

[root@TestXenServerHost ~]# _
```

The example above shows the virtual machine named “deb_nfs” has 3 virtual disks – “User Data Disk deb_nfs”, “Swap Disk deb_nfs” and “System Disk deb_nfs”.

2. To create a snapshot of “User Data Disk deb_nfs”, run the `xe vdi-snapshot` command. The command will create a snapshot of the disk and return a read-only VDI which is the snapshot disk.

```
xe vdi-snapshot uuid=<UUID of VDI to be snapshot>
```



```
C:\WINDOWS\system32\cmd.exe - ssh root@10.70.0.41
[root@TestXenServerHost ~]# xe vdi-snapshot uuid=0cad5fd7-f294-4297-9c79-749bd0c78274
96659f4c-97d8-459f-81d2-3bae3651f27d
[root@TestXenServerHost ~]# _
```

In the above example, “User Data Disk deb_nfs” is snapshot by passing its UUID to the `xe vdi-snapshot` command. The returned UUID is for the newly created VDI snapshot disk.

The CLI to snapshot a VDI on a NetApp Data ONTAP SR takes in additional arguments. As noted earlier, the NetApp Data ONTAP SR comprises of a collection of FlexVols which in turn contain the LUNs that represent each individual VDI. Thus, cloning a VDI in this scenario entails generating a Snapshot of the FlexVol and then creating a LUN clone backed off the Snapshot.

When generating a VM snapshot, an administrator must snapshot each of the VMs disks in sequence. Since all the disks are expected to be located in the same FlexVol, and the FlexVol Snapshot operates on all LUNs in the same FlexVol, it makes sense to re-use an existing Snapshot for all subsequent LUN clones. By default, if no snapshot hint is passed into the NetApp Data ONTAP adapter, it will generate a random ID with which to name the FlexVol Snapshot. There is a CLI override however for this value, passed in as an *epochhint*. The first time the epochhint value or 'cookie' is received, the backend will generate a new Snapshot based on the cookie name. Any subsequent snapshot requests with the same *epochhint* value will be backed off the existing Snapshot:

```
xe vdi-snapshot uuid=<VALID_VDI_UUID> driver-params:epochhint=<COOKIE>
```


3. Rename and give a description for the newly created and returned VDI snapshot disk to a meaningful name. This will help identify it a snapshot disk later on. The renaming can be achieved by using the `xe vdi-param-set` command.

```
xe vdi-param-set name-label=<New name label> name-description=<New description> uuid=<UUID of VDI>
```

```
C:\WINDOWS\system32\cmd.exe - ssh root@10.70.0.41
[root@TestXenServerHost ~]# xe vdi-param-set name-label="User Data Disk deb_nfs snapshot 20081231.235959.0001" name-description="Snapshot of user data disk for deb_nfs 001" uuid=96659f4c-97d8-459f-81d2-3bae3651f27d
[root@TestXenServerHost ~]# _
```

Notice that the snapshot disk has been renamed with a time reference that will help indicate when the snapshot was taken.

4. To verify that the snapshot disk exists in the same storage repository as the original disk, in XenCenter, click on the XenServer host. Then click on the NFS storage repository in which the VDI for the virtual machine exists, and click the *Storage* tab. There should be the original disk, as well as the snapshot disk.



BACKING UP STORAGE REPOSITORIES ON NETAPP FILER

Utilize NetApp's SnapMirror® technology to backup the FlexVols that make up the SR. For more information on use of SnapMirror, please refer to the SnapMirror Administration guide at <http://now.netapp.com/NOW/knowledge/docs/ontap/rel706/html/ontap/onlinebk/mirror.htm>.

The steps to restore a FlexVol that has been backed up using SnapMirror are:

1. Quiese the SnapMirror relationship.
2. Break the SnapMirror relationship.
3. Once the SnapMirror relationship has been broken, the mirrored FlexVol is automatically used as the primary SR FlexVol

To restore the SnapMirror relationship:

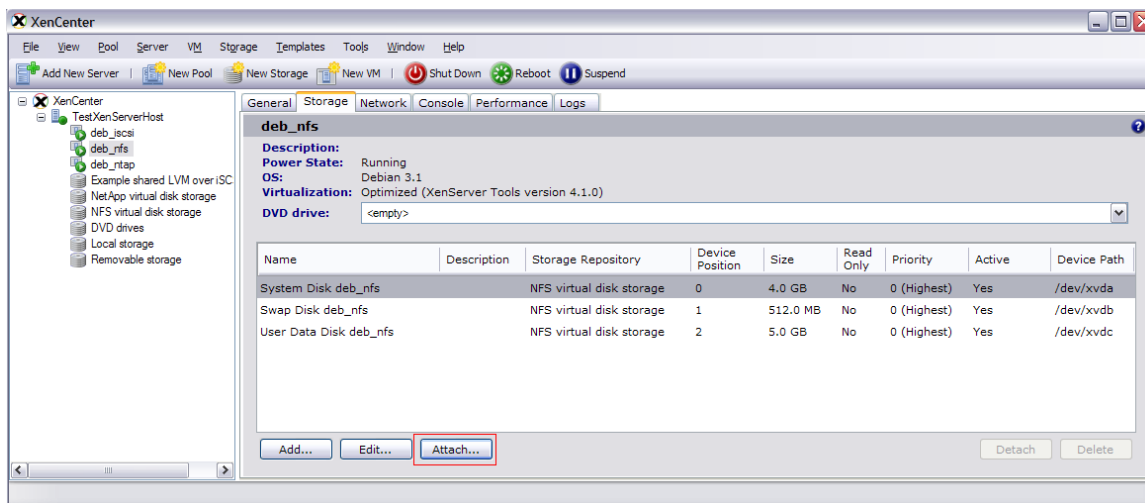
1. In the NetApp FilerView, click the *SnapMirror* tab
2. Select the relevant relationship, and click on *Resync* to put the relationship back into a Snapmirrored state.

RECOVERY OF SNAPSHOT VIRTUAL DISKS

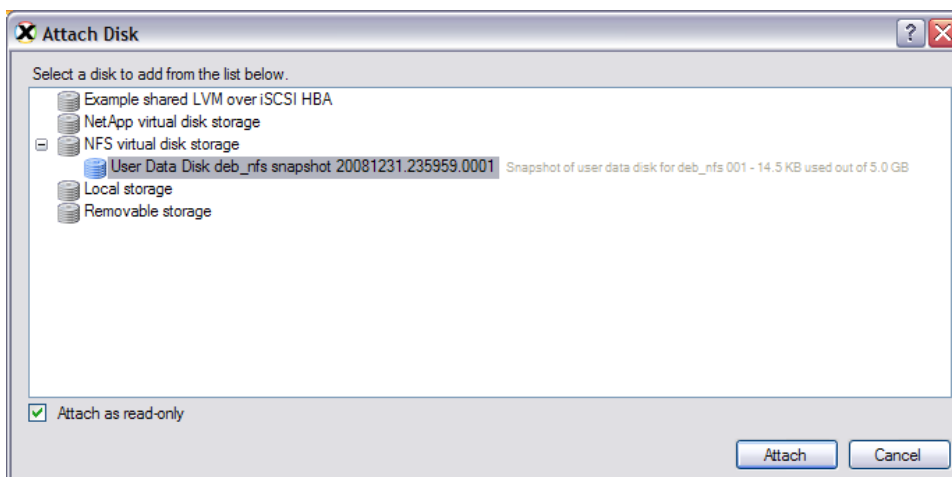
To recover data from the snapshot disks, recover the exported NFS volume or the iSCSI LUNs being used for NetApp Data ONTAP storage repository as listed in the previous section.

Once the volume has been restored on the NetApp filer, perform the following steps to recover data. Using the above for virtual machine "deb_nfs", these steps illustrate how to recover the data for "User Data Disk deb_nfs"

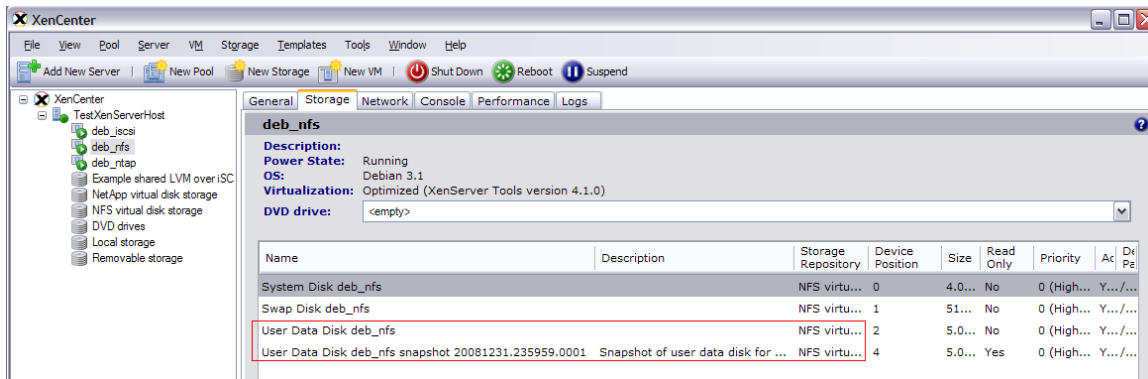
1. In the XenCenter console, click on the XenServer host and then the virtual machine that you want to recover data on. Click the *Storage* tab and then *Attach*.




2. In the *Attach Disk* wizard, select the snapshot VDI that will be used to restore data off, and check *Attach as read-only*. Click *Attach* to hot-plug the snapshot disk to the virtual machine.



3. Both the original disk and the snapshot disk should now be attached to the virtual machine. From this point on, data can be copied from the snapshot disk back into the original disk.



4. Once the data has been copied over, go back to XenCenter, and click on the snapshot disk in the *Storage* tab for the virtual machine. Click *Detach* to hot-unplug the snapshot disk from the virtual machine. Click *Yes* on the confirmation dialog box.



SUMMARY

While other virtualization products consume valuable server resources to run proprietary management tools to manage storage services, Citrix XenServer enables each hardware component to do what it does best: NetApp storage systems can deliver optimal storage features and performance, and servers can direct their power to business-critical application workloads.

Above all, by incorporating the best practices for configuration, allocation and administration into the NetApp Data ONTAP Adapter, Citrix XenServer makes optimal management of storage in virtualized environment seamless – allowing organizations to unify virtualization management for increased agility and reduced management costs.



APPENDIX A

SCRIPT TO PERFORM SNAPSHOT

The script below is an example of how a virtual machine's disks may be snapshot for backup and recovery. . .

```
#!/usr/bin/env python
# Copyright (c) 2008 Citrix Systems, Inc.
#
# Permission to use, copy, modify, and distribute this software for any
# purpose with or without fee is hereby granted, provided that the above
# copyright notice and this permission notice appear in all copies.
#
# THE SOFTWARE IS PROVIDED "AS IS" AND THE AUTHOR DISCLAIMS ALL WARRANTIES
# WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF
# MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR
# ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES
# WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN
# ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF
# OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

import sys, time
import XenAPI

def main(session, vmname):

    vms = []
    vms = session.xenapi.VM.get_by_name_label(vmname)
    if vms == []:
        print "Could not find any VM with name %s" % vmname
        sys.exit(1)

    vbds = session.xenapi.VM.get_VBDs(vms[0])
    epoch = time.strftime("%Y%m%d-%H%M%S", time.localtime())
    for vbd in vbds:
        vbd_record = session.xenapi.VBD.get_record(vbd)
        if vbd_record['type'] == 'Disk':
            vdi_record = session.xenapi.VDI.get_record(vbd_record['VDI'])
            vdi_snapshot = session.xenapi.VDI.snapshot(vbd_record['VDI'], \
                {"epochhint":epoch})

            vdi_clone = session.xenapi.VDI.clone(vdi_snapshot)
            session.xenapi.VDI.set_name_label(vdi_clone, \
                vdi_record['name_label'] + "-" + epoch)
            print "Created snapshot for %s => %s" % \
                (vdi_record['name_label'], session.xenapi.VDI.get_name_label(vdi_clone))
            session.xenapi.VDI.destroy(vdi_snapshot)

    session.xenapi.session.logout()

if __name__ == "__main__":
    if len(sys.argv) <> 2:
        print "Usage:"
        print sys.argv[0], "<name of VM to snapshot>"
        sys.exit(1)
    vmname = sys.argv[1]
    session = XenAPI.Session("http://_var_xapi_xapi", transport=XenAPI.USTransport())
    session.xenapi.login_with_password("username", "password")
    try:
        main(session, vmname)
    except Exception, e:
        print str(e)
        raise
```

APPENDIX B

SECURITY CONSIDERATIONS

To allow XenServer administrators to login into the filer with root privileges and create the FlexVols for the NetApp Storage Repository, use the following guidelines.

Equating Windows 'Domain\Administrator' privileges to UNIX 'root' in a multi-protocol environment

1. To equate Windows 'Domain\Administrator' privileges to UNIX 'root' in a multi-protocol environment, on the filer enter:

```
filer> options wafl.nt_admin_priv_map_to_root on
```

2. Authorizing a Unix User Account to Login As Root on the Filer

Data ONTAP uses the `/etc/usermap.cfg` file to map user names. In its simplest form, each `/etc/usermap.cfg` entry contains a pair of names: the Windows name and the UNIX name. Data ONTAP can translate the Windows name to the UNIX name or vice versa.

When a connection is started, if the `/etc/usermap.cfg` file is missing, a default file is created. It contains commented-out sample map entries that are useful for improving security. When Data ONTAP receives a connection request from a user, it searches the `/etc/usermap.cfg` file to see whether an entry matches the user's Windows domain name and user name. If an entry is found, Data ONTAP uses the UNIX name specified in the entry to look up the UID and GID from the UNIX password database. If the UNIX name is a null string, Data ONTAP denies access to the user.

If an entry is not found, Data ONTAP converts the Windows name to lowercase and considers the UNIX name to be the same as the Windows name. Data ONTAP uses this UNIX name to look up the UID and GID from the UNIX password database. Data ONTAP <http://www.citrix.com/English/partners/partner.asp?partnerID=950197> scans the file sequentially. It uses the first matching entry for mapping.

For information about character coding of the `/etc/usermap.cfg` file, see the information about the contents of the `/etc` directory in the Storage Management Guide.

Specify each entry using the following format:

```
[IP_qualifier:] Windows_name [direction] [IP_qualifier:] UNIX_name
```

where

IP_qualifier field is an IP address that qualifies the user name by narrowing the match.

Windows_name field consists of a Windows domain name, which is optional, and a Windows user name.

Direction field indicates the direction of the mapping.

UNIX_name field is a UNIX name in the UNIX password database.

You can embed comments in the file by beginning the comment lines with `#`. Comments at the end of an entry are also allowed if preceded by `#`. Blank lines are ignored.

The way in which Data ONTAP interprets a domain name in the `/etc/usermap.cfg` file that contains a dot depends on whether storage system is in a Windows NT domain or a Windows Active Directory domain. Follow some guidelines to keep entries simple and easy to understand, and add several entries to the `/etc/usermap.cfg` file to prevent unauthorized users from accessing the storage system.

ACKNOWLEDGEMENTS

Chaitanya Upadhyay, *Principal Partner Solutions Manager, Virtualization and Management Division, Citrix Systems, Inc.*

Julian Chesterfield, *Senior Software Engineer, Virtualization and Management Division, Citrix Systems, Inc.*

Wim Colgate, *Senior Software Engineer, Virtualization and Management Division, Citrix Systems, Inc.*

Alan Oehler, *Technical Writer, Virtualization and Management Division, Citrix Systems, Inc.*

Christopher Chandler, *Lead Reference Architect, Virtualization and Grid Infrastructure Business Unit, NetApp*

Preetom Goswami, *Reference Architect, Virtualization and Grid Infrastructure Business Unit, NetApp*

Eric Forgette, *Reference Architect, Virtualization and Grid Infrastructure Business Unit, NetApp*

Rachel Zhu, *Reference Architect, Virtualization and Grid Infrastructure Business Unit, NetApp*