



# Citrix XenServer and NetApp Storage Best Practices

Citrix Systems, Inc. and NetApp, Inc.

February 2011 | TR-3732 | Rev 3.0

**Table of Contents**

The Challenge of Today’s Enterprise..... 3

Citrix XenServer for Enterprise-Ready Virtualization..... 4

NetApp Storage Solutions for Enterprise-Ready Virtualization..... 5

Overview of XenServer storage ..... 5

    Storage repositories (SRs)..... 5

    Virtual Disk Images (VDIs)..... 5

    Managing Storage..... 6

    XenServer Shared Storage Options ..... 7

Configuration and Setup ..... 16

    NetApp Device Configuration Basics ..... 16

Fixing Misaligned Windows Guest Operating System ..... 85

    Resolution ..... 88

Backup and Recovery ..... 89

    Snapshot for Virtual Machines on NetApp Data ONTAP and NFS Storage Repositories..... 89

    Backing up VM Metadata Information on NetApp Device..... 94

    Backing up Storage Repositories on NetApp Device ..... 97

Summary ..... 113

Appendix A ..... 114

    Script to perform Snapshot and Recovery of VMs..... 114

Appendix B ..... 119

    Security Considerations ..... 119

Appendix C ..... 120

    Enhancements in ONTAP 7.3.1 ..... 120

    References ..... 121

    Citrix Worldwide ..... 121

## 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 XenApp, Citrix XenDesktop 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 reimplementation. 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
- Transparent Storage Cache Sharing to increase I/O performance
- Flash Cache to help reduce virtual desktop storm activities

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.

NOTE: Make sure your configuration is supported by contacting NetApp sales engineer to check [NetApp IMT tool](#) at NetApp [NOW](#) site.

## 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 host. 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 device, there are 4 options; the Citrix StorageLink feature, an NFS file share, an iSCSI LUN or a Fibre Channel LUN.

### Virtual Disk Images (VDIs)

There are two fundamental VDI types (Files and LUNs) that can be accessed with a NetApp device as the backend over 4 different SR driver types:

- NetApp managed LUNs: Managed NetApp LUNs are accessible via the StorageLink feature included in Citrix Essentials for XenServer, and are hosted on a NetApp device running a version of Data ONTAP 7.0 or greater. LUNs are allocated on demand via StorageLink and mapped dynamically to the host via the StorageLink service while a VM is active. All the thin provisioning and fast clone capabilities of the device are exposed via StorageLink.
- VHD files. The VHD format can be used to store VDIs in a sparse format. Being sparse, the image file grows proportionally to the number of writes to the disk by the Virtual Machine (VM), so large portions of the disk which are typically unused do not consume unnecessary space. VHD on NFS, iSCSI, or Hardware HBA storage repositories can be shared among all hosts in a pool.

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

## Managing Storage

There are four XenServer 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. SRs are long-lived, and may be shared among XenServer hosts or moved between them.
- 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. There are several classes of XenServer SRs available for use with NetApp storage:
  - NFS. With an NFS SR, a new directory on an existing NetApp NFS share is created for the storage of VDIs in VHD format.
  - iSCSI and Fibre Channel. The creation of iSCSI or Fibre Channel (Hardware HBA) SRs involves erasing any existing data on a specified LUN. A LUN will need to be created on the NetApp storage before creating the XenServer SR. Volume management is performed via LVM (Logical Volume Manager), and the underlying VDI storage on an iSCSI or FC SR is VHD.
  - Direct StorageLink-NetApp adapter SR. First introduced in XenServer 4.1, this type of SR is created within XenCenter. With this type of SR, FlexVol resources are allocated on the NetApp device and communication is performed via the iSCSI protocol.
  - StorageLink Gateway. New with XenServer 5.6, this type of SR is initially created from the StorageLink Manager and visible within XenCenter thereafter. The StorageLink Gateway service runs on a general-purpose Windows server or VM. Once brokered by the StorageLink Gateway, the data path between the virtual machine and storage system can be performed via either iSCSI or Fibre channel. It is important to note that the StorageLink Gateway only serves as a broker and a control path; it is not part of the data path between a XenServer host and the NetApp storage infrastructure.
- 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.

- 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.




Figure 1 Graphical Overview of Storage Repository and Related Objects

## XenServer Shared Storage Options

When using a NetApp device as your networked, backend storage array, it is recommended with XenServer v5.6 onwards to use the StorageLink Gateway SR type. The architecture of the StorageLink Gateway allows XenServer to utilize the capabilities of the NetApp device to provide data efficiency, high performance and ensure compatibility with existing ONTAP device 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 active-active controller configuration, 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 with the StorageLink Gateway

The StorageLink Gateway is a feature offered with Citrix Essentials for XenServer 5.5, Citrix XenServer 5.6 and Citrix XenServer 5.6 Feature Pack 1, Enterprise Edition. It acts as a broker between XenServer and the NetApp storage system. The StorageLink Gateway service runs on a general-purpose Windows server or VM, and this system can typically be shared with other Citrix Essentials for XenServer components such as Workload Balancing and Workflow Studio. Once brokered by the StorageLink Gateway, the data path between the virtual machine and storage system can be performed via either iSCSI or Fibre channel. It is important to note that the StorageLink Gateway only serves as a broker and as part of the control path; it is not part of the data path between a XenServer host and the NetApp storage infrastructure.

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.



### Shared Storage using the direct StorageLink-NetApp adapter

The direct StorageLink adapter for NetApp was first introduced in XenServer 4.1. While it is still available in XenServer 5.5, XenServer 5.6 and XenServer 5.6 Feature Pack 1, it is primarily intended to maintain backward compatibility with pre-5.5 and 5.6 and 5.6 Feature Pack 1 deployments. Very small environments and Linux-centric deployments may also find this StorageLink option to be useful. For XenServer 5.6 and XenServer 5.6 FP 1 deployments, it is recommended to use the StorageLink Gateway.

The following table provides a comparison of the two StorageLink options for NetApp.

	StorageLink Gateway	Direct StorageLink-NetApp Adapter
Deployment	On separate Windows VM or server	Included within XenServer control domain; no additional VMs or servers are required
VM / LUN relationship	LUN per VDI	LUN per VDI
Protocol support	iSCSI and Fibre Channel	iSCSI
Interoperability with Hyper-V	Yes	No
Interoperability with other vendor storage arrays	Yes	Limited

This legacy StorageLink-NetApp adapter for NetApp Data ONTAP uses the Zephyr API (ZAPI) interface to the device to create a group of FlexVol®s which corresponds to a XenServer SR. VDIs are created as virtual LUNs on the device, 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. Similar to the StorageLink Gateway, the legacy StorageLink adapter exposes LUNs to the XenServer host as VDIs, and there is a one-to-one mapping between the Virtual Machine disks in the XenServer environment and the NetApp disk data abstraction.

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 2 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 device 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 3 Shared NFS storage

## The Linux Host Utilities

The Host Utilities provide software programs and documentation that you can use to connect your XenServer host to NetApp storage systems running Data ONTAP. The software is available as an .rpm file that you can download from the [NOW](#) site. It is a best practice to install Linux Host Utility on your XenServer.

### The Host Utilities include the following components:

- The SAN Toolkit

The toolkit is installed automatically when you install the Host Utilities. This kit provides the following key tools:

Note: This toolkit is common across all configurations and protocols of the Linux Host Utilities. As a result, some of its contents apply to one configuration, but not another. Having unused components does not affect your system performance.

- The `sanlun` utility, which helps you to manage LUNs and HBAs.
- The `san_version` command, which displays the versions of the Host Utilities.

Note: Previous versions of the Host Utilities also included diagnostics programs. These programs have been replaced by the nSANity Diagnostic and Configuration Data Collector and are no longer installed with the Host Utilities. The nSANity program is not part of the Host Utilities. You should download, install, and execute it only when requested to do so by technical support.

## Protocols and configurations supported by Host Utilities

The Host Utilities provide support for Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI connections to the storage system using direct-attached, fabric-attached, and Ethernet network configurations.

These protocols enable the host to access data on storage systems. The storage systems are targets that have storage target devices called LUNs. The protocol enables the host to access the LUNs to store and retrieve data.

## 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 XenServer LVHD, a technology that combines the Logical Volume Manager (LVM) and Virtual Hard Disk (VHD) standards. Virtual machine VDIs are stored on an iSCSI LUN created on the NetApp storage system. Shared storage with iSCSI is a good choice for general purpose virtualization deployments, though it may not be as suitable for demanding workloads or deployments.




Figure 4 Shared iSCSI storage using Open iSCSI initiator




Figure 5 Shared iSCSI storage using iSCSI HBA

### Shared Fibre Channel Storage

XenServer hosts support Fibre Channel SANs using an 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 LVHD technology with the same benefits as iSCSI storage, just utilizing a different data I/O path.




Figure 6 Shared FC storage

## Configuration and Setup

### NetApp Device Configuration Basics

This section covers the best practices for configuring the NetApp active-active controller configuration.

#### Configuring Target Portal Groups for iSCSI connections

Considerations when using Target Portal Groups: (for more information, please consult the Data ONTAP admin guide located at <http://now.netapp.com>)

1. Adding or removing a NIC may change the target portal group assignments. Be sure to verify the target portal group settings are correct after adding or removing hardware, especially in active-active controller configuration.
2. When used with multi-connection sessions, the Windows iSCSI software initiator creates a persistent association between the target portal group tag value and the target interfaces. If the tag value changes while an iSCSI session is active, the initiator will be able to recover only one connection for a session. To recover the remaining connections, you must refresh the initiator's target information.

If you want to employ multi-connection sessions to improve performance and reliability, you must use target portal groups to define the interfaces available for each iSCSI session.

Create a target portal group that contains all of the interfaces you want to use for one iSCSI session. However, note that you cannot combine iSCSI hardware-accelerated interfaces with standard iSCSI storage system interfaces in the same target portal group. When you create a target portal group, the specified interfaces are removed from their current groups and added to the new group. Any iSCSI sessions using the specified interfaces are terminated, but the initiator should automatically reconnect. However, initiators that create a persistent association between the IP address and the target portal group will not be able to reconnect.

1. Open console connection to the NetApp active-active controller configuration, and run the *iSCSI tpgroup* command

```
iscsi tpgroup create [-f] tpgroup_name [-t tag] [interface ...]
```

*-f* forces the new group to be created, even if that terminates an existing session using one of the interfaces being added to the group.

*tpgroup\_name* is the name of the group being created (1 to 32 characters, no spaces or non-printing characters).

*-t tag* sets the target portal group tag to the specified value. In general you should accept the default tag value. User-specified tags must be in the range 1 to 256.

*interface ...* is the list of interfaces to include in the group, separated by spaces.

For example, the following command creates a target portal group named *xenstorage\_iscsi\_group* that includes interfaces *e0c* and *e0d*.

```
iscsi tpgroup create xenstorage_iscsi_group e0c e0d
```

2. Enable target portal group interfaces to accept iSCSI connections

```
iscsi interface enable e0c
iscsi interface enable e0d
```

## Configuring a Dynamic Multimode Virtual Interface (VIF)

Throughput can be improved by creating a multimode VIF. With a multimode VIF, all interfaces in the VIF are active and share a single IP address. This logical aggregation of interfaces can be effectively used for NFS connections.

The following prerequisite must be met before creating a multimode VIF:

1. Identify or install a switch that supports link aggregation over multiple port connections in your network, configured according to your switch vendor's instructions.
2. Decide on a case-sensitive name for the VIF that meets the following criteria:
  - a. It must begin with a letter.
  - b. It must not contain a space.
  - c. It must not contain more than 15 characters.
  - d. It must not already be in use for a VIF.
3. Decide on the interfaces you want the VIF to consist of.
4. Configure all interfaces that will be included in the VIF to be down using the *ifconfig* command.

To create a dynamic multimode (LACP) VIF for NFS datapaths:

1. Open a console session to the NetApp active-active controller configuration, and run the *ifconfig down* command for interfaces that you want to create the vif on.

```
ifconfig down e0c
ifconfig down e0d
```

2. Run the *vif* command to create the new vif
 

```
vif create [multi|lacp] <vif_name> -b ip [<interface_list>]
```

It is recommended to use the *lacp* option and use IP address-based load balancing algorithm.

```
vif create lacp nfs_bond -b ip e0c e0d
```

3. Run the *ifconfig* command on the newly created vif to assign it an IP address
 

```
ifconfig nfs_bond 192.168.2.203 netmask 255.255.255.0
```
4. Create a corresponding interface on the partner controller and add the following lines to the */etc/rc* file. This will allow the configuration to be persistent in case of reboot. The *partner* command makes the configured IP Address be taken over by the cluster partner in case of failover
 

```
vif create lacp nfs_bond -b ip e0a e0c
ifconfig nfs_bond 192.168.2.203 netmask 255.255.255.0 partner nfs_bond
```

## Configuring an Aggregate


An aggregate consists of disk drives; you must specify a number of disks to be added to the new aggregate. Aggregates are the means by which the total IOPs available to all of the physical disks are pooled as a resource. NetApp recommends that whenever possible a small aggregate should be used as the root aggregate. The root aggregate stores the files required for running and providing GUI management tools for the FAS system. The remaining storage should be placed into a small number of large aggregates. On smaller FAS arrays, it may not be practical to have more than a single aggregate, due to the restricted number of disk drives on the system. In these cases, it is acceptable to have only a single aggregate.

Note: Data ONTAP 8.x has support for 64-bit aggregates, which allow for the creation of larger aggregate sizes. Please consult <http://now.netapp.com> for the maximum aggregate size for your storage system model.


NetApp System Manager is used throughout this document and only the SnapMirror section uses the NetApp FileView. To download the System Manager please visit <http://now.netapp.com>.

Data aggregates should have RAID groups of no less than twelve.


1. Open the NetApp System Manager, and click *Aggregates->Create* to add a new aggregate on the device.




- 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. NetApp RAID-DP® is an advanced RAID technology that is provided as the default RAID level on all FAS systems. RAID-DP protects against the simultaneous loss of two drives in a single RAID group. It is very economical to deploy; the overhead with default RAID groups is a mere 12.5%. This level of resiliency and storage efficiency makes data residing on RAID-DP safer than data stored on RAID 5 and more cost effective than RAID 10. NetApp recommends using RAID-DP on all RAID groups that store XenServer data.



- Select the disk that will be used for the aggregate. *Automatic* is selected by default for this section.
- Choose the disk size to be used in the aggregate. By default *Any Size* will be selected.



5. 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

Thick provisioning preallocates storage; thin provisioning provides storage on demand. Thin provisioning allows the admin to overprovision storage in order to avoid running out of storage and the associated application downtime when expanding the provisioned storage. Although no system can be run at 100% storage utilization, there are methods of storage virtualization that allow administrators to address and oversubscribe storage in the same manner as with server resources (such as CPU, memory, networking, and so on). This form of storage virtualization is referred to as *thin provisioning*.

The value of thin-provisioned storage is that storage is treated as a shared resource pool and is consumed only as each individual VM requires it. This sharing increases the total utilization rate of storage by eliminating the unused but provisioned areas of storage that are associated with traditional storage. The drawback to thin provisioning and oversubscribing storage is that (without the addition of physical storage) if every VM requires its maximum possible storage at the same time, there will not be enough storage to satisfy the requests 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

NetApp recommends that when you enable NetApp thin provisioning, you also configure storage management policies on the volumes that contain the thin-provisioned LUNs. These policies aid in providing the thin-provisioned LUNs with storage capacity, as they require it. The policies include automatic sizing of a volume and automatic Snapshot deletion.

Volume Auto Size is a policy-based space management feature in Data ONTAP that allows a volume to grow in defined increments up to a predefined limit when the volume is nearly full. For Citrix environments, NetApp recommends setting this value to 'on'. Doing so requires setting the maximum volume and increment size options.

1. Log in to NetApp console.
2. Set Volume Auto Size Policy

```
Vol autosize <vol-name> [-m <size>[k|m|g|t]] [-i <size>[k|m|g|t]] on.
```

Snapshot Auto Delete is a policy-based space-management feature that automatically deletes the oldest

Snapshot copies on a volume when that volume is nearly full. NetApp recommends setting this value to delete Snapshot copies at 5% of available space. In addition, you should set the volume option to have the system attempt to grow the volume before deleting Snapshot copies.

1. Log in to NetApp console.
2. Set Snapshot Auto Delete Policy
3. snap autodelete <vol-name> commitment try trigger

```
snap autodelete <vol-name> commitment try trigger volume target_free_space 5
delete_order oldest_first.
```

## Fractional Reservation

LUN Fractional Reserve is a policy that is required when you use NetApp Snapshot copies on volumes that contain XenServer LUNs. This policy defines the amount of additional space reserved to guarantee LUN writes if a volume becomes 100% full. For XenServer environments where Volume Auto Size and Snapshot Auto delete are in use, NetApp recommends setting this value to 0%. Otherwise, leave this setting at its default of 100%. To disable Fractional Reserve, follow these steps.

1. Log in to NetApp console.
2. Set Volume Snapshot Fractional Reserve:

```
vol options <vol-name> fractional_reserve 0.
```

## Using NetApp De-duplication

NetApp deduplication 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.

Deduplication is enabled on a per flexible volume basis. It can be enabled on any number of flexible volumes in a storage system. It can be run one of four different ways:

- Scheduled on specific days and at specific times
- Manually, by using the command line
- Automatically, when 20% new data has been written to the volume
- Automatically on the destination volume, when used with SnapVault®

Only one deduplication process can run on a flexible volume at a time.

Up to eight deduplication processes can run concurrently on eight volumes within the same NetApp active-active controller configuration.

Beginning with Data ONTAP 7.3.1, deduplication checkpoint restart allows a deduplication process that was interrupted to continue from the last checkpoint. Prior to Data ONTAP 7.3.1, an interrupted deduplication process would result in a restart of the entire deduplication process. If the system is restarted while deduplication is in process, when the system is once again online, the deduplication process automatically restarts from the last checkpoint.

Below are some results on the testing NetApp deduplication

- Supported by all storage data access types; iSCSI, FCP & NFS
- >70% with virtualization environment
- Saves up to 95% for full backups; 25% to 55% for most data sets.

If NetApp deduplication is going to be enabled for the FlexVol® from XenCenter or StorageLink, then note that the volume size should match the maximum supported deduplication limit for the device.

The maximum flexible volume size limitation for deduplication varies based on the platform (this number depends primarily on the amount of system memory). When this limit is reached, writes to the volume fail just as they would with any other volume after it is full.

This could be important to consider if the flexible volumes are ever moved to a different platform with a smaller maximum flexible volume size. For current volume limits, please consult TR-3505: NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide, available at <http://www.netapp.com>

The maximum shared data limit per volume for deduplication is 16TB, regardless of the platform type. Once this limit is reached, there is no more deduplication of data in the volume, but writes to the volume continue to work successfully until the volume is completely full.

## 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 active-active configuration. The recommended storage configuration is to utilize the StorageLink Gateway SR type since it provides the most optimized performance as well as maximum flexibility in configuration and protocol selection.

### Storage Networking

1. Configure physical network infrastructure so that different traffic is on different subnets.
2. Configure a new network to make the bond and/or VLAN available
3. Create a management interface to use the new network

The next few sections will cover how multipathing and NIC bonding can be used for storage networks. The above three steps will be enumerated in each section.

### Storage Multipathing

XenServer 5.0 onwards Active/Active multipathing for iSCSI and FC protocols for I/O data paths is introduced. It is recommended to use the following diagram as a guide to configure multipathing (Dynamic MultiPathing or DMP)




Figure 7 Graphical representation of DMP configuration


Dynamic multipathing uses a round-robin mode load balancing algorithm, so both routes will have active traffic on them during normal operations. Multipathing can be enabled via XenCenter or on the command line. However, before attempting to enable multipathing, verify that multiple targets are available on your storage server.

XenServer 5.6 and XenServer 5.6 FP 1 support ALUA, asymmetric logical unit access. ALUA is a relatively new multipathing technology for asymmetric arrays. NetApp is ALUA compliant. Note ALUA is only available with fiber channel.


## Dedicate NICs for software iSCSI storage traffic

Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1


1. From *XenCenter*, click on the pool (or standalone server) in the *Resources* tab. Click on the *Network* option followed by the *Configure* button




2. On the *Management interfaces* section, click on the *New Interface* button.



3. Give the new interface a recognizable name, and select the *Network* you want the dedicated interface on.



4. Click on the *Use these IP settings:* radio button and enter a starting IP address for the NICs in the *Network*.



5. Repeat the above steps for each NIC dedicated to storage, and click *OK*
6. Open a console session to the XenServer pool master, and run the *iscsiadm* command with *discovery* and *sendtargets* option.

```
[root@antisana-v5 ~]# iscsiadm -m discovery --type sendtargets --portal 192.168.2.200
```

```
192.168.2.200:3260,1 iqn.1992-08.com.netapp:sn.101173424
192.168.3.200:3260,1 iqn.1992-08.com.netapp:sn.101173424
```

The command returns multiple targets on the NetApp active-active controller configuration showing that the backend storage device is configured properly.

## Configuring Multipathing for iSCSI

1. Modify `/etc/multipath-enabled.conf` file to include the following in the `devices` section.

```
device
{
  vendor                "NETAPP"
  product               "LUN"
  path_grouping_policy  group_by_prio
  getuid_callout        "/sbin/scsi_id -g -u -s /block/%n"
  prio_callout          "/sbin/mpath_prio_ontap /dev/%n"
  features               "1 queue_if_no_path"
  path_checker          directio
  failback              immediate
}
```

## Configuring Multipathing for Fiber Channel

1. Open a console session to the NetApp active-active controller configuration, set `igroup` to ALUA.

```
igroup set <initiator_group> alua yes
igroup show -v
```


2. Modify `/etc/multipath-enabled.conf` file to include the following in the `devices` section

```
defaults
{
  user_friendly_names no
  max_fds max
  queue_without_daemon no
}
devices
{
  device{
  vendor "NETAPP"
  product "LUN"
  getuid_callout "/sbin/scsi_id -g -u -s /block/%n"
  prio_callout "/sbin/mpath_prio_alua /dev/%n"
  features "1 queue_if_no_path"
  hardware_handler "0"
  path_grouping_policy group_by_prio
  failback immediate
  path_checker directio
  flush_on_last_del yes
  }
}
```


## Enable Multipathing at XenCenter

Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1

1. From *XenCenter*, right click the server in the pool from the *Resources* pane, and select the option *Enter Maintenance Mode...*




Select *Enter maintenance mode* from the dialog box. Choose any other server node in the pool as the temporary master.




There is a short delay while *XenCenter* migrates any active virtual machines and unplugs the existing storage; if the server is a pool master, it will be disconnected and may disappear from the *Resources* pane temporarily while a new pool master is assigned.

2. When connection is restored, click on the server and then the *General* tab.




Click the *Properties* button and then click on the *Multipathing* tab. Select the *Enable multipathing on this server* check box, and click *OK*.



There is a short delay while *XenCenter* saves the new storage configuration.

3. Take the server back out of Maintenance mode by right clicking on the server in the *Resources* pane and click *Exit Maintenance Mode*.



4. Repeat the above steps for all servers in the pool that multipathing needs to be enabled on.

**NIC Bonding**

Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1

NIC bonds can improve XenServer host resiliency by using two physical NICs as if they were one. If one NIC within the bond fails the host's network traffic will automatically be routed over the second NIC. NIC bonds work in an Active/Active mode, with traffic balanced between the bonded NICs on XenServer 5.6. On XenServer 5.6 Feature Pack 1, NIC bonds can work in either an active/active mode, with traffic balanced between the bonded NICs, or in an active/passive mode.

As such, XenServer host NICs should be bonded for NFS traffic to the NetApp active-active controller configuration. It is recommended to use the following diagram as a guide to configure NIC bonding.






Figure 8 Graphical representation of NIC bonding configuration


1. From XenCenter, click on the pool in the *Resources* pane (or server if this is a standalone server), and click on the *Network* tab.




2. Click on the *Add Network...* button. Select the *Bonded Network* radio button, and click *Next>*.




3. Select the NICs that need to be bonded from the *Available NICs* pane. Click *Finish*.




4. This will create a new network named *Bond 0 + 1*.




- Select *Bond 0 + 1*, and click the *Properties* button to give the new network a more meaningful name and description.



- On the *Management interfaces* section, click on the *Configure*, and select *New Interface* button



7. Click on *New Interfaces* in the *Management interfaces* wizard.
8. Give the new interface a name, and choose the *Network* created in previous steps from the drop down. Click *OK*.




### Integrated StorageLink SR configuration

You can now use XenCenter to create new Storage Repositories (SRs) that use the existing StorageLink storage infrastructure, allowing you to access a range of different native storage services using any of the leading architectures and protocols, including DAS, NAS, SAN, iSCSI and Fibre Channel. XenCenter's new StorageLink SR management capabilities allow you to create, remove and use StorageLink SRs without having to install the StorageLink Manager. From XenCenter, you can:

- View StorageLink servers and storage architecture – storage systems, pools, volumes, and SRs – in the Resources pane. Here, you can see information about arrays including capabilities, name, and serial number, and also see the amount of free space and capacity on your storage volumes and pools.
- View, add and remove Storage credentials.
- Connect to a StorageLink license server.
- Create and destroy SRs on storage managed by StorageLink.

Installation and configuration of the StorageLink service must be done outside of XenCenter; for more information, please refer to the StorageLink documentation.



## Shared Storage using StorageLink Gateway


Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1

Use of the StorageLink Gateway for shared storage is the recommended approach for deployments of XenServer with NetApp. This type of storage repository is available with Citrix XenServer, Enterprise Edition, Platinum Edition and requires the installation of the StorageLink Gateway service on a Windows Server virtual machine or physical server. In addition to XenServer, StorageLink Gateway also supports Hyper-V hosts and provides storage interoperability between the two platforms. As mentioned before, the StorageLink Gateway only serves as a broker, and is not part of the data path between a XenServer host and the NetApp storage infrastructure.


If installing the StorageLink Gateway on a XenServer virtual machine, the virtual machine disks will need to be connected via a standard NFS, iSCSI, or Hardware HBA storage repository type. For highly available configurations, it is recommended that StorageLink be implemented in an isolated resource pool and using an iSCSI or Hardware HBA storage repository type (ideally the same SR used for the XenServer HA heartbeat disk). The StorageLink Gateway services do not need to be available while VMs are running, however many operations such as VM start require the service for a short period of time.

Once the StorageLink Gateway Service and Manager (management console) have been installed on a virtual machine or physical server, the following process can be followed to create the storage repository.


1. Open the StorageLink Manager console.




- Click "Connect to the StorageLink Gateway" in the center pane, and enter the credentials entered during installation.




3. Add the Hypervisor hosts. Use the name or IP address of the XenServer resource pool master and the XenServer root credentials. (Note: if you receive a networking error during this step, you may need to temporarily remove NIC bonds from your XenServer host(s). The bonds can be recreated once the hosts have been added to the StorageLink Manager.)




4. Once all of the XenServer hosts have been enumerated, your screen should look something like this:




- Next, add the NetApp storage active-active controller configuration. Select the “Storage Infrastructure” tab and then “Add Storage System.”




6. After the storage system has been discovered, your screen should look similar to this. You will see the any aggregates on the NetApp active-active controller configuration within the StorageLink Manager console (in this case XS1 and Vol0)
7. Create the Storage Repository. Click on the “Create Storage Repository” link in the middle pane.






- After clicking “Finish” an SR named “StorageLink Gateway” (or whatever name you specified) will be visible within both the StorageLink Manager as well as XenCenter, similar to what is shown below.




- In addition, a new volume will be created on the NetApp aggregate selected in step 7. The name of this new volume will start with the name of the SR.
- At this point, the StorageLink Manager console will not be required for regular administration activities. The StorageLink Gateway service does need to be running for many VM operations within XenCenter.

- When new VMs are created on the storage repository, the StorageLink Gateway will automatically provision LUNs on the NetApp active-active controller configuration. For example, after using XenCenter to create a new Debian VM with two disks (one 4 GB and another 512 MB), the following two LUNs will be visible in NetApp FilerView or via NetApp System Manager:



- If a virtual machine snapshot is performed from XenCenter, additional LUNs will be visible in FilerView or NetApp System Manager with a “snapXX” appended to the name, similar to below:




## Shared Storage using the direct StorageLink Adapter for NetApp

Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1


The direct StorageLink adapter for NetApp was first introduced in XenServer 4.1. While it is still available in XenServer 5.5, XenServer 5.6, and now on XenServer 5.6 Feature Pack 1, it is primarily intended to maintain backward compatibility with pre-5.5, 5.6 and 5.6 FP1 deployments. For the new XenServer 5.6 FP 1 deployments, it is recommended to use the StorageLink Gateway.

The direct StorageLink-NetApp 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 device to be able to provision FlexVols and generate Snapshots on the device.

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



2. Provide the name of the device (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 device manages CHAP authentication credentials per host IQN. This must be managed and configured directly on the device if required.


3. The NetApp Data ONTAP adapter will poll the device 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 deduplication 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, NetApp deduplication 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 FAS 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.

Deduplication is enabled on a volume, and the amount of data deduplication realized is based on the commonality of the data stored in a deduplication-enabled volume. For the largest storage savings, NetApp recommends grouping similar operating systems and similar applications into one volume. For deduplication best practices, including scheduling and performance considerations, see TR 3505 NetApp FAS Dedupe: Data Deduplication Deployment and Implementation Guide.


When we consider how many volumes we create, we need to consider the total VM number, backup strategy and space saving. The more same OS VMs in one volume, the more deduplication storage saving. NetApp recommends grouping VMs with similar backup requirement in other words VMs with the same snapshot schedule in one volume.

For example, a 1000 virtual desktops environment, we can put the 250 VMs in one volume and have 4 volumes to host the 1000 VMs. And use SnapMirror to backup gold image volume to disaster recovery site storage. Another example, we can use one volume to host 100 VMs if all the VMs have same snapshot backup requirement.

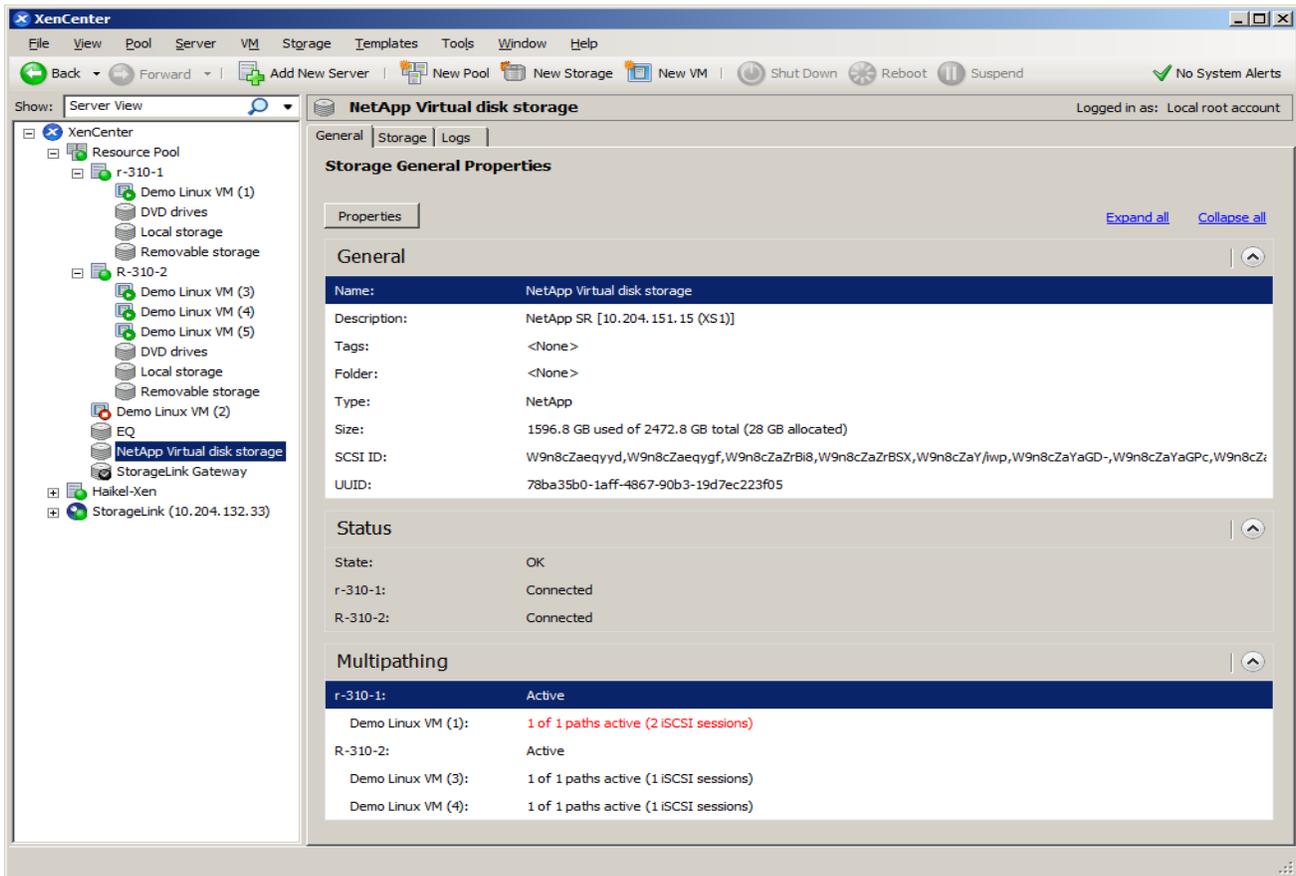


4. Click *Finish* for the storage repository to be created using Data ONTAP for the control path.

NOTE: If multipathing is enabled, the NetApp SR in XenCenter will show multipathing as *Not Active*. The reason is that the NetApp SR type is a LUN-per-VDI model, and no LUNs or VDIs exist when the SR is created.



However, once a VM has been created with VDI's attached and is started up, XenCenter will show multipathing as *Active*. Henceforth, the SR will remain in multipath mode.



## Identifying XenServer objects on the NetApp active-active controller configuration

Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1

The NetApp SR is mapped to a FlexVol in the NetApp active-active controller configuration. Each VDI is represented as a LUN within a FlexVol, with VDI's of the same VM being within a single FlexVol.

When the NetApp SR is created in Direct StorageLink-NetApp Adapter, the FlexVols created on the NetApp active-active controller configuration follow a naming pattern:

`XenStorage_UUID_FVn`

where `UUID` is the UUID of the NetApp SR in XenServer, and `n` is the FlexVol number

The LUNs that are created within the FlexVol use as their name the UUID of the VDI that they represent.

When the NetApp SR is created in the StorageLink gateway, the FlexVol created on the NetApp active-active controller configuration follows a naming pattern: `<CustomizedName>_UUID`.


To view the FlexVols that correspond to the NetApp SR

1. In XenServer console, get the UUID of the NetApp SR using the `xe sr-list` command.

```

[root@r-310-1 ~]# xe sr-list type=netapp
uuid ( RO)                : 78ba35b0-1aff-4867-90b3-19d7ec223f05
name-label ( RW)         : NetApp Virtual disk storage
name-description ( RW)   : NetApp SR [10.204.151.15 (XS1)]
host ( RO)               : <shared>
type ( RO)               : netapp
content-type ( RO)      :
    
```

2. On the NetApp System Manager, click on *Volumes*.



In the NetApp System Manager above, *XenStorage\_78ba35b0\_1aff\_4867\_90b3\_19d7ec223f05\_FV0* to *XenStorage\_78ba35b0\_1aff\_4867\_90b3\_19d7ec223f05\_FV7* are the FlexVols that make up the NetApp SR which has UUID 78ba35b0-1aff-4867-90b3-19d7ec223f05.

On the NetApp active-active controller configuration, the FlexVol can also be displayed using the ONTAP CLI command.

```
vol status
```

```

ISV3140> vol status
Volume State      Status      Options
vol0 online       raid_dp, trad root, nosnap=on
XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV0 online       raid_dp, flex  nosnap=on, guarantee=None
sis
XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV1 online       raid_dp, flex  nosnap=on, guarantee=None
sis
XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV2 online       raid_dp, flex  nosnap=on, guarantee=None
sis
XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV3 online       raid_dp, flex  nosnap=on, guarantee=None
sis
XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV4 online       raid_dp, flex  nosnap=on, guarantee=None
sis
XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV5 online       raid_dp, flex  nosnap=on, guarantee=None
sis
XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV6 online       raid_dp, flex  nosnap=on, guarantee=None
sis
XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV7 online       raid_dp, flex  nosnap=on, guarantee=None
sis
    
```

## To view the LUNs associated with the VDIs in the NetApp SR

1. In XenServer console, get the UUID of VDIs created on the NetApp SR using the `xe vdi-list` command.

```
[root@r-310-1 ~]# xe vdi-list sr-name=label=NetApp\ Virtual\ disk\ storage param
s=uuid,is-a-snapshot
uuid ( RO)                : ab58c2df-8e5b-4e2f-90f2-c47388dabb9b
is-a-snapshot ( RO): false


uuid ( RO)                : 292b393c-f63c-42d0-a306-db5a881de616
is-a-snapshot ( RO): false

uuid ( RO)                : 251d91d6-5073-4f80-af57-26c9c7502dde
is-a-snapshot ( RO): false

uuid ( RO)                : dd620e83-09cf-4549-b9a7-e633cd5c01d1
is-a-snapshot ( RO): false

uuid ( RO)                : 7da4e227-83b4-48fc-9934-432979cf2c38
is-a-snapshot ( RO): false
```

2. In the NetApp System Manager, click on *LUNs->Manage*.



In the above example, /vol/XenStorage\_78ba35b0\_1aff\_4867\_90b3\_19d7ec223f05\_FV4 is the LUN that represents the VDI with UUID ab58c2df-8e5b-4e2f-90f2-c47388dabb9b.

On the NetApp active-active controller configuration, the LUNs can also be displayed via ONTAP CLI command.

```
lun show
```


```
ISV3140> lun show
/vol/XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV0/e0e9e8b9-da34-4d45-ad0b-9b1440c7bd3e 4g (4294967296) (r/w, online, mapped)
/vol/XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV0/f61bf64d-1904-4c0b-b468-3822490c576f 512m (536870912) (r/w, online, mapped)
/vol/XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV2/292b393c-f63c-42d0-a306-db5a881de616 10g (10737418240) (r/w, online)
/vol/XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV4/251d91d6-5073-4f80-af57-26c9c7502dde 4g (4294967296) (r/w, online)
/vol/XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV4/7da4e227-83b4-48fc-9934-432979cf2c38 512m (536870912) (r/w, online)
/vol/XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV4/ab58c2df-8e5b-4e2f-90f2-c47388dabb9b 4g (4294967296) (r/w, online, mapped)
/vol/XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV4/dd620e83-09cf-4549-b9a7-e633cd5c01d1 512m (536870912) (r/w, online, mapped)
/vol/XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV5/5e691437-9b4e-4b40-b7da-cb140d91b7d9 4g (4294967296) (r/w, online, mapped)
/vol/XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV5/6a061e38-e8d2-4cec-b6b8-96ee910e1659 512m (536870912) (r/w, online, mapped)
```

### Configuration Shared NAS using NFS


Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1

To use the NetApp active-active controller configuration 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 NetApp System Manager and discover the NetApp Filer.
2. Click on *Volumes*, and then *Create* 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. Select NAS for Storage type and choose the appropriate aggregate.
6. Set the size required for the NFS SR. If snapshots are not required, you can set 0% snapshot reserve to save space




7. Under the Space settings tab, for storage saving, you can select “None” for Space Guarantee




8. Click *Create* to create the FlexVol.




9. To create an NFS export on the above volume, select *Shared Folder->Shares/Exports*. Click on *Create*.




10. Click browse to choose from a list of Shares/Exports.



11. Select the “XenServer\_NFS\_Storage” export created in step8 and click ok and click Next.






12. Select the NFS check box and give it a new export name, in this case “/vol/XenServer\_NFS\_Storage1” and click Next.




- 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 the appropriate root access and click Next to finalize the export creation..

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.






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



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




By default, the VDI that is created in the NFS SR is thin-provisioned.


## Creating a Fully Provisioned VDI in the NFS SR

Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1

1. In XenCenter, click on the newly created NFS SR in the *Resources* pane, and click the *Storage* tab.



2. Click on *Add Disk...* button, and enter details for the size of VDI you want. Make sure the newly created NFS SR is highlighted, and click *Add*.



- From a XenServer console, find the mount point for the newly created NFS SR using the `mount` command.

```
[root@r-310-1 ~]# mount
/dev/sda1 on / type ext3 (rw)
none on /proc type proc (rw)
none on /sys type sysfs (rw)
none on /dev/pts type devpts (rw)
none on /dev/shm type tmpfs (rw)
/opt/xensource/packages/iso/XenCenter.iso on /var/xen/xc-install type iso9660
(ro,loop=/dev/loop0)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
10.204.151.15:/vol/XenServer_NFS_Storage/e5f7e999-f067-25b4-aa93-23025aad13bb on /var/run/sr-
mount/e5f7e999-f067-25b4-aa93-23025aad13bb type nfs
(rw,soft,timeo=133,retrans=2147483647,tcp,noac,addr=10.204.151.15)
```

In the example above, `/var/run/sr-mount/e5f7e999-f067-25b4-aa93-23025aad13bb` is the mount point.

- Change directory to the mount point, and run the `vhd-util` command to write zeroes into the `.vhd` file that represents the newly created VDI that needs to be fully provisioned

```
[root@r-310-1 ~]# vhd-util
```

```
usage: vhd-util COMMAND [OPTIONS]
```

```
COMMAND := { create | snapshot | query | read | set | repair | resize | fill | coalesce | modify | scan | check |
revert }
```

In the example VDI above, the command is

```
[root@r-310-1 e5f7e999-f067-25b4-aa93-23025aad13bb]# vhd-util fill -n e5f7e999-f067-25b4-aa93-23025aad13bb.vhd
```

You can now attach the VDI to a virtual machine as a fully provisioned virtual disk.

## Configuring iSCSI Storage

Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1

To set up the NetApp active-active controller configuration for an iSCSI SR from XenServer, NetApp System Manager will be used. The first step is to create an Initiator Group and then assign the XenServer IQN initiators to that group.

1. Select LUNs under the Storage option, click on the Initiator Groups tab and click Add.




2. 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*, the operating system as *Xen* and click Add.




3. Click Add on the screen below to assign the XenServer initiators to the newly created Initiator Group.



4. The Add Initiator ID box will appear. Select iSCSI for the Group Type, the Initiator Group that was created in step 2 for Group Name and the IQN initiator from XenServer in Initiator Name. Click Add.



The IQN of a XenServer host can be seen from the *General* tab for the host in XenCenter (in the screen shot 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 *Properties* button at the top right of the *General* tab, and modify the *iSCSI IQN* field (see the screen shot below where IQN is changed to iqn.2008-01.com.example:testdc).




- Repeat Step 4 for all of the XenServer hosts in the pool.


After the Initiator Group has been created and all of the initiator IQNs added, the LUN can be created using the Create LUN Wizard. The wizard gives the ability to use previously created volumes or create both the volume and LUN in a single wizard. In the example below System Manager will create both the LUN and Volume.

Note: SAN Volumes have a 0% Snap Reserve. For further information please refer to the Data ONTAP Block Access Management Guide for iSCSI and FC.


- Select LUNs under the Storage option, click on the LUN Management tab and click on Create.




- Under General Properties, add the name of the LUN, which in this example will also be the name of the Volume, the description of the LUN, the size and select Xen as the type. Click Next to continue.



8. Under LUN Container, choose Automatically create a new volume.



- Under Initiator Mapping, select the Initiator Group that was created in step 2 and click on the right arrow to add the Initiator Group to Hosts to connect.




- Review the selections and click Next to create the Volume and LUN.
- Click on Finish to close the Create LUN Wizard.


## Creating an iSCSI SR Using the iSCSI Software Initiator

Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1


1. To create the storage repository, in XenCenter, choose *New Storage*. Select *iSCSI*.



2. Enter the *Target Host* as the hostname or IP address of the NetApp active-active controller configuration in which the LUN was set up in steps 1-9. If *CHAP Authentication* was set up for iSCSI security on the device, 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 active-active controller configuration. 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*.



3. 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.



## Creating an iSCSI SR Using an iSCSI HBA

Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1

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 1-9 at the beginning of this section that covers LUN creation. However, instead of using the XenServer's IQN, specify the IQN of the various ports of the HBA. Do this for every XenServer host in the pool.

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.

1. Get the IQN of the ports of the iSCSI HBA from the output of *iscli* command.

```

root@Haikel-Xen:/opt/QLogic_Corporation/SANsurferiCLI
Current HBA/Port Information: HBA Alias:
HBA: 0 Port: 0 HBA Port Index: 1 HBA Model: QLA4052C
IP Address: 192.168.0.200 Link: Up
Port iSCSI Name: iqn.2000-04.com.qlogic:qla4052c.gs10649a26116.1
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: QLA4052C
   HBA Serial Number: (GS10649A26116) FW Version: 2.0.0.45 Type: Copper
   IP Address: 192.168.0.200
   Alias:
   iSCSI Name: iqn.2000-04.com.qlogic:qla4052c.gs10649a26116.1
2. HBA: 0 Port: 1 HBA Port Index: 2 HBA Model: QLA4052C
   HBA Serial Number: (GS10649A26116) FW Version: 2.0.0.45 Type: Copper
   IP Address: 192.168.0.201
   Alias:
   iSCSI Name: iqn.2000-04.com.qlogic:qla4052c.gs10649a26116.2
Press the Enter key to continue.

```

- Set the IP address for the HBA. In the control domain, use the *iscli* CLI to do so. Choose option 4, then option 2 to enter *Port Network Settings Menu*. Enter option 4 to *Select HBA Port* and then option 2 to *Configure IP Settings*.

```

root@Haikel-Xen:/opt/QLogic_Corporation/SANsurferiCLI
main Interactive Menu

-----
Program Version: 1.2.00.29 Driver Version: 5.02.03.00.05.07-c2 IC: 2
FW Version: 2.0.0.45 Type: Copper
Current HBA/Port Information: HBA Alias:
HBA: 0 Port: 0 HBA Port Index: 1 HBA Model: QLA4052C
IP Address: 192.168.0.200 Link: Up
Port iSCSI Name: iqn.2000-04.com.qlogic:qla4052c.gs10649a26116.1
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: 4
  
```

```

root@Haikel-Xen:/opt/QLogic_Corporation/SANsurferiCLI
iSNS Port Number      : 3205
iSNS Server Conn Status : closed
Press the Enter key to continue.
Port Network Settings Menu

-----
Program Version: 1.2.00.29 Driver Version: 5.02.03.00.05.07-c2 IC: 2
FW Version: 2.0.0.45 Type: Copper
Current HBA/Port Information: HBA Alias:
HBA: 0 Port: 0 HBA Port Index: 1 HBA Model: QLA4052C
IP Address: 10.204.151.65 Link: Up
Port iSCSI Name: iqn.2000-04.com.qlogic:qla4052c.gs10649a26116.1
Port iSCSI Alias:
-----

1. Display Network Settings
2. Configure IP Settings
3. iSNS Settings
4. Select HBA Port
5. Save changes and reset HBA (if necessary)
6. Refresh
7. Exit
enter selection: 2
DHCP to obtain TCP Information: [off] :
IP Address [10.204.151.65] :
IP_Subnet_Mask [255.255.255.0] :
IP_Gateway [10.204.151.1] :
Port Network Settings Menu
  
```

3. Create an Initiator Group and then assign the IQNs of the HBA to that group.
4. From the NetApp System Manager, select LUNs under the Storage option, click on the Initiator Groups tab and click on Add.




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



6. Click Add on the screen below to assign the iSCSI HBA IQN to the newly created Initiator Group.



7. The Add Initiator ID box will appear. Select iSCSI for the Group Type, the Initiator Group that was created in step 5 for Group Name and the IQN from iSCSI in Initiator Name. Click Add.




8. After the Initiator Group has been created and all of the IQNs added, the LUN can be created using the Create LUN Wizard. The wizard gives the ability to use previously created volumes or create both the volume and LUN in a single wizard. In the example below System Manager will create both the LUN and Volume.

Note: SAN Volumes have a 0% Snap Reserve. For further information please refer to the Data ONTAP Block Access Management Guide for iSCSI and FC.


9. Select LUNs under the Storage option, click on the LUN Management tab and click on Create.




10. Under General Properties, add the name of the LUN, which in this example will also be the name of the Volume, the description of the LUN, the size and select Linux as the type. Click Next to continue



11. Under LUN Container, choose Automatically create a new volume.



12. Under Initiator Mapping, select the Initiator Group that was created in step 2 and click on the right arrow to add the Initiator Group to Hosts to connect.




13. Review the selections and click Next to create the Volume and LUN.
14. Click on Finish to close the Create LUN Wizard.
15. The newly created LUN now needs to be zoned in to the XenServer host and will appear as a SCSI device. The 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>
```

16. Now add a persistent target to the HBA. The target iSCSI IQN can be retrieved from the NetApp FilerView by clicking *Configuration->Protocols->iSCSI->Service*.

```
/opt/QLogic_Corporation/SANsurferiCLI/iscli -pa 0 <iSCSI_target_IP_address>
```



```
root@Haikel-Xen:/opt/QLogic_Corporation/SANsurferiCLI
[root@Haikel-Xen SANsurferiCLI]# iscli -pa 0 10.204.151.65 -INAME iqn.1992-08.com.netapp:sn.151699572
Using config file: /opt/QLogic_Corporation/SANsurferiCLI/iscli.cfg
[root@Haikel-Xen SANsurferiCLI]# iscli -pa 0 10.204.152.65 -INAME iqn.1992-08.com.netapp:sn.151699572
Using config file: /opt/QLogic_Corporation/SANsurferiCLI/iscli.cfg
[root@Haikel-Xen SANsurferiCLI]#
```

Note that the above command was run for 2 ports of the iSCSI HBA, each port connecting to a different subnet (as indicated in the multipathing configuration at the beginning of the XenServer Storage Configuration section). This command will assign the port the specified iSCSI IP address and specific target IQN.


17. Use the *xe sr-probe* command to force a scan of iSCSI 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>
```

```
root@Haikel-Xen:~  
[root@Haikel-Xen ~]# xe sr-probe type=lvmoaha host-uuid=71990813-2eb5-4a5d-9bdf-50f558734912  
Error code: SR_BACKEND_FAILURE_107  
Error parameters: , The SCSIid parameter is missing or incorrect, <?xml version="1.0" ?>  
<Devlist>  
  <BlockDevice>  
    <path>  
      /dev/sdb [sdc]  
    </path>  
    <numpaths>  
      2  
    </numpaths>  
    <SCSIid>  
      360a9800057396e38635a623439723435  
    </SCSIid>  
    <vendor>  
      NETAPP  
    </vendor>  
    <serial>  
      W9n8cZb49r45  
    </serial>  
    <size>  
      10737418240  
    </size>  
    <adapter>  
      3  
    </adapter>  
    <channel>  
      0  
    </channel>  
    <id>  
      64  
    </id>  
    <lun>  
      0  
    </lun>  
    <hba>  
      qllogic  
    </hba>  
  </BlockDevice>
```


To validate that the device path is for the newly created LUN on the device, match the serial number from the `<serial>` field of the `xe sr-probe` output with the serial number of the LUN in FilerView. Note that there are 2 paths to the LUN indicating that multipathing is active.

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.




18. Repeat the above steps for all hosts in the pool.

19. To create the iSCSI SR over HBA on this LUN, in XenCenter click *New Storage* and select *Hardware HBA* option. Click *Next*.




There is short delay while XenServer probes for available LUNs.

20. Select the appropriate LUN. Give the SR an appropriate *Name* and click *Finish*.



21. A warning is displayed that the LUN will be formatted and any data present will be destroyed. Click *Yes* to format the disk.



Note that multipathing is active by clicking on the newly created SR in the *Resources* pane, and then the *General* tab.

The screenshot shows the XenCenter web interface. The left-hand navigation pane displays a tree view of the environment, including a Resource Pool named 'Haikel-Xen' with several VMs (RHEL 5, W2K3, WinServe2K3) and various storage types like CIFS ISO library, DVD drives, and 'ISCSI HBA SR on NetApp'. The main content area is titled 'ISCSI HBA SR on NetApp' and shows the 'Storage General Properties' for this storage type. The 'General' tab is active, displaying the following information:

Storage General Properties	
<b>General</b>	
Name:	ISCSI HBA SR on NetApp
Description:	Hardware HBA SR [NETAPP - /dev/sdb [sdc]]
Tags:	<None>
Folder:	<None>
Type:	Hardware HBA
Size:	4 MB used of 10 GB total (0 B allocated)
SCSI ID:	360a9800057396e38635a623439723435
UUID:	a458877b-ba48-e93f-55a3-67d8ccf6ea6d
<b>Status</b>	
State:	OK
Haikel-Xen:	Connected
<b>Multipathing</b>	
Haikel-Xen:	2 of 2 paths active

## Configuring Fibre Channel Storage

To set up the NetApp active-active controller configuration to be used for an FCP SR from XenServer, NetApp System Manager will be used. The first step is to create an Initiator Group and then assign the XenServer World Wide Port Names (WWPNs) to that group.

1. Select LUNs under the Storage option, click on the Initiator Groups tab and click on Add.




2. 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*, the operating system as *Linux* and click Add. Note: Do not use Operating system as “Xen” in FCP, see NetApp burt ID 479889.



3. Click Add on the screen below to assign the XenServer WWPNs to the newly created Initiator Group.



4. The Add Initiator ID box will appear. Select FCP for the Group Type, the Initiator Group that was created in step 2 for Group Name and the WWPN from XenServer in Initiator Name. Click Add.



The WWPN for the HBAs in the XenServer hosts can be found by running a tool such as SANsurfer. 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 screen shot below shows the WWPN for the ports 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:
  
```

5. Repeat Step 4 for all of the XenServer hosts in the pool.


After the Initiator Group has been created and all of the WWPNs added, the LUN can be created using the Create LUN Wizard. The wizard gives the ability to use previously created volumes or create both the volume and LUN in a single wizard. In the example below System Manager will create both the LUN and Volume.

Note: The snapshot reserve of the newly created Volume will be 0% and snapshots will not be scheduled. If a reserve of greater than 0% is required, right click on the Volume after its creation and select Snapshot > Configure.


- Select LUNs under the Storage option, click on the LUN Management tab and click on Create.




- Under General Properties, add the name of the LUN, which in this example will also be the name of the Volume, the description of the LUN, the size and select Linux as the type. Click Next to continue.



- Under LUN Container, choose Automatically create a new volume.



- Under Initiator Mapping, select the Initiator Group that was created in step 2 and click on the right arrow to add the Initiator Group to Hosts to connect.



- Review the selections and click Next to create the Volume and LUN.
- Click on Finish to close the Create LUN Wizard.
- 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. The 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=lvmhba 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 device, match the serial number from the `<serial>` field of the `xe sr-probe` output with the serial number of the LUN in System Manager. To determine the LUN serial number from System Manager click *LUNs* and select LUN Management. Click on the newly created LUN. The serial number is available towards the bottom of the window.



- Repeat the above steps for all hosts in the pool.
- To create the FC SR over HBA on this LUN, follow the same steps used to create the iSCSI SR over HBA.

## Fixing Misaligned Windows Guest Operating System

Disks use geometry to identify themselves and their characteristics to the upper layer operating system. The upper layer operating system uses the disk geometry information to calculate the size of the disk and partition the disk into predetermined addressable blocks. Just as with physical disks, logical disks (LUNs) report disk geometry to the host so that it can calculate space and partition the LUN into addressable blocks.

NetApp uses 4KB blocks ( $4 \times 1,024 = 4,096$  bytes) as its basic storage building block. Writes can consume no less than a single 4KB block and can consume many 4KB blocks depending on the size of the write operation. Files that are smaller than 4KB are actually stored in the inode that contains their metadata. When a LUN is created in Data ONTAP, a certain amount of space is carved out of the disk in the form of 4KB blocks. These are then reported to the host OS with a specified geometry, depending on the LUN type selected. The host OS then takes the presented LUN and partitions it appropriately.

The problem of misaligned LUN I/O occurs when the partitioning scheme used by the host OS does not match the block boundaries inside the LUN.

Note: This problem is not peculiar to NetApp. All SAN vendors have the potential to experience misalignment issues. Misalignment issues can be caused because the partition size of the blocks in the host OS use something other than 4KB blocks. Misalignment issues can be caused if the host OS imposes an offset that has not been compensated for. NetApp compensates for offsets by identifying the OS and then adjusting the offset. This is done during LUN creation when the user enters the LUN type.

For a deeper discussion of disk geometry, consider attending the NGS SAN Fundamentals class or reading [ntapcs7976: "FCP SCSI Geometry FAQs."](#)




Figure 9 Properly aligned and unaligned I/O.

When aligning the partitions of virtual disks for use with NetApp active-active controller configuration, the starting partition offset must be divisible by 4,096. The recommended starting offset value for Windows 2000, 2003, and XP operating systems is 32,768. Windows 2008 and Vista default at 1,048,576 and do not require any adjustments.

To verify this value, we need to run msinfo32.exe from the Windows command console. And you will typically find that the VM is running with a default starting offset value of 32,256. To run msinfo32, you select start > All Programs > Accessories > System Tools > System Information. Notice the partition starting offset is 32,256 bytes in Figure 10. This indicates disk misalignment.




Figure 10 Using system information to identify the starting partition offset.

## Resolution

Virtual disks can be formatted with the correct offset at the time of creation by simply booting the VM before installing an operating system and manually setting the partition offset. For Windows guest operating systems, consider using the Windows Preinstall Environment boot CD or alternative tools such as Bart's PE CD. To set up the starting offset, follow these steps:

1.	Boot the VM with the WinPE CD.
2.	Select Start > Run and enter DISKPART..
3.	Enter Select Disk0.
4.	Enter Create Partition Primary Align=32.
5.	Reboot the VM with WinPE CD.
6.	Install the operating system as normal.



Figure 11 Running diskpart to set a proper starting partition offset.

Before running *diskpart.exe* for Windows guests using the NFS SR, run the *vhd-util* to fully provision the VHD VDI, as detailed in the section describing the configuration of the NFS SR.

For more information on misalignment, read [Best Practice for File System Alignment in Virtual Environment](#).

## Backup and Recovery

There are various elements for backup and recovery:

- Snapshot the virtual machines time to time,
- Backup VM metadata
- Backup the snapshots and metadata, and
- Recover data/virtual machines from snapshots and metadata when needed.

### Snapshot for Virtual Machines on NetApp Data ONTAP and NFS Storage Repositories

Creating VM snapshots for those VMs whose VDIs are resident on the NetApp Data ONTAP SR utilizes the NetApp active-active controller configuration Snapshot technology directly by invoking the Snapshot at the device 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 VMs using the standard NFS, iSCSI, and Hardware HBA SRs, however, do not invoke any NetApp active-active controller configuration capability. It uses the VHD capability for its VDIs 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.

XenServer 5 provides a convenient snapshot mechanism that can take a snapshot of a VM's storage and metadata at a given time. Where necessary, IO is temporarily halted while the snapshot is being taken to ensure that a self consistent disk image can be captured.

Snapshot operations result in a snapshot VM that is similar to a template. The VM snapshot contains all the storage information and VM configuration, including attached VIFs, allowing them to be exported and restored for backup purposes.

The snapshot operation is a 2 step process:

- Capturing metadata as a template.
- Issuing a VM snapshot.


Two types of VM snapshots are supported: regular and quiesced.

## Regular Snapshots


Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1

Regular snapshots are crash consistent and can be performed on all VM and Storage Repository types, including Linux VMs.


1. Within XenCenter, select the VM in the left-hand pane and then the “Snapshot” tab. VMs snapshots can be created for both running and powered down VMs.



2. After clicking “Take Snapshot” and providing a name, the snapshot will appear in XenCenter



3. Snapshots can be used to restore a VM. Right click the snapshot and select “Create new VM from Snapshot.” Note that the MAC address (es) will need to be manually changed to make the new VM identical to the original. Snapshots can also be used to create templates, either for creating new VMs or for backup.



## Quiesced Snapshots

Quiesced snapshots are a special case that take advantage of the Windows Volume Snapshot Service (VSS) for services that support it, so that a supported application (for example Microsoft Exchange or SQLServer) can flush data to disk and prepare for the snapshot before it is taken.

Quiesced snapshots are therefore safer to restore, but can have a greater performance impact on a system while they are being taken. They may also fail under load so more than one attempt to take the snapshot may be required.

It is essential to install the Xen VSS provider in the Windows guest in order to support VSS. This is done via the *install-XenProvider.cmd* script provided with the Windows PV drivers.

1. In the Windows VM that needs to be snapshot in quiesced mode, open a command window and change directory to where the Citrix XenTools are installed (this is by default in %ProgramFiles%\Citrix\XenTools. Run the *install-Provider.cmd*.

```

C:\Program Files\Citrix\XenTools>install-XenProvider.cmd
The Virtual Disk Service service is not started.
More help is available by typing NET HELPMSG 3521.
The Volume Shadow Copy service is not started.
More help is available by typing NET HELPMSG 3521.
The Microsoft Software Shadow Copy Provider service is not started.
More help is available by typing NET HELPMSG 3521.
The Citrix Tools for Virtual Machines Service service is stopping.
The Citrix Tools for Virtual Machines Service service was stopped successfully.
Microsoft (R) Windows Script Host Version 5.6
Copyright (C) Microsoft Corporation 1996-2001. All rights reserved.
Microsoft (R) Windows Script Host Version 5.6
Copyright (C) Microsoft Corporation 1996-2001. All rights reserved.
The Virtual Disk Service service is starting.
The Virtual Disk Service service was started successfully.
The requested service has already been started.
More help is available by typing NET HELPMSG 2182.
The Microsoft Software Shadow Copy Provider service is starting.
The Microsoft Software Shadow Copy Provider service was started successfully.
The Citrix Tools for Virtual Machines Service service is starting.
The Citrix Tools for Virtual Machines Service service was started successfully.

C:\Program Files\Citrix\XenTools>_
    
```

2. Open a console session to the XenServer master and run the *xe vm-snapshot-with-quiesce* command to snapshot the VM.
   
`xe vm-snapshot-with-quiesce vm=<vm_name> new-name-label=<vm_snapshot_name>`

```


root@Haikel-Xen:~
new-name-description= new-name-label=
[root@Haikel-Xen ~]# xe vm-snapshot-with-quiesce vm=WinServe2K3 new-name-label=Win_2k3_quiesce_snap_Jan2011_5edb3ee4-9d9f-6f81-59b4-f81303533bea
    
```

- Once a quiesced snapshot has been created, it will appear in the VM’s “snapshot” tab within XenCenter.


## Backing up VM Metadata Information on NetApp active-active controller configuration

Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1


- Open an *xsconsole* session to the XenServer master from XenCenter (if in a shell prompt, type *xsconsole* and press *<Enter>*).



- Scroll to the *Backup, Restore and Update* option and hit *Enter*. Choose the *Backup Virtual Machine Metadata* option and press *<Enter>*.




3. Enter the login credentials for the XenServer host if prompted.
  
4. Select the NetApp SR, and press *<Enter>*.



There will be a short delay while the VM metadata is backed up to the NetApp active-active controller configuration

5. A success dialog box shows up indicating that the backup of the metadata was successful



## Backing up Storage Repositories on NetApp active-active controller configuration

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>

Some considerations when using SnapMirror:

1. For SnapMirror volume replication, the destination storage system must use a version of Data ONTAP that is the same as or later than that of the SnapMirror source storage system. For volume SnapMirror to support replication for the purpose of disaster recovery, both the source and destination storage systems must use the same version of Data ONTAP.
2. For SnapMirror volume replication, the capacity of the destination volume must be greater than or equal to the capacity of the source volume
3. To optimize performance, stagger Snapshot copy update schedules so that SnapMirror activity does not begin or end at the exact minute a *snap sched* command operation attempts to create a Snapshot copy. If the SnapMirror feature is scheduled to perform Snapshot management at the same time as a *snap sched* activity, then the Snapshot management operations scheduled using the *snap sched* command might fail with syslog messages: *Skipping creation of hourly snapshot* and *Snapshot already exists*.
4. For optimum SnapMirror volume replication performance, ensure that the SnapMirror source volume and destination volume contain disks of the same size.

- To achieve maximum space savings on the destination volume, scan the entire file system to recreate the deduplication metadata for the destination volume. Use the `sis start -s` command to do so. **Note:** The destination volume is accessible for read-write operations when the deduplication scan is in progress. The `sis start` command, without the `-s` option, has the potential for space savings on the destination volume is reduced because only the new data written to the volume will be scanned for deduplication.

Note: Starting with Data ONTAP 8.0.1 7-Mode, FlexClone volumes can be replicated using volume SnapMirror without the need for additional capacity on the destination system as long as the parent of the FlexClone volume is also replicated. Please consult TR-3446: SnapMirror Async Overview and Best Practices Guide available at <http://www.netapp.com>.

- For Firewall setting, allow a range of TCP ports from 10565 to 10569.

It is recommended to use the below diagram configuration as a guide to setup replication between the primary and DR site.




Figure 12 Graphical representation of replication setup

## Creating the SnapMirror Relationship

(For purposes of illustration, the NetApp SR type is used for the backup and recovery steps. The same steps can be used for other SR types given the FlexVol)

1. Determine the SR UUID from XenServer host console (as indicated in previous sections). For NFS/iSCSI/LVMoHBA SR on the NetApp active-active controller configuration, the FlexVol is the name determined when creating the FlexVol from the FilerView FlexVol creation wizard.
2. Open a console session to the NetApp active-active controller configuration on the primary site. Convert the FlexVols in the NetApp SR to use ucode using the `vol options` command

```
vol options <vol-name> <option-name> <option-val>
```

```
ISV3140> vol options XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV0 create_ucode on
ISV3140> vol options XenStorage_78ba35b0_1aff_4867_90b3_19d7ec223f05_FV0 convert_ucode on
```

3. Repeat the above step for all FlexVols that make up the NetApp SR
4. Ensure SnapMirror is turned on using the `options snapmirror.enable` command

```
ISV3140> options snapmirror.enable
snapmirror.enable      on
```

If not enabled, turn snapmirror on using the same command with `on` option


```
options snapmirror.enable on
```

5. Authorize DR site host to replicate the primary site storage using the `options snapmirror.access` command

```
options snapmirror.access host=<IP address or DNS name of DR site device>
```



```
ISV3140> options snapmirror.access host=10.204.132.16
```



- On the primary site NetApp SR, disable any system Snapshots that may be on by default. In FilerView, click *Volumes->Snapshots->Configure* and select the FlexVols that make up the NetApp SR. Click *Unselect All*, and set the *Number of Scheduled snapshots to Keep* to 0. Ensure that *Scheduled Snapshots* is unselected as well. Click *Apply*



- Open the NetApp FilerView for the NetApp active-active controller configuration at the DR site, and create an aggregate that is larger than the aggregate on the primary site that has the NetApp SR FlexVols. The aggregate can be created using steps outlined in the NetApp active-active controller configuration configuration section.

8. Create FlexVols within the aggregate created in the previous step, such that the FlexVols have the same name as the FlexVols that make up the NetApp SR at the primary site. FlexVols can be created using the steps outlined earlier in the document. Ensure that the space guarantee for the FlexVol is set to *none* and the size is at least 5% greater than the FlexVol in the NetApp SR at the primary site. Also set the *Snapshot Reserve* to 0%.






9. Set the ucode options on the FlexVols.
10. Turn off Snapshot on the FlexVols using the `vol options` command.
 


```
vol options <flexvol name> nosnap on
```
11. If Data Deduplication is turned on at the primary site, turn on dedupe on the DR site FlexVols as well using the `sis on` command.
 

```
sis on <path>
```


12. Restrict the FlexVols by clicking on *Volumes->Manage*, selecting the FlexVols and then clicking the *Restrict* button.




13. Give access to the primary site NetApp active-active controller configuration to the DR site NetApp active-active controller configuration and vice-versa. From FilerView, click *SnapMirror->Remote Access->Add* and add the DNS name or IP address of the filer will request access.



14. Set up SnapMirror from FilerView by clicking *SnapMirror->Add*.
15. Select the first FlexVol and click *Next*.



16. Enter the primary site NetApp active-active controller configuration name or IP address and FlexVol for the first FlexVol that needs to be replicated.



17. Leave the *Restart Mode* and *Maximum Transfer Rate* as the default. The data transfer limit may be changed depending on the network and storage configuration at the deployment site.
  - a. Calculate a TCP window size that works well with the network, using the following formula:

Window Size = (Round Trip Delay) × (Desired Rate)


b. Adjust the TCP window size by entering the following command

```
options snapmirror.window_size rate
```

*rate* = the desired TCP window size.


18. Set the schedule for SnapMirror to occur as determined by the business needs. Having a small window will affect disk I/O on both the primary and DR sites, depending on the rate of change of the data blocks.

19. Commit the changes to create the new SnapMirror.




20. Click on *SnapMirror->Manage* and click *Advanced* for the SnapMirror just created.

21. Click on *Initialize* to setup the initial base line transfer.



22. Once all FlexVols have been initialized, the SnapMirror *State* should show *snamirrored*. To verify, in FilerView, click on *SnapMirror->Manage*.



## Configuring XenServer at DR site at time of disaster

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

1. Quiesce the SnapMirror relationship.
2. Break the SnapMirror relationship.
3. Once the SnapMirror relationship has been broken, the mirrored FlexVol is attached to the backup XenServer host at the DR site.

### Quiesce the SnapMirror relationship


1. In FilerView, click on *SnapMirror->Manage*. Click *Advanced* for the FlexVol that needs to be quiesced.



2. Click *Quiesce* and then *OK* in the confirmation dialog box.

## Break the SnapMirror Relationship

1. Once the mirror has quiesced, click *Break* and *OK* in the confirmation dialog box. The state of the mirror should show *broken-off*.




Repeat the quiesce and break-off process for each FlexVol that makes up the mirrored NetApp SR..


### Attaching the mirror'ed SR to XenServer host at DR site

Please note that some images appear slightly different in XenServer 5.6 than XenServer 5.6 Feature Pack1


1. From XenCenter, connect to the XenServer host at the DR site..
2. Click on *New Storage*, and choose the *Advanced StorageLink technology*> *NetApp* option.
3. Enter the name/IP address of the NetApp active-active controller configuration at the DR site that has the mirror'ed FlexVols that make up the NetApp SR.
4. Click on *Reattach on the existing SRs on the filer* and choose the SR that was mirror'ed. Click *Finish*.




5. Click *Yes* to confirm attaching the SR..
6. Open the *xsconsole* to the XenServer host at the DR site. Scroll to *Backup, Restore and Update* and press <Enter>.




7. Scroll to *Restore Virtual Machine Metadata* and press <Enter>.



8. If prompted, provide XenServer login credentials to perform the metadata restoration.
9. Select the attached NetApp SR from the previous steps and press <Enter>. There is a slight delay while XenServer looks for the backup VDI.
10. Once the backup VDI is found, select the appropriate metadata to restore from and press <Enter>.



11. Select the *Only VMs on this SR* option to import all the VM metadata into the XenServer host at the DR site.



There is a slight delay while all metadata is imported.

### Restoring SnapMirror relationship to bring primary site back up

Once the primary site has recovered from the disaster, use the below steps to restore information back to primary site.

To restore the SnapMirror relationship:

1. On primary site, destroy all data by re-initializing the primary NetApp active-active controller configuration

2. Set up a SnapMirror from the DR NetApp active-active controller configuration to the primary site NetApp active-active controller configuration:
  - a. Initialize a base line transfer
  - b. Quiesce the SnapMirror relationship
  - c. Break the SnapMirror relationship
  - d. Delete the SnapMirror configuration

This restores all the updated data from the DR site back to primary site.

3. Delete the SnapMirror configuration on the DR site NetApp active-active controller configuration
4. Repeat the steps in creating the SnapMirror relationship to set up replication between the primary site and DR site.

## 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 and Recovery of VMs


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

Usage:

Copy/paste the script below to a file with .py extension. Copy the .py file to your XenServer host, and execute it.

Sample run output:

```
python snapshot_restore.py
```



```

root@Haikel-Xen:~
[root@Haikel-Xen ~]# python snapshot_restore.py

Citrix XenServer VM Snapshot/Restore Utility
Copyright (c) 2008 Citrix Systems, Inc.

VMs available:

1. W2K3 (1)
2. RHEL 5 (64-bit) (Dev)
3. WinServe2K3 (64-bit) (LicenseServer)

Please choose VM to snapshot or recover [1 - 3]: 1
Do you want to snapshot|snapshot with quiesce (Windows guests only)|recover?
(input "s" or "sq" or "r")

```

```

#!/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 snapshot_vm(session, vm_rec, snap_choice):
    epoch = time.strftime("%Y%m%d-%H%M%S", time.localtime())
    vm_snap_name = "%s - %s" % (vm_rec['name_label'], epoch)
    print "\nNow creating snapshot for VM %s..." % vm_rec['name_label']
    if snap_choice == 'sq':
        vm_snap = session.xenapi.VM.snapshot_with_quiesce(\
            session.xenapi.VM.get_by_uuid(vm_rec['uuid']), \
            vm_snap_name)
    else:

```

```

vm_snap = session.xenapi.VM.snapshot(\
    session.xenapi.VM.get_by_uuid(vm_rec['uuid']), \
    vm_snap_name)
print "VM snapshot created as : %s" % vm_snap_name

def recover_vm(session, vm_rec):
    snaps = vm_rec['snapshots']
    if len(snaps) == 0:
        print "No snapshots available for VM %s" % vm_rec['name_label']
    else:
        snap_num = 0
        snap_dict = {}
        print "List of snapshots available for VM %s:\n" % vm_rec['name_label']
        for snap in snaps:
            snap_rec = session.xenapi.VM.get_record(snap)
            snap_num += 1
            snap_dict[snap_num] = snap
            print "\t%i. %s" % (snap_num, snap_rec['name_label'])
        if (snap_num == 1):
            snap_num_choice = 1
            raw_input("\nPress Enter to restore from this snapshot")
        else:
            snap_num_choice = 0
            while (snap_num_choice == 0):
                snap_num_choice_s = raw_input ('\nPlease choose a snapshot to
Restore to [1 - %i]: ' % snap_num)
                if (snap_num_choice_s == ""):
                    print "\nInvalid entry. Please try again...\n"
                    continue
                snap_num_choice = int(snap_num_choice_s)
                if (snap_num_choice > snap_num):
                    print ("\nInvalid entry. Please try again...\n")
                    snap_num_choice = 0
                    continue
            snap_choice = snap_dict[snap_num_choice]
            snap_choice_rec = session.xenapi.VM.get_record(snap_choice)

            # Steps to restore
            # 1. Force shutdown VM
            # 2. Destroy VBDs
            # 3. Clone VDI's of snapshot
            # 4. Sync state (VBD, VIF, VCPU, Memory, etc)
            # 5. Destroy subsequent snapshots

            print "\nCleanup VM..."
            vm_name = vm_rec['name_label']
            vm_power_state = vm_rec['power_state']
            if not(vm_power_state == 'Halted'):
                print "Halting VM..."

            session.xenapi.VM.hard_shutdown(session.xenapi.VM.get_by_uuid(vm_rec['uuid']))

            vbds = vm_rec['VBDs']
            for vbd in vbds:
                vbd_rec = session.xenapi.VBD.get_record(vbd)
                if vbd_rec['type'] == 'Disk':
                    vdi = vbd_rec['VDI']
                    session.xenapi.VDI.destroy(vdi)
            vifs = vm_rec['VIFs']
            for vif in vifs:
                session.xenapi.VIF.destroy(vif)

```

```

print "Restoring VM..."

# Restore VIFs
snap_vifs = snap_choice_rec['VIFs']
for snap_vif in snap_vifs:
    snap_vif_rec = session.xenapi.VIF.get_record(snap_vif)
    restore_vif = { 'device' : snap_vif_rec['device'],
                    'network' : snap_vif_rec['network'],
                    'VM' : session.xenapi.VM.get_by_uuid(vm_rec['uuid']),
                    'MAC' : snap_vif_rec['MAC'],
                    'MTU' : snap_vif_rec['MTU'],
                    "other_config" : snap_vif_rec['other_config'],
                    "qos_algorithm_type" : snap_vif_rec['qos_algorithm_type'],
                    "qos_algorithm_params" :
snap_vif_rec['qos_algorithm_params'] }
    session.xenapi.VIF.create(restore_vif)

# Restore VBDs
snap_vbds = snap_choice_rec['VBDs']
for snap_vbd in snap_vbds:
    snap_vbd_rec = session.xenapi.VBD.get_record(snap_vbd)
    if snap_vbd_rec['type'] == 'Disk':
        snap_vdi = snap_vbd_rec['VDI']
        restore_vdi = session.xenapi.VDI.clone(snap_vdi, {})
        restore_vbd = { 'VM' :
session.xenapi.VM.get_by_uuid(vm_rec['uuid']),
                        'VDI' : restore_vdi,
                        'userdevice' : snap_vbd_rec['userdevice'],
                        'bootable' : snap_vbd_rec['bootable'],
                        'mode' : snap_vbd_rec['mode'],
                        'type' : snap_vbd_rec['type'],
                        'unpluggable' : snap_vbd_rec['unpluggable'],
                        'empty' : snap_vbd_rec['empty'],
                        'other_config' : snap_vbd_rec['other_config'],
                        'qos_algorithm_type' :
snap_vbd_rec['qos_algorithm_type'],
                        'qos_algorithm_params' :
snap_vbd_rec['qos_algorithm_params'] }
        session.xenapi.VBD.create(restore_vbd)

# Restore config information
vm_rec['platform'] = snap_choice_rec['platform']
vm_rec['VCPUs_max'] = snap_choice_rec['VCPUs_max']
vm_rec['recommendations'] = snap_choice_rec['recommendations']
vm_rec['user_version'] = snap_choice_rec['user_version']
vm_rec['HVM_shadow_multiplier'] = snap_choice_rec['HVM_shadow_multiplier']
vm_rec['memory_dynamic_max'] = snap_choice_rec['memory_dynamic_max']
vm_rec['HVM_boot_policy'] = snap_choice_rec['HVM_boot_policy']
vm_rec['PV_legacy_args'] = snap_choice_rec['PV_legacy_args']
vm_rec['actions_after_shutdown'] = snap_choice_rec['actions_after_shutdown']
vm_rec['xenstore_data'] = snap_choice_rec['xenstore_data']
vm_rec['actions_after_crash'] = snap_choice_rec['actions_after_crash']
vm_rec['PCI_bus'] = snap_choice_rec['PCI_bus']
vm_rec['tags'] = snap_choice_rec['tags']
vm_rec['other_config'] = snap_choice_rec['other_config']
vm_rec['VCPUs_at_startup'] = snap_choice_rec['VCPUs_at_startup']
vm_rec['actions_after_reboot'] = snap_choice_rec['actions_after_reboot']
vm_rec['PV_bootloader'] = snap_choice_rec['PV_bootloader']
vm_rec['resident_on'] = snap_choice_rec['resident_on']
vm_rec['HVM_boot_params'] = snap_choice_rec['HVM_boot_params']
vm_rec['blocked_operations'] = snap_choice_rec['blocked_operations']
vm_rec['PV_args'] = snap_choice_rec['PV_args']
    
```

```

vm_rec['affinity'] = snap_choice_rec['affinity']
vm_rec['PV_ramdisk'] = snap_choice_rec['PV_ramdisk']
vm_rec['memory_static_min'] = snap_choice_rec['memory_static_min']
vm_rec['PV_bootloader_args'] = snap_choice_rec['PV_bootloader_args']
vm_rec['PV_kernel'] = snap_choice_rec['PV_kernel']
vm_rec['memory_dynamic_min'] = snap_choice_rec['memory_dynamic_min']

if not(vm_power_state == 'Halted'):
    print "Powering up VM..."
    session.xenapi.VM.start(session.xenapi.VM.get_by_uuid(vm_rec['uuid']),
False, False)

print "Destroying orphaned snapshots..."
for snap in snaps:
    snap_rec = session.xenapi.VM.get_record(snap)
    if snap_rec['snapshot_time'] > snap_choice_rec['snapshot_time']:
        vbds = snap_rec['VBDs']
        for vbd in vbds:
            vbd_rec = session.xenapi.VBD.get_record(vbd)
            if vbd_rec['type'] == 'Disk':
                vdi = vbd_rec['VDI']
                session.xenapi.VDI.destroy(vdi)
            session.xenapi.VM.destroy(snap)

def snapshot_recover_vm(session, vm_rec):
    while (True):
        choice = raw_input("Do you want to snapshot|snapshot with quiesce (Windows
guests only)|recover?\n(input \"s\" or \"sq\" or \"r\") ")
        if (choice == 's') or (choice == 'sq'):
            snapshot_vm(session, vm_rec, choice)
            break
        elif (choice == 'r'):
            recover_vm(session, vm_rec)
            break
        else:
            print "\nInvalid choice. Please try again...\n"
            continue

def main(session):
    vm_records = session.xenapi.VM.get_all_records()
    vm_num = 0
    vm_dict = {}
    print "VMs available:\n"
    for vm_record in vm_records:
        vm_rec = vm_records[vm_record]
        if not(vm_rec['is_a_template']) and not(vm_rec['is_control_domain']):
            vm_num += 1
            vm_dict[vm_num] = vm_rec
            print "\t%i. %s" % (vm_num, vm_rec['name_label'])
    print ""
    if (vm_num == 0):
        print ("No VMs available for snapshot/recover")
        return
    elif (vm_num == 1):
        vm_num_choice = 1
    else:
        vm_num_choice = 0
        while (vm_num_choice == 0):
            vm_num_choice_s = raw_input('Please choose VM to snapshot or recover [1 -
%i]: ' % vm_num)
            if (vm_num_choice_s == ""):
                print "\nInvalid entry. Please try again...\n"

```

```

        continue
        vm_num_choice = int(vm_num_choice_s)
    if (vm_num_choice > vm_num):
        print "\nInvalid entry. Please try again...\n"
        vm_num_choice = 0
        continue

    vm_rec_choice = vm_dict[vm_num_choice]

    snapshot_recover_vm(session, vm_rec_choice)

if __name__ == "__main__":
    print "\n\tCitrix XenServer VM Snapshot/Restore Utility"
    print "\tCopyright (c) 2008 Citrix Systems, Inc.\n"

    session = XenAPI.Session("http://_var_xapi_xapi", transport=XenAPI.USTransport())
    session.xenapi.login_with_password("root", "xensource")

    try:
        main(session)
    except Exception, e:
        print str(e)
        raw_input ("\n\n--Press Enter to exit--")
        raise
    raw_input ("\n\n--Press Enter to exit--")

```

## Appendix B

### Security Considerations

To allow XenServer administrators to login into the device 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 device enter:

```
device> options wafl.nt_admin_priv_map_to_root on
```

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

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.

## Appendix C

### Enhancements in Data ONTAP

Enhancements in Data ONTAP 7.3.1 affect the ability to do a LUN clone split when using NetApp's ZAPI. ONTAP has two ways (styles) to do the LUN clone split. Beginning with Data ONTAP 7.3.1, it will default to the new style. Citrix XenServer's NetApp Adapter will attempt to use the old style LUN clone split.

You will see the following error message if you are affected by this issue:

```
“Error parameters: General IO error [opterr=creating flexvol snapshot:Snapshot operation not allowed during LUN clone split. (13023)]”
```

If you are affected, the issue can be verified and fixed with the following commands.

```
> options lun.clone_restore  
> options lun.clone_restore {off | on}
```

Setting to 'off' will mimic the old behavior; the system defaults to 'on'.

Data ONTAP 8.0 7-Mode introduces 64-bit aggregates which are no longer limited to 16TB. Consult TR-3786: A Thorough Introduction to 64-Bit Aggregates for complete details and sizing information, available at <http://www.netapp.com>.

Data ONTAP 8.0.1 7-Mode FlexClone volumes can be replicated using volume SnapMirror without the need for additional capacity on the destination system as long as the parent of the FlexClone volume is also replicated. Please consult TR-3446: SnapMirror Async Overview and Best Practices Guide available at <http://www.netapp.com>.

## References

XenServer 5.6 Administration Guide: <http://support.citrix.com/article/CTX124887>

XenServer 5.6 Virtual Machine Installation Guide:  
<http://support.citrix.com/article/CTX124888>

XenServer 5.6 Feature Pack 1 Administration Guide:  
<http://support.citrix.com/article/CTX127321>

XenServer 5.6 Feature Pack 1 Virtual Machine Installation Guide:  
<http://support.citrix.com/article/CTX127323>

Citrix StorageLink 2.3 User Guide: <http://support.citrix.com/article/CTX126304>

Data ONTAP® 7.3 System Administration Guide:  
<http://now.netapp.com/NOW/knowledge/docs/ontap/rel73/pdfs/ontap/sysadmin.pdf>

Data ONTAP® 7.3 Data Protection Online Backup and Recovery Guide:  
<http://now.netapp.com/NOW/knowledge/docs/ontap/rel73/pdfs/ontap/onlinebk.pdf>

## Citrix Worldwide

Worldwide headquarters  
Citrix Systems, Inc.  
851 West Cypress Creek Road  
Fort Lauderdale, FL 33309  
USA  
T +1 800 393 1888  
T +1 954 267 3000

### Regional headquarters

Americas  
Citrix Silicon Valley  
4988 Great America Parkway  
Santa Clara, CA 95054  
USA  
T +1 408 790 8000

Europe  
Citrix Systems International GmbH  
Rheinweg 9  
8200 Schaffhausen  
Switzerland  
T +41 52 635 7700

Asia Pacific  
Citrix Systems Hong Kong Ltd.  
Suite 3201, 32nd Floor  
One International Finance Centre  
1 Harbour View Street  
Central  
Hong Kong  
T +852 2100 5000

Citrix Online division  
5385 Hollister Avenue  
Santa Barbara, CA 93111  
USA  
T +1 805 690 6400

[www.citrix.com](http://www.citrix.com)

### About Citrix Systems, Inc.

Citrix Systems, Inc. (NASDAQ: CTXS) is the global leader in access infrastructure and the most trusted name in secure access. More than 200,000 organizations around the world use the Citrix Access Platform to provide secure, well-managed access to business information wherever it lives —on demand. Citrix customers include 100 percent of the Fortune 100 companies, 99 percent of the Fortune 500, and 97 percent of the Fortune Global 500, in addition to hundreds of thousands of smaller businesses and individuals. Based in Fort Lauderdale, Florida, Citrix has offices in 22 countries and approximately 6,200 channel and alliance partners in more than 100 countries. Annual revenue in 2007 was \$1.4 billion.

### About NetApp, Inc.

NetApp creates innovative storage and data management solutions that accelerate business breakthroughs and deliver outstanding cost efficiency. Discover our passion for helping companies around the world go further, faster at [www.netapp.com](http://www.netapp.com)

