Whitepaper

Deploying Citrix XenServer 5.0 with Dell EqualLogic PS Series storage

Table of contents

Dell EqualLogic PS Series Storage 4

XenServer 5.0 Dell Edition Overview 5
What's new in XenServer 5.0 Dell Edition? 5
Expanded hardware and systems management support 5
Faster and more powerful 5
Simplified management 6
Disaster proof 6
The only dynamic provisioning game in town 6
EqualLogic PS Series and XenServer 7
Optimized Utilization 7
Reduced Deployment Time and Effort 7
Reduced Management Complexity 7
High Availability 8
High Scalability 8
Increased Flexibility 8
XenServer Storage Model and Definitions 9
Dell EqualLogic PS Series Feature Definitions 11
Volume 11
Thin Provisioning 11
Volume Snapshots 11
Volume Clones 12
Volume Replicas 12
iSCSI Storage Repositories in XenServer 5.0 13
LVM over iSCSI vs. Dell EqualLogic SR usage considerations 14
Setup and configuration 15
Typical PS Series-XenServer Setup 15
Multi-pathing with XenServer and PS Series 16
Sample XenServer host and PS Series network configuration 16
Networking Best Practices for PS Series and XenServer 17
Creating a Storage Repository on PS Series Group 18
Creating a storage repository of type LVM over iSCSI (lvmoiscsi) 18
Creating a storage repository of type Dell EqualLogic (equal) 19
Identifying XenServer storage objects on a PS Series Group 21

Typical lifecycle tasks for VMs on Dell EqualLogic SR.....22
Create a Virtual Machine 22
Create a VDI 22
Delete a VDI.....23
Resize a VDI 23
Copy or Clone a VM.....24
Deploy from Template 24
Snapshot a VM 25
Restore a VM from a Snapshot 26
Delete a VM 26
Delete a Snapshot 26
Using PS Series Replication with XenServer Portable SRs 27
Using replication for SR of type LVM over iSCSI 27
Using replication for SR of type Dell EqualLogic 28
References 29

Dell EqualLogic PS Series storage

Dell® EqualLogic™ PS Series iSCSI arrays simplify storage deployment by offering high performance, reliability, intelligent automation, and seamless virtualization of a single pool of storage. The foundation of an EqualLogic storage solution is a PS Series Group: an iSCSI SAN including one or more PS Series storage arrays (members) connected to an IP network and managed as a single system. Each array has fully redundant hardware and up to three active network connections for maximum bandwidth. Integrated virtualization firmware provides:

- Seamless scalability
- Automatic RAID configuration and spare disk configuration
- Automatic network, performance, and capacity load balancing

The Dell EqualLogic PS5000E line spans from 2 to 16 TB in capacity in a single array, utilizing SATA drives ranging from 250GB to 1TB. The PS5000X arrays offer 6.4TB or raw storage capacity utilizing 400GB, 10,000 RPM Serial Attached SCSI (SAS) disk drives. The PS5000XV arrays use 146GB or 300GB 15,000 RPM Serial Attached SCSI (SAS) disk drives to deliver 2.3TB or 4.8TB of raw capacity, respectively. The PS5500E array utilizes 48 7200 RPM 500GB or 1TB SATAII drives to deliver 24 or 48 TB of raw capacity in a 4U enclosure.

Performance and capacity may easily be scaled by adding additional PS Series arrays. All PS Series models have dual controllers with three 1GbE ports per controller, for a total of six 1GbE ports.

XenServer 5.0 Dell Edition overview

With the 64-bit open-source Xen hypervisor at its core, Citrix® XenServer™ Dell Edition is a powerful virtualization solution that enables efficient resource consolidation, utilization, dynamic provisioning, high availability and integrated systems management. XenServer Dell Edition has a small foot print and is optimized to run from an internal flash storage in Dell PowerEdge™ servers. Dell and Citrix have partnered to bring pre-qualified and virtualization-ready platforms for today's dynamic and growing data centers.

What's new in XenServer 5.0 Dell Edition?

Expanded hardware and systems management support

Expanded Server Support—XenServer 5.0 is now supported on an expanded list of Dell PowerEdge servers. PowerEdge 1950 III, 2900III, 2950III, R805, R900, R905, M805 and M905 are now fully qualified and supported to run XenServer 5.0.

Expanded Storage Support—Storage arrays supported by Dell with XenServer 5.0 include Dell PowerVault™ MD1000, MD1120, MD3000, MD3000i and Dell EqualLogic PS5000 arrays. XenServer 5.0 also includes a new Dell EqualLogic PS Series Storage Adapter that not only simplifies virtual machine storage management, but also enables advanced PS Series features to be used by XenServer hosts.

OpenManage 5.5—With Dell OpenManage™ 5.5 integrated in XenServer 5.0, customers can easily monitor and manage server hardware and direct attached storage from a simple to use web user interface or CLI. Dell Remote Access Controller (DRAC) 5 provides out of band, including full KVM and virtual media, access to the server.

Faster and more powerful

Powered by Xen—Boasting faster out-of-the-box performance, Citrix XenServer 5.0 gets more virtual machines, with more users and faster applications onto Dell's 64-bit PowerEdge servers. The Xen® 3.2 hypervisor includes control domain performance tuning, updated paravirtualization drivers and VHD performance tuning.

Tuned for Windows—XenServer adds support for Windows Server 2008, has been optimized to run Windows-based workloads and has been certified on both 32-bit and 64-bit versions of Windows Server through Microsoft's Server Virtualization Validation Program (SVVP).

Tuned for XenApp—Citrix XenApp™ users benefit from XenServer's memory management performance enhancements for XenApp workloads. More users and faster applications lead to more efficient application delivery with XenApp.

Simplified management

Usability enhancements—XenServer 5.0 includes 30 new ease-of-use enhancements to make it simpler for customers to use the XenServer platform.

Search, sort and tag—The XenCenter™ management console allows administrators to tag and search virtual infrastructure with ease to keep track of virtual machines as they proliferate across large datacenters. Web 2.0 style tagging and searching capabilities allow IT professionals to assign metadata and virtual tags to workloads, either pre-defined or customized to each organization's needs.

Performance monitoring and alerts—XenCenter includes performance monitoring, trending and alerting. Easily track key performance statistics for virtual and physical servers to identify hot spots, balance computing resources and alert administrators about potentially detrimental server behavior or changes.

Wizard driven—XenServer includes new Wizard-driven utilities such as XenConvert which transforms physical servers into virtual machines in minutes. The XenCenter management console also includes new wizards for advanced storage and server networking configurations.

Disaster proof

Automated high availability—XenServer adds automated high availability (HA) with resource-based placement of virtual machines in the event of server failures. Automated HA includes dynamic fail-over planning based on available resources to help ensure that virtual machines always restart on the appropriate physical server.

XenCenter self-healing—XenCenter combines new automated HA with its distributed management architecture to create a self-healing, hard to kill management console. Every server in a resource pool can be a fail-over target for the XenCenter management console.

Disaster recovery—XenServer includes enhanced support for SAN based remote replication and auto-backup of virtual machine configuration to remote sites.

The only dynamic provisioning game in town

Windows Server 2008 and Hyper-V support—XenServer now includes the ability to stream Windows Server 2008 workloads to physical servers or even Hyper-V (and associated workloads) to bare metal servers for dynamic provisioning of server workloads.

Additional Microsoft integration—Role based administration using Active Directory, virtual hard disk (VHD) compatibility and out-of-the-box integration with Microsoft SQL Server are just a few of the many Microsoft related enhancements in XenServer's dynamic provisioning capabilities.

EqualLogic PS Series and XenServer

Dell EqualLogic PS Series iSCSI arrays simplify storage deployment by offering high performance, reliability, intelligent automation, and seamless virtualization of a single pool of storage. A PS Series SAN provides intelligent automation of storage management as well as virtualized storage assets. With single console management and ease of storage provisioning, an EqualLogic PS Series SAN increases the power and flexibility of a virtual infrastructure. The arrays present a single virtualized pool of storage resources to attached servers. Additional arrays can be seamlessly added into an existing SAN to automatically increase storage pool resources without disruption of application or data availability.

Dell EqualLogic storage brings unique benefits to the Citrix XenServer virtualization solution.

Optimized utilization

While XenServer provides an excellent virtualization platform to consolidate and optimize server resource usage, EqualLogic PS Series storage arrays optimize storage utilization by dynamically balancing loads among multiple storage arrays as the usage changes. Additionally one can configure thin-provisioned volumes to grow on-demand only when additional storage is needed for those volumes. Thin-provisioning can increase the efficiency with which the storage resources are utilized.

The XenServer EqualLogic Storage Adapter allows thin volumes to be created from the XenCenter interface or using XenServer CLI, thus optimizing storage requirement and simplifying storage deployment for virtual machines.

Reduced deployment time and effort

EqualLogic PS Series arrays reduce the time and effort required to deploy a SAN. Inside the EqualLogic PS Series Group Manager, one can easily create new volumes and assign them to XenServer hosts, which can immediately access those volumes as storage repositories. Using the XenServer EqualLogic Storage Adapter, a volume on the PS Series Group is automatically created as a virtual hard disk for a virtual machine and access controls are automatically set for the volume. Using the volume cloning feature of PS series, new virtual machines can be instantly provisioned by cloning an existing virtual machine.

Reduced management complexity

EqualLogic PS Series SANs simplify storage management by consolidating physical storage and providing a single-pane management view of the entire virtualized storage pool in the Group. The PS Series intelligently balances workloads across the available arrays without need for human intervention, automatically adapting to changes in workload.

The XenServer EqualLogic Storage Adapter integrates the PS series storage management directly into the XenServer and XenCenter interface. Storage volumes for virtual machines can be managed right from XenCenter or XenServer CLI. Thus virtual machine storage deployment and management are greatly simplified when using EqualLogic storage with XenServer.

High availability

XenServer 5.0 introduces an automated local High-Availability feature that protects virtual machines hosted on shared storage against host failures. The virtual machines running on a failed host are restarted on remaining servers in the XenServer pool. NIC bonding available in XenServer protects against failures in either network cards or paths for storage, virtual machine or management traffic. EqualLogic PS Series arrays offer best in class reliability, with hot-swappable and fully redundant components, RAID 5/10/50 support and hot spares.

High scalability

As one adds more PS Series array members to a PS Series group, the storage capacity scales along with the performance without disrupting application or data availability. Plus the SAN dynamically balances the load to optimize storage resources.

Using XenServer Resource Pools, once can easily add new XenServer hosts to an existing pool, increasing the pool's capacity to meet increased workload demands.

Increased flexibility

The net result of all these features and capabilities is not only a reduced TCO, but increased flexibility. The ease and speed with which you can add new virtual servers or expand physical storage, combined with single-pane management and the security of a high-availability infrastructure, allows IT administrators to respond quickly and flexibly to enterprise demands and initiatives.

SAN features like tiering, load balancing, advanced features in replication and snapshots provide a higher level of functionality that you can't find in standard stand-alone direct attached storage.

XenServer storage model and definitions

XenServer host defines a container called a Storage Repository (SR) to describe a particular storage target, in which Virtual Disk Images (VDIs) are stored. A VDI is a disk abstraction which contains the contents of a virtual disk. Figure 1 illustrates the overall storage model in XenServer.

The interface to storage hardware allows VDIs to be supported on a large number of SR types. With built-in support for IDE, SATA, SCSI and SAS drives connected locally, and iSCSI, NFS and Fibre Channel connected remotely, the XenServer host SR is very flexible. The SR and VDI abstractions allows advanced storage features such as sparse provisioning, VDI snapshots, and fast cloning to be exposed on storage targets, such as Dell EqualLogic PS Series, that support them. For storage subsystems that do not inherently support advanced operations directly, a software stack is provided based on Microsoft's Virtual Hard Disk (VHD) specification which implements these features.

Each XenServer host can use multiple SRs and different SR types simultaneously. These SRs can be shared, or dedicated between hosts. For shared SR, shared storage is pooled between multiple hosts within a defined resource pool. A shared SR must be accessible to each host, and thus must be on iSCSI, NFS, Fibre Channel or Direct Attached for clusters (such as Dell PowerVault MD3000) storage. Finally, all hosts in a single resource pool must have at least one shared SR in common.

The following four object classes that are used to describe, configure, and manage storage:

Storage Repositories (SRs) are storage targets containing homogeneous virtual disks (VDIs). SR commands provide operations for creating, destroying, resizing, cloning, connecting and discovering the individual Virtual Disk Images (VDIs) that they contain. A storage repository is a persistent, on-disk data structure. So the act of "creating" a new SR is similar to that of formatting a disk—for most SR types, creating a new SR involves erasing any existing data on the specified storage target. The exception is NFS SRs, which create a new directory on the filer leaving existing NFS SRs as they were. SRs are long-lived, and may in some cases be shared among XenServer hosts, or moved between them.

Physical Block Devices (PBDs) represent the interface between a physical server and an attached SR. PBDs are connector objects that allow a given SR to be mapped to a XenServer host. PBDs store the device configuration fields that are used to connect to and interact with a given storage target. For example, NFS device configuration includes the IP address of the NFS server and the associated path that the XenServer host mounts. PBD objects manage the run-time attachment of a given SR to a given XenServer host.

Virtual Disk Images (VDIs) are on-disk representation of virtual disks provided to VMs. 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 connector objects (similar to the PBD described above) that allow 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.

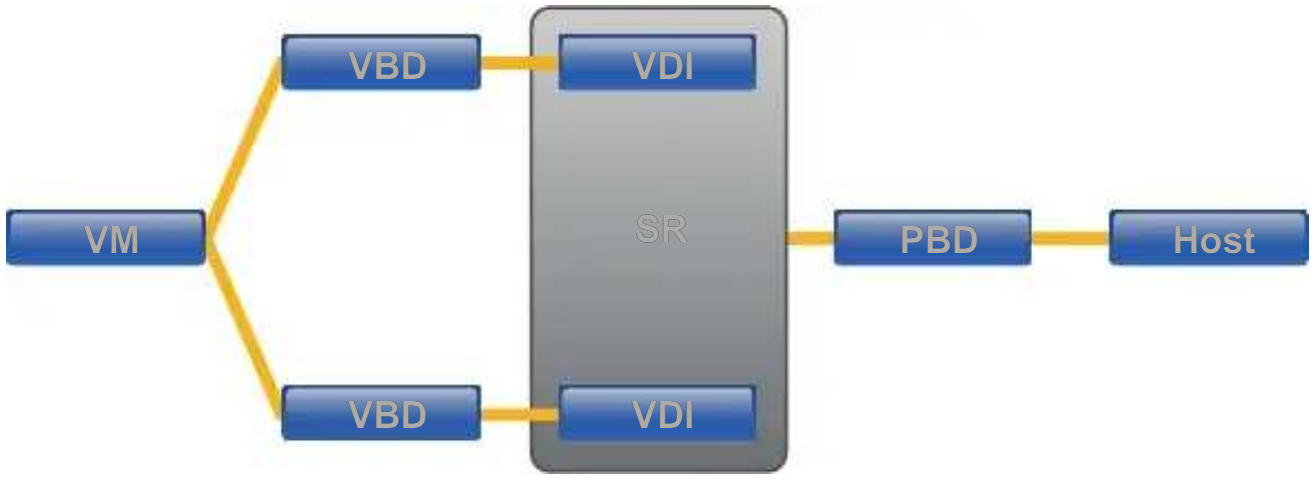

Figure 1: XenServer Storage Model

Dell EqualLogic PS Series feature definitions

This section describes some important PS series definitions relevant to this paper. Refer to PS Series online help for detailed information on various definitions and features.

Volume

A *volume* is a logical device representing a portion of storage pool space and seen on the network as an iSCSI target. Volume data is divided among pool members. Volumes are protected by access control records.

Thin provisioning

Thin provisioning technology allows one to more efficiently provision storage, while still meeting application and user storage needs. A volume has two capacity sizes associated with it:

- The reported size is the size that is seen by iSCSI initiators. Hosts connected to the volume will associate this amount of space with the volume.
- The volume reserve is the amount of space that is allocated to the volume.

For a volume without thin provisioning enabled, the reported size and the volume reserve are the same. For volume with thin-provisioning enabled, the reported size is greater than (or equal to) the volume reserve, because the volume is not fully allocated.

A thin-provisioned volume is initially allocated only a portion of the volume size. As data is written to the volume, more space is automatically allocated (if available), and the volume reserve increases up to a specified limit. Regular event messages occur as space is used, giving you the opportunity to make adjustments.

Because thin-provisioned volumes are only allocated the storage space that they actually use, you can more efficiently use storage resources, with no difficult resize operations on the computer.

Thin provisioning is most effective when you know how a volume usage will increase over time, the rate of growth is predictable, and users do not need immediate, guaranteed access to the full volume size. Administrators must monitor thin-provisioned volumes and be prepared to increase storage space by adding or moving arrays to the pool containing the volume. If you must guarantee the full volume size, thin provisioning is not recommended.

Volume snapshots

A *snapshot* is a space-efficient point-in-time representation of volume data. Snapshot creation does not disrupt access to the volume.

Like volumes, snapshots on the network look like iSCSI targets and can be set online and accessed by computers with iSCSI initiators. You can recover volume data by restoring a volume from a snapshot or by cloning a snapshot, which creates a new volume.

To create a snapshot, one must first reserve snapshot space for the volume. Snapshot reserve is taken from the pool containing the volume, and is based on a percentage of the current volume reserve. Group level defaults can be set to reserve a specified percentage of volume size for snapshot reserve.

PS Series group snapshot technology uses a reallocate-on-write technique. When users make changes to the base volume (or to the snapshot), the group stores the changes in the snapshot reserve.

Volume clones

A *volume clone* is an exact copy of a volume created from either the original volume or from a *snapshot* or *replica*. Cloning a volume creates a new volume with a new name and iSCSI target, but the same size, contents, and thin provisioning settings as the original volume at the time of the cloning. The new volume is created in the same pool as the original volume and is available immediately. Cloning a volume does not change the original volume.

A cloned volume consumes 100% of the original volume size from free space in the pool. If you create snapshots or replicas of the new volume, they will use additional pool space. You cannot select a different storage pool (if available) during the clone operation, but you can move the new volume to another pool later.

Volume replicas

Replication technology in the PS Series firmware lets you copy volume data between groups, thereby protecting the data from a variety of failures, ranging from the destruction of a volume to a complete site disaster, with no effect on data availability or performance. Similar to a snapshot, a *replica* is a copy of a volume at a specific point in time. A *replica set* is the set of replicas for a volume created over time.

A volume and its replica *set* are always stored in different groups. By separating the groups geographically, volume data is protected against a complete site disaster. To replicate data between two PS Series groups, you must configure the groups as *replication partners*. The groups can be in the same building or an unlimited distance apart.

Each replication partner plays a specific role in the replication of a volume. The volume is stored on the *primary group*, while the volume replica set is stored on the *secondary group* in space that the secondary group delegates to the primary group. The primary group sends (replicates) the data and the secondary group receives the data. Mutual authentication offers security between partners.

The first replication of a volume sends the entire contents of the volume to the secondary group. In subsequent replications, only the data that changed since the previous replication is sent to the secondary group. The longer a replication takes, or the longer the interval between scheduled replications, the more the data that might have to be transferred to the partner.

To recover volume data from replicas in the secondary group, you can clone an individual replica to create a new volume. You can also host a volume on the secondary group, either temporarily, for example to do backups, or permanently (e.g. to replace a failed primary group or to change the roles of the primary and secondary groups).

If you are hosting the volume temporarily on the secondary group, you can later fail back to the primary group and use the original replication configuration.

iSCSI storage repositories in XenServer 5.0

XenServer 5.0 hosts can access iSCSI storage using the built in open-iscsi software initiator or an iSCSI HBA. XenServer 5.0 includes support for two types of iSCSI storage repositories:

LVM over iSCSI or lvmoiscsi SR: utilizes the Linux Volume Manager (LVM) to create a logical volume per virtual disk image (VDI) on the iSCSI LUN. The VDIs residing in LVM over iSCSI SRs do not provide support for sparse provisioning or fast cloning. The SR utilizes the entire LUN specified at creation time and may not span more than one LUN.

Dell EqualLogic or equal SR: utilizes the XenServer EqualLogic Storage Adapter to manage virtual machine storage on PS Series Group. For each SR of type *Dell EqualLogic*, a small management volume is created on PS Series Group. For each XenServer virtual disk image (VDI) in the SR, a corresponding volume on PS series group is created. This allows for advanced VM life cycle operations such as snapshots, fast clones, thin provisioning, volume resizing etc.

Both *LVM over iSCSI* and *Dell EqualLogic* SRs are capable of supporting VM agility using XenMotion: VMs can be moved across XenServer hosts in a pool without noticeable downtime. CHAP support is provided for client authentication, during both the data path initialization and the LUN discovery phases.

LVM over iSCSI vs. Dell EqualLogic SR usage considerations

When using EqualLogic PS Series to provide iSCSI storage, you should consider storage and feature requirements to decide whether to use the specialized Dell EqualLogic SR, or to use the generic LVM over iSCSI SR. The specialized XenServer EqualLogic Storage Adapter enables XenServer to communicate with the PS Series group to manage storage for VMs and use advanced array features such as thin provisioning, fast cloning and snapshots. Table 1 below summarizes comparison between the LVM over iSCSI and EqualLogic SR capabilities.

Features/ Requirements	LVM over iSCSI SR	Dell EqualLogic SR
Shared SR for XenServer Pool	Yes	Yes
XenMotion	Yes	Yes
Resize Virtual Disk Image (VDI)	Yes	Yes
Thin Provisioning	No	Yes
Fast Cloning	No	Yes
Virtual Machine Snapshots	No	Yes
PS Series volume management from XenServer/XenCenter	No	Yes
Number of volumes on PS Series per SR	One for SR; each VDI is a logical volume	One for SR plus one for each VDI
iSCSI connections per host to PS series group	One per SR	One per active VDI
Connection load balancing for an SR on PS series group	From a host, connection to an SR through a single iSCSI port on a member	Connections to VDI volumes load balanced across member iSCSI ports

Table 1: LVM over iSCSI Vs EqualLogic SR

Setup and configuration

A simple setup utility lets you quickly configure a PS Series array as a member of a new group. After you choose the RAID type, the array is automatically configured, and the storage is immediately ready for use. Data and network I/O are automatically load balanced across the disks and network interfaces—with no impact on data availability.

To increase SAN capacity and performance, connect another array to the network and add it to the group. Load balancing across the members occurs automatically, as needed. Management overhead remains the same, regardless of the group size.

For instructions on initial PS series array configuration, please refer to the PS Series Quick Start Guide at http://www.equallogic.com/support

Typical PS Series-XenServer setup

This section describes a typical multi XenServer host setup with EqualLogic PS Series Group. As shown in Figure 2, Citrix XenServer Dell Edition runs on Dell PowerEdge servers, configured in a resource pool. All XenServer hosts are connected via an Ethernet fabric and share storage on the Dell EqualLogic PS5000 array(s). The iSCSI traffic is physically isolated from other network traffic using separate physical Ethernet switches. Two NICs on each host are bonded to provide high availability and load balancing for host management and virtual machine traffic and high availability for IP storage traffic.

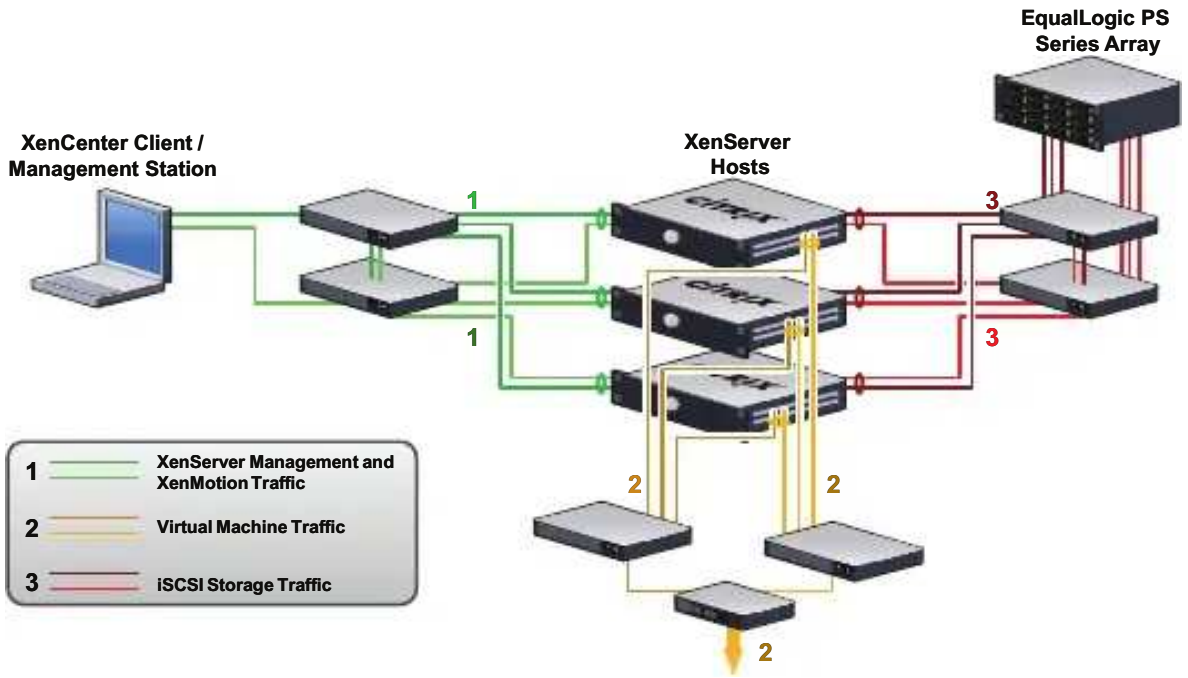

Figure 2: A sample of typical XenServer Pool Configuration using PS Series storage

This configuration illustrates traffic segregation using separate physical switches; however, isolation can also be achieved using VLANs. XenCenter can be installed on any workstation or server machine that has access to the host management network.

Note: Refer to XenServer Dell Edition Solution Guide for more information on supported Dell storage configurations with XenServer.

Multi-pathing with XenServer and PS Series

Multi-pathing in XenServer for iSCSI traffic is achieved using a highly available bond device on the XenServer host. In case of failure of a network port, the other network port in the bond takes over the network traffic. As shown in Figure 2, it is recommended to use redundant Ethernet switches to prevent a path failure at the switch level. Note that load balancing for network traffic over a bond is supported only for virtual machine traffic and not for IP storage traffic. At any given time, only one NIC, out of two in a bond, is used for iSCSI traffic. If the active NIC in a bond fails, the other NIC will take over. Refer to the XenServer 5.0 Reference Manual for more information on creating and managing NIC bonds.

The EqualLogic PS arrays have two controllers; only one controller is active at a given time. Each controller has 3 network ports for iSCSI connectivity. The XenServer hosts connect to the PS Series Group using the well known address or the Group IP address rather than the IP address of an iSCSI port on a member array. When an initiator connects to the PS Series Group using the Group IP address, the group redirects the initiator (via iSCSI redirect functionality) to log-in to an iSCSI port on a member. The PS series group intelligently redirects incoming connections to load-balance iSCSI traffic across iSCSI ports on group members. In case a port on a controller fails, the connections are redirected to other available ports. In case the active controller fails, the passive controller on a member becomes active and takes over IP settings of the failed controller.

Note: XenServer 5.0 natively uses dm-multipath for iSCSI traffic where the iSCSI storage arrays expose multiple target portals. However this is not the case with EqualLogic PS arrays. For this reason, enabling multi-pathing for a XenServer host using XenCenter or CLI does not use dm-multipath functionality.

Sample XenServer host and PS Series network configuration

This section discusses a sample network configuration for a XenServer host and a PS Series group to create a configuration as illustrated in Figure 2. Figure 3 shows the XenServer host network configuration. The host has four Gigabit LOM ports and two single port Gigabit network adapters. Using 2 NIC ports each, three bond devices (for host management, IP storage and virtual machines) are created using the XenCenter interface. In XenServer, the host management interface is used for the IP traffic by default. However, as a best practice it is recommended to use a separate host interface for IP storage traffic. This is achieved by creating another host interface on a subnet separate from the host management interface.

1. Using XenCenter interface, create a bond device for host management and virtual machine network. Active/active bonding in XenServer 5.0 provides both high availability and load balancing for virtual machine and host management traffic. Figure 4 shows the host management interface as seen in XenCenter interface.
2. To provide high availability for storage traffic, create a bond using two NIC ports.
3. Create a new host interface for storage traffic: in XenCenter interface, dick on the host and select *"Management Interfaces."* Create a new interface and configure the IP settings such that the storage management interface is on a separate subnet from the host management interface, and on the same subnet as the PS Series Group. Figure 4 shows storage interface as seen in the XenCenter interface. Note that the storage interface must be on a separate subnet from the host management interface to segregate IP storage traffic from host management traffic.

Example XenServer Host IP configuration:

Host Management Interface: 172.17.40.71/255.255.255.0

Host Storage Interface (bond):172.17.50.115/255.255.255.0

Name	Description	NIC	VLAN	Auto	Link Status	MAC
Host Management	XenServer API Management and XenMotion	Bond 0+2	-	No	Connected	00:1e:4f:fb:d8:96
IP Storage	iSCSI Storage Network	Bond 1+4	-	No	Connected	00:1e:4f:fb:d8:98
Network 0 (Slave)		NIC 0	-	Yes	Connected	00:1e:4f:fb:d8:96
Network 1 (Slave)		NIC 1	-	Yes	Connected	00:1e:4f:fb:d8:98
Network 2 (Slave)		NIC 2	-	Yes	Connected	00:1e:4f:fc:00:6e
Network 3 (Slave)		NIC 3	-	Yes	Connected	00:1e:4f:fc:00:70
Network 4 (Slave)		NIC 4	-	Yes	Connected	00:10:18:37:44:ee
Network 5 (Slave)		NIC 5	-	Yes	Connected	00:10:18:37:43:6c
Virtual Machine	Virtual Machine Network	Bond 3+5	-	Yes	Connected	00:1e:4f:fc:00:70

Figure 3: Sample host management and storage networks on a XenServer host

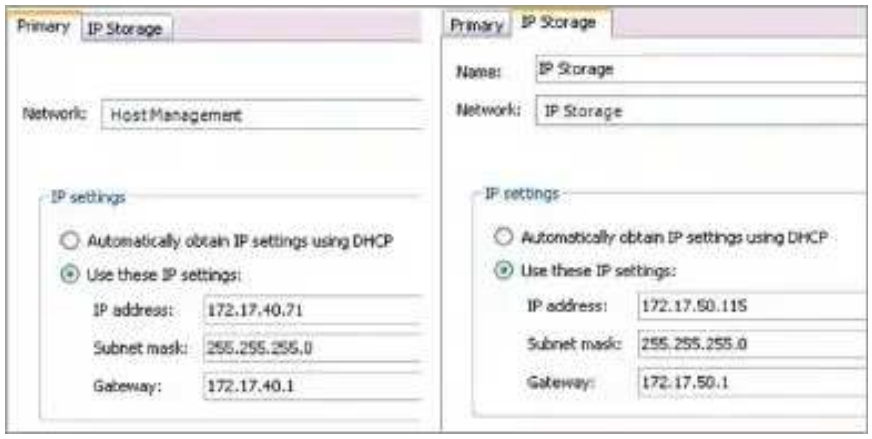

Figure 4: XenServer Host Networking: Management and Storage Interfaces

Note: As a best practice, it is recommended to use a static IP address for host management and storage interface.

Example IP configuration for a PS Series group with one member:

PS Series Group and member IP configuration:

PS Series Group IP: 172.17.50.40/255.255.255.0

PS Series Member Storage Port (eth0): 172.17.50.41/255.255.255.0

PS Series Member Storage Port (eth1): 172.17.50.42/255.255.255.0

PS Series Member Storage Port (eth2): 172.17.50.43/255.255.255.0

Note: PS Series network configuration requires that the Group IP and the IPs for the member storage ports (eth0, eth1 and eth2) must all be on the same subnet.

If using a separate management IP (other than the Group IP) to manage a PS Series group, then the Group IP and storage ports eth0 and eth1 need to be on same subnet. The management IP and eth2 for the members should be on the same subnet. For more information on configuring a management interface (other than the group interface) for a PS Series group, refer to the PS series online help.

Networking best practices for PS Series and XenServer

Refer to the PS Series Array Network Performance Guidelines document for information on general network configuration best practices for EqualLogic PS Series arrays.

Note: Jumbo frames are not supported in XenServer 5.0.

Note: XenServer 5.0 natively uses dm-multipath for iSCSI traffic where the iSCSI storage arrays expose multiple target portals. However this is not the case with EqualLogic PS arrays. For this reason, enabling multi-pathing for a XenServer host using XenCenter or CLI does not use dm-multipath functionality.

Creating a storage repository on PS Series Group

Creating a storage repository of type *LVM over iSCSI (lvmoiscsi)*

1. Following the instructions available in the PS Series online help, configure the PS Series array, create a group, set a member RAID policy, create a volume and set appropriate access control method (using XenServer host IQN or management IP).
2. If creating an SR for a pool, make sure that the storage volume is enabled for shared access from multiple initiators. To enable shared access, in the PS series group manager user interface, right-click on the volume and select *"Modify Volume Settings."* Select the Advanced tab and check *"Enable shared access to the iSCSI target from multiple initiators"* (See Figure 5).

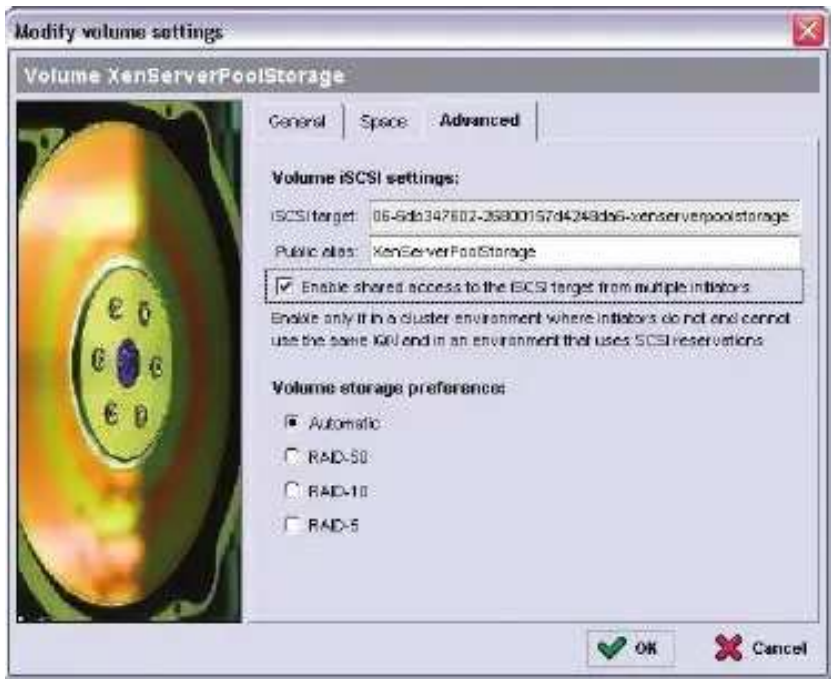

Figure 5: Enable shared access to the SR volume when SR is intended to be shared

3. Using the XenCenter interface, right-dick on the XenServer host or the pool for which you would like to create SR for and select *"New Storage Repository."* Choose the iSCSI radio button for the type of new storage, and select *"Next."*
4. Enter the name for the new SR in the *name* field, PS series Group IP address in the *Target Host* field, and 3260 in the *port* field. Provide the chap username and password information if configured for the volume. Then click on *"Discover IQNs"* to perform target discovery. Each volume on a PS Series array has a unique target name with LUN ID set to "0." Select the appropriate IQN and then click on *"Discover LUNs"* to discover the LUN or volume associated with the target. Select the volume and click *"Finish"* to create a new SR.

Creating a storage repository of type *Dell EqualLogic (equal)*

Note: To use an SR of type *Dell EqualLogic*, the firmware version of PS Series arrays must be 4.0.1 or higher.

To create a storage repository of type *Dell EqualLogic* or *equal*, follow the steps below:

1. Following the instructions available in the PS Series online help, configure the EqualLogic array, create a group, set a member RAID policy, and create a volume.
2. Using the PS Series Group Manager User Interface, make sure that SSH access is enabled on the group members. Click on the *"Group Configuration"* and select the *"Administration"* tab. If not already enabled, check the *"Enable SSH (secure shell)"* access checkbox in the CLI access options for the group and save the configuration.
3. Using the XenCenter interface, right-click on the XenServer host or the pool for which you would like to create SR for and select *"New Storage Repository."*
4. In the *"Choose the type of new storage window"*, select *"Dell EqualLogic"* (see Figure 6) and click *"Next."*

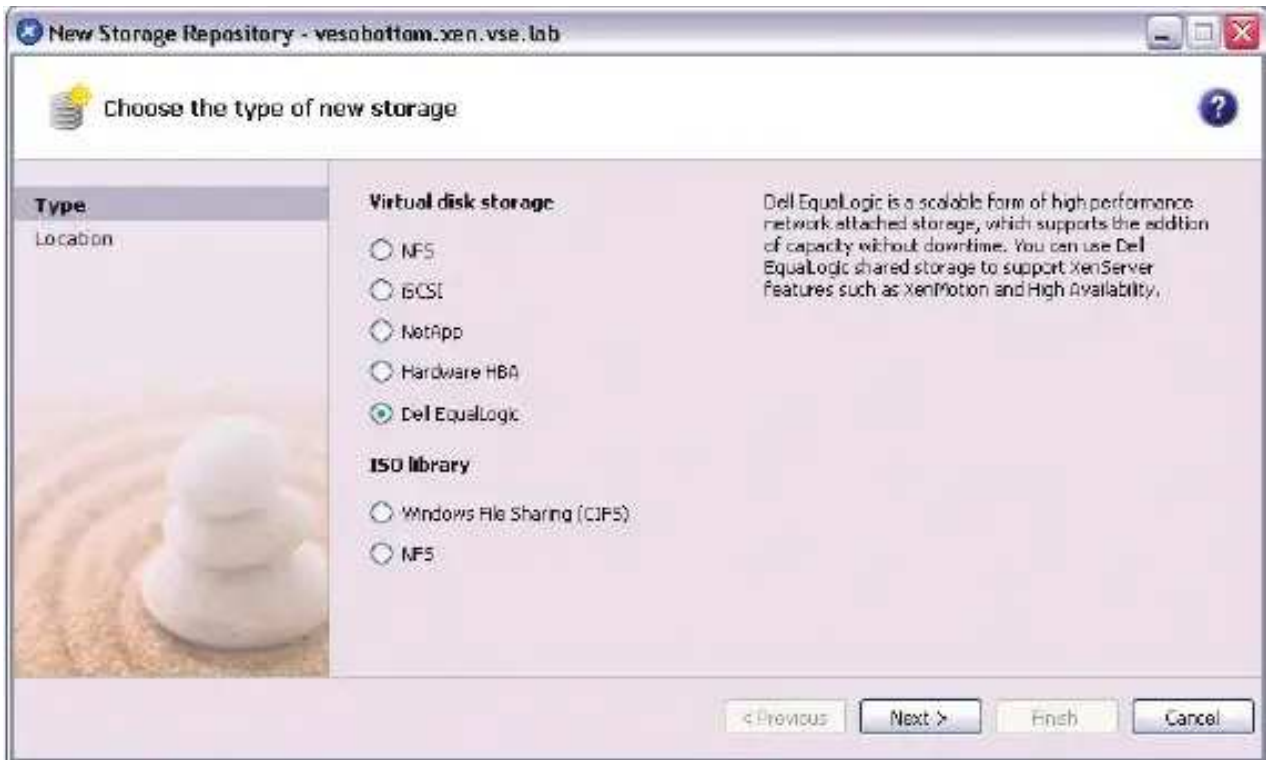

Figure 6: Select Dell EqualLogic to create an SR of type Dell EqualLogic

5. In the *"Enter the Dell EqualLogic filer details"* window, provide a name for the SR in the name field, the Group IP of your PS Series group, and the user name and password for an account with administrative privileges for group or a pool. Optionally, provide CHAP credentials if set on the PS Series group and click *"Next."*
6. The following screen provides two options: Reattach an existing SR or create a new SR. Choose the *"Create a new SR on the following storage pool"* radio button. You can see the details on the available storage pool including the size, free available space, number of existing volumes and members in the pool.

   If you like all VDIs in the SR to be thin provisioned, select the *"Thin Provisioning"* checkbox when creating the new SR (see Figure 7).

   Click *"Finish"* to create a new SR.

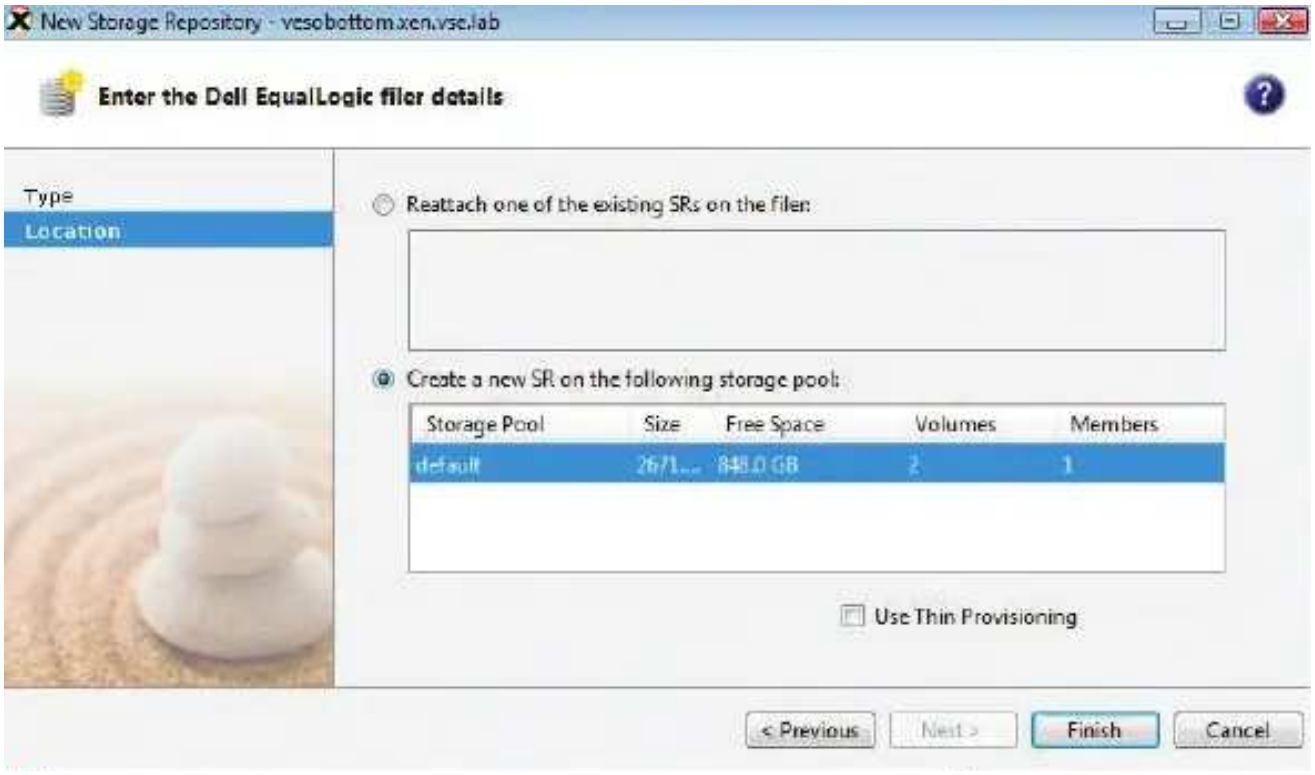

Figure 7: Select the appropriate storage pool and thin provisioning settings for EqualLogic SR

As shown in Figure 8, upon completion of the SR creation process, a corresponding volume of size 20MB is created on the PS Series Group. When adding an SR to the pool, access control records for all XenServer hosts in the pool are added to the volume. The access control is specified by the IQN of the XenServer host.

Note: Volume name and description fields are used by XenServer EqualLogic Storage Adapter to manage volumes on the PS Series Group. Do not manually change the name and description fields for any volume automatically created by a XenServer host.

Figure 8: SR volume as shown in PS Series Group GUI

Identifying XenServer storage objects on a PS Series Group

For *Dell EqualLogic* SR, an SR is just a logical grouping of VDIs or volumes on the PS Series Group. The naming scheme for volumes automatically created by XenServer EqualLogic Storage Adapter is as follows:

At the point the SR is created, a management volume is provisioned on the PS Series Group. This serves to identify the persistent SR presence on the group, as well as providing a location to store SR metadata as required by XAPI. The management volume is named as:

"XenStorage<SR_UUID>MANAGEMENT"

where SR_UUID is the UUID of the storage repository.

For example, in Figure 9, the volume
`XenStorageb93f1728dd0c6077e96dbc968a047fb2MANAGEMENT`
corresponds to an SR with UUID `b93f1728dd0c6077e96dbc968a047fb2`

Each volume corresponding to a VDI within an SR is identified by the identification string "XenStorage" followed by the SR UUID, and a cropped version of the VDI UUID. The full VDI UUID is stored in the Volume description field.

For example, as shown in Figure 9 the volume with name
`XenStorageb93f1728dd0c6077e96dbc968a047fb22332cdcae4fd4a3d9bbc`
corresponds to a VDI. The VDI is related to SR with UUID `b93f1728dd0c6077e96dbc968a047fb2` and the cropped UUID of VDI is `2332cdcae4fd4a3d9bbc`. The full UUID of VDI is available in the volume description.

Figure 9: A VDI volume as shown in PS Series Group GUI

Note: The XenServer EqualLogic Storage Adapter uses volume name and description fields to reference volumes on PS Series Group. Manually changing either the volume name or description fields to values other than those assigned by the XenServer storage adapter may lead to failure in storage management tasks.

Typical lifecycle tasks for VMs on Dell EqualLogic SR

This section describes typical virtual machine life cycle operations for virtual machines created and managed using the XenServer EqualLogic Storage Adapter.

Create a virtual machine

The process to create a virtual machine with storage on Dell EqualLogic PS series group remains the same as with any other storage backend. While creating a VM, just select the appropriate SR created on your PS Series Group. For more information on creating virtual machines, refer to *XenServer Virtual Machine Installation Guide*.

Create a VDI

When using an SR of type *Dell EqualLogic*, for each Virtual Disk Image (VDI) created on a XenServer host, a corresponding volume is created on the PS Series Group. Figure 9 shows a volume, with name `XenStorageb93f1728dd0c6077e96dbc968a047fb22332cdcae4fd4a3d9bbc`, that has been automatically created by the XenServer EqualLogic Storage Adapter.

The size of the volume on the array is same as the VDI size specified by the user.

Note: When a VDI is created using XenCenter, the corresponding volume created on the PS series Group inherits the default group volume and thin provisioning settings.

To create a VDI using CLI use the xe vdi-create command:

```
# xe vdi-create virtual-size=10GiB name-label=testvdi
sr-uuid=<SR UUID> type=user sm-config:allocation=<thin | thick>
sm-config:snap-reserve-percentage= <Percentage Integral Value>
sm-config:snap-depletion=<delete-oldest | volume-offline>
```

where `SR_UUID` is the UUID of the SR of type Dell EqualLogic

sm-config:allocation controls whether the VDI volume is provisioned as a thin volume or not. Setting `sm-config:allocation=thin` will create a volume with thin provisioning enabled. Setting `sm-config:allocation=thick` will create a volume with thin provisioning disabled. If type of allocation is not specified, the default allocation for the SR is used to provision the VDI volume.

`sm-config:snap-reserve-percentage` specifies amount of space, in terms of percentage of volume, to reserve for volume snapshots. If `snap-reserve-percentage` is not specified, the group default value for snapshot reserve is used.

`sm-config:snap-depletion` specifies the snapshot space recovery policy:
action taken when the space reserved for snapshots has been exceeded. Setting `sm-config:snap-depletion =delete-oldest` deletes the oldest snapshots until sufficient space is recovered (the default). Setting `sm-config:snap-depletion=volume-offline` sets the volume and snapshots offline. Active iSCSI connections will be terminated before a snapshot is automatically deleted.

For example, the following command to create a new VDI

```
# xe vdi-create virtual-size=10GiB name-label=vdi1 type=user
sr-uuid=b93f1728-dd0c-6077-e96d-bc968a047fb2 sm-config:allocation=thin
sm-config:snap-reserve-percentage=50 sm-config:snap-depletion=delete-oldest
```

provisions a volume, on PS Series Group, of size 10GB, with a snapshot reserve of 50% of volume size (10GB), thin provisioning enabled and snapshot space recovery policy to delete oldest snapshot.

Delete a VDI

Deleting a VDI deletes the associated volume on the PS Series Group. However if a VDI has any associated snapshot, deleting a VDI from XenServer host will not delete the parent volume on PS series Group until the last existing snapshot is deleted.

A VDI can be deleted or destroyed using the XenCenter interface or using a **`vdi-destroy`** command.

```
# xe vdi-destroy uuid=<VDI_UUID>
```

Resize a VDI

To resize a VDI from XenCenter, right click on the VDI and select *Edit Virtual Disk*. Modify the size of the VDI and click *OK*.

To resize a VDI using CLI, use the xe vdi-resize command:

```
# xe vdi-resize uuid=<VDI_UUID> disk-size=<new size>
```

where **`VDI_UUID`** is the UUID of the VDI under consideration

`new size` the new size of the VDI.

For example, the following command sets the size of the VDI to 50GB:

```
# xe vdi-resize uuid=6be5ed34-7600-4218-a68e-741aeb4ee815 disk-size=50GiB
```

When a VDI is resized using XenCenter or CLI, the size of the corresponding volume on PS series group is also set to the new size value specified for the VDI. The snapshot reserve for the volume is also automatically adjusted to comply with the predefined percentage of volume size.

Copy or clone a VM

A copy or clone of a VM can be created in two ways:

Fast Clone—Uses PS Series volume clone feature to instantaneously create clones of volumes corresponding to VDIs attached to the virtual machine.

Full Copy—Creates a copy of the source virtual machine storage volumes by copying all data over to new volumes. Depending on the size of data to be copied, a full copy operation may take from a few to several minutes to complete.

The XenServer fast cloning feature utilizes PS Series volume cloning to provision new virtual machines within seconds. To provision a new VM using fast cloning, right click on the source VM and select *"Copy VM."* Provide a name and description for new VM and select *"Fast Clone"* as the copy mode (see Figure 10). Click *"Copy"* to finish creating a new VM that is a copy of a specified source VM.

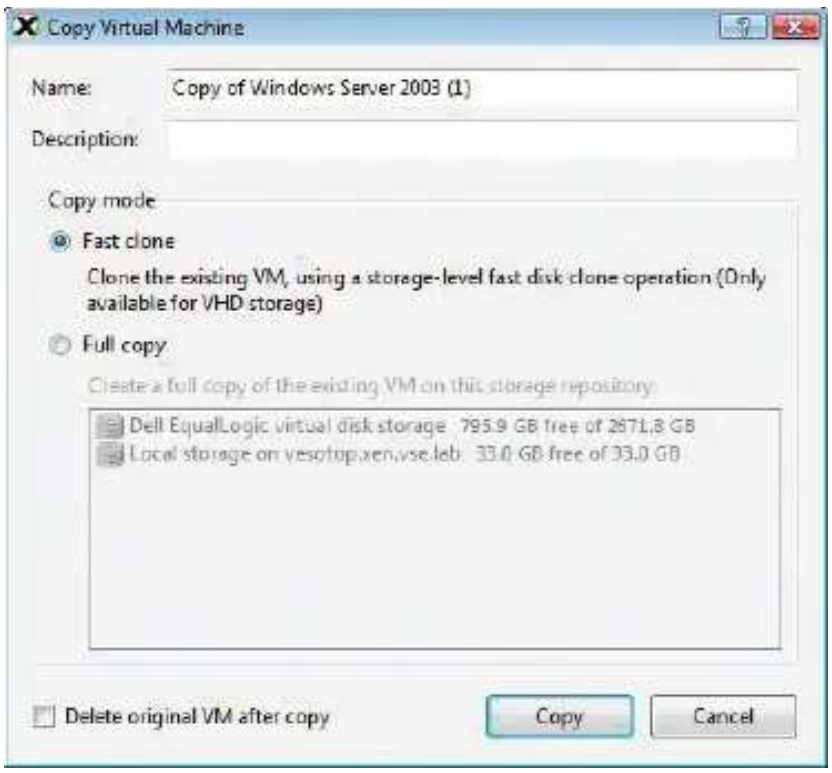

Figure 10: Fast cloning feature uses PS Series volume cloning to instantaneously provision new VMs

Deploy from Template

A template is a "gold image" that contains all the configuration settings to instantiate a specific VM. With XenServer one can create VMs, configure them in standard forms for particular needs, and save a copy of each as a template for future use in VM deployment.

If all the VDIs of the template VM reside on an SR of type *Dell EqualLogic*, when deploying a new VM, in the same SR, from the template, XenServer uses the PS Series cloning feature to instantaneously provision a new virtual machine. However if a VM is deployed to an SR other than the SR where template resides, XenServer provisions the new VM by copying.

Snapshot a VM

XenServer provides a convenient VM snapshot mechanism that can take a snapshot or capture a point-in-time copy of a VM's storage and metadata. Where necessary, I/O 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. XenServer 5.0 also includes a Xen VSS provider (installed as part of Citrix XenTools) to enable quiesced snapshots. Quiesced snapshots take advantage of the Windows Volume Snapshot Service (VSS) for services that support it, so that a supported application (for example Microsoft Exchange or SQL Server) can flush data to disk and prepare for the snapshot before it is taken.

For virtual machines residing on SR of type *Dell EqualLogic*, when a snapshot operation is issued, XenServer EqualLogic Storage Adapter instructs the PS Series Group to take a snapshot of the volumes associated with the virtual machine's VDIs. Snapshots are supported for only *Dell EqualLogic* SRs and not for an SR of type *LVM over iSCSI*.

A snapshot on the PS series is treated as a separate target and gets its own IQN. When virtual machine snapshots are created from a XenServer host, access control records (specified by host IQN) for the snapshots are automatically created to allow the XenServer or a group of XenServer hosts in a pool to access the snapshots.

Figure 11 shows a virtual machine *win2003_vm1* which has two VDIs. Under the Virtual Disks tab on XenCenter, snapshots for the VDIs appear for each VDI. The corresponding view on PS Series GUI is shown in Figure 12.

Figure 11: VM snapshots as shown in XenCenter Interface

Figure 12: VM snapshots as shown in PS Series Group UI

XenServer treats snapshots as templates. As seen in Figure 11, the snapshot *win2k3_vm1_snap1* of virtual machine *win2k3_vm1* is also listed as a custom template. Just like regular templates, one can instantaneously deploy new virtual machines from such snapshots. When a new VM is deployed from a snapshot, XenServer uses the snapshot clone feature of the PS series arrays to instantaneously create clones of snapshots associated with the source VM. Access control records (specified by host IQN) for newly cloned volumes are automatically created to allow a XenServer or a group of XenServer hosts in a pool to access the volumes.

Refer to the XenServer Reference Guide for more information on regular and quiesced snapshots.

Restore a VM from a Snapshot

A virtual machine can be restored to a snapshot state if needed. One can restore the VDI volumes on PS series from snapshots; however one must restore all volumes attached to a VM. Another way to restore a VM from the snapshot state is to create a new virtual machine from a snapshot and deactivate or delete the old VM. Follow the steps in the VM Rollback section in XenServer 5.0 Reference Guide for instructions on how to restore a VM from a snapshot.

Delete a VM

When a VM is uninstalled, all the volumes, on the PS Series Group, associated with the virtual machine are deleted. When the volumes are deleted, the storage space is marked as free and returned to the storage pool.

If a virtual machine has any associated snapshots, deleting a virtual machine does not delete the volumes associated with virtual machine VDIs. When the last available snapshot associated with the virtual machine is deleted, the parent volume gets deleted as well.

Delete a Snapshot

When a snapshot of a virtual machine is deleted or uninstalled, the corresponding snapshots of all volumes (on PS Series) for VM VDIs are deleted from the PS Series Group. To delete or uninstall a snapshot, click on the template corresponding to the snapshot and select *Uninstall Template.*

Using PS Series replication with XenServer Portable SRs

XenServer 5.0 introduces the concept of Portable Storage Repositories (Portable SRs). Portable SRs contain all of the information necessary to recreate all the Virtual Machines (VMs) with Virtual Disk Images (VDIs) stored on the SR after re-attaching the SR to a different host or pool. Portable SRs can be used when regular maintenance or disaster recovery requires manually moving a SR between pools or standalone hosts.

Portable SRs work by creating a dedicated metadata VDI within the specified SR. The metadata VDI is used to store copies of the pool or host database as well as the metadata describing each VM's configuration. As a result the SR becomes fully self-contained, or portable, allowing it to be detached from one host and re-attached to another as a new SR. Once the SR is re-attached, a restore process is used to recreate all of the VMs on the SR from the metadata VDI.

For disaster recovery the metadata backup can be scheduled to run regularly to ensure that the metadata SR is current. For more information on XenServer 5.0 portable SRs and restrictions on usage, refer to the XenServer 5.0 Reference Manual.

Using replication for SR of type *LVM over iSCSI*

When using replication to protect storage repositories of type *LVM over iSCSI*, one needs to configure replication for volumes backing the SRs.

1. Configure a replication partner (secondary) group for the primary PS Series group.
2. Configure replication for the appropriate SR volume(s). Replication schedules to create replicas at specific intervals can be created for the volumes accessible to XenServer hosts.
3. On the primary XenServer host or pool, use the pool master console interface to schedule VM metadata backup on the protected Storage Repository. Make sure that replication is configured for the SR that holds pool or host metadata.

   Note: For suspended VMs, please make sure to replicate the designated suspend SR.

On the secondary site, the replicas may be used to perform backup operations or for disaster recovery operations should the primary site becomes unavailable. In order to attach the SR on XenServer host/ pool on the secondary site:

1. Once the replicas for the replicated SR volume are in the *Ready* state, promote the replica as a volume or create a replica clone. Note that by default, replica clones have thin provisioning enabled.
2. For the replica clone, add appropriate access control records for XenServer hosts on the secondary site to access the SR volume(s). For shared SRs, make sure to enable shared access for the SR volume(s).
3. On the XenServer host or pool master on the secondary site, attach the replica clone or promoted volume as an existing SR.
4. Using the *Restore Virtual Machine Metadata* option from XenServer local console, restore the VM metadata to recover all virtual machines on the replicated SR volume.

Using replication for SR of type Dell EqualLogic

See CTX122163 - How to Use Replication for Storage Repository Type Dell EqualLogic

References

XenServer Dell Edition Documentation: http://support.dell.com/support/edocs/software/Citrix/

Dell EqualLogic PS Series Online help: http://psonlinehelp.equallogic.com/V4.0/

XenServer 5.0 User Documentation: http://www.citrix.com/xenserver/dell

PS Series Array Network Performance Guidelines:
http://www.equallogic.com/uploadedfiles/Resources/Tech_Reports/tr-network-guidelines-TR1017.pdf

About Citrix

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 Dell

Dell Inc. (NASDAQ: DELL) listens to customers and delivers innovative technology and services they trust and value. Uniquely enabled by its direct business model, Dell is a leading global systems and services company and No. 34 on the Fortune 500. For more information, visit www.dell.com, or to communicate directly with Dell via a variety of online channels, go to HYPERLINK "http://www.dell.com/" www.dell.com/ conversations. To get Dell news direct, visit HYPERLINK "http://www.dell.com/RSS" www.dell.com/RSS.

©2008 Citrix Systems, Inc. All rights reserved. Citrix®, the Citrix logo, Citrix XenApp™, XenCenter™ and XenServer™ are trademarks or registered trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. Dell, the Dell logo, Dell OpenManage™, EqualLogic™, PowerEdge™ and PowerVault™ are trademarks or registered trademarks of Dell in the U.S. and/or other countries. All other trademarks and registered trademarks are the property of their respective owners.

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