

VMware[®] ESX Server 3.x Considerations, Configuration and Operation Using an EqualLogic PS Series SAN

ABSTRACT

This Technical Report shows how VMware ESX Server can work with an EqualLogic PS Series SAN to provide dynamic, scalable and highly available virtual environments.

TR1025

v2.0

Copyright © 2007-2008 EqualLogic, Inc.

January 2008

EqualLogic is a registered trademark of EqualLogic, Inc.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

Possession, use, or copying of the documentation or the software described in this publication is authorized only under the license agreement specified herein.

EqualLogic, Inc. will not be held liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change. Your performance can vary.

PS Series Firmware Version 3.2

EqualLogic, Inc.
110 Spit Brook Road
Building ZKO2
Nashua, NH 03062
Tel: 603.579.9762
Fax: 603.579.6910

Table of Contents

Revision Information	5
Introduction.....	6
Why VMware Virtualization	6
Benefits of Virtualization.....	6
Server Consolidation	6
Facilities and Physical infrastructure.....	7
Rapid provisioning	7
Testing and Development.....	7
High Availability	7
Disaster Recovery.....	7
Virtual Infrastructure 3	7
VMware ESX Server.....	8
VMware VirtualCenter	8
VMware VMotion	9
VMware DRS	9
VMware HA	10
VMware Storage VMotion (Only available in 3.5/3i).....	11
Site Recovery Manager	11
VMware Consolidated Backup.....	11
VMware File System.....	11
EqualLogic iSCSI Arrays for Virtualized Storage.....	12
VMware + EqualLogic = An Adaptable Virtual Environment	12
Optimized Asset Utilization	12
Reduced Deployment Time and Effort.....	12
Reduced Management Complexity	12
High Availability	13
High Scalability	13
Increased Flexibility	13
Planning Considerations	13
Boot Options.....	13
Server and Storage Connectivity	14
Volume Connectivity and Data Formats	14
Volume size considerations	15
Storage Array RAID Levels	15

Server and Storage Redundancy.....	16
MPIO	16
MPIO with ESX 3.x Networking	18
Data Protection	18
Replication	18
Snapshots	19
Templates and Rapid Provisioning.....	19
Deployment.....	19
VirtualCenter	19
EqualLogic PS Series SAN	19
Deploy ESX 3.x.....	20
Connect.....	20
Provision.....	20
Conclusion	20
Appendix A: Accessing a PS Series SAN using the ESX software iSCSI Initiator	21
Appendix B: Accessing a PS Series SAN Using an HBA hardware Initiator	31
EqualLogic Documentation and Customer Support	36

REVISION INFORMATION

The following table describes the release history of this Technical Report.

Report	Date	Document Revision
1.0	06/2006	Initial Release
1.1	10/2006	Updated for ESX Server 3.0.1 and Virtual Center 2.0.1
2.0	01/2008	Second Revision, Updated for ESX 3.0.2/3.5/3i and Virtual Center 2.0.2/2.5

The following table shows the software and firmware used for the preparation of this Technical Report.

Vendor	Model	Software Revision
VMware	ESX Server	3.0.2/3.5/3i
EqualLogic	PS Series Firmware	Version 3.2 or higher

All EqualLogic Technical Reports are available on the Customer Support site at:
<https://www.equallogic.com/support/>

INTRODUCTION

This technical report describes how VMware Virtual Infrastructure 3 can work together with EqualLogic PS Series iSCSI SANs to provide dynamic, scalable and highly available virtual environments. This report also provides guidelines for planning, designing and implementing a virtual environment appropriate to your data center needs. With the release of ESX 3.5/3i there are some new added features and functionality. This technical report will address Virtual Infrastructure 3 as a whole and where needed, will point out which version of ESX the function is available for.

WHY VMWARE VIRTUALIZATION

In today's data center environments, virtualization makes sense to make use of the powerful systems that are available today. Virtualization is the technology that separates the physical computing resources from the software that utilizes them. In this way you can run multiple systems on one hardware platform maximizing the utilization and efficiency of that platform.

A primary goal of virtualization technology is to increase resource management and flexibility. VMware's virtualization approach makes it possible to run multiple instances of Virtual Machines (VMs) on each hardware platform, sharing hardware resources and thereby getting higher utilization of those platforms.

Each VM runs its own operating system and applications, and is isolated from all other VMs. They cannot communicate with each other except through normal networking channels. This isolation prevents a problem in one virtual machine from bringing down the entire physical server or any of the other Virtual Machines.

Because VMware utilizes virtual drivers and hardware, the operating system in the VM is not hardware dependent. This makes it much easier to migrate to different systems when you obtain newer hardware or in DR scenarios.

Another benefit is that all of the Virtual Machines are encapsulated. The operating system, configuration and memory are all contained in files making it easy to migrate and move these files to other areas.

BENEFITS OF VIRTUALIZATION

Server Consolidation

Virtualization can be utilized to meet any number of challenges in the modern data center. You can use this technology for server consolidation and optimization. With the high powered hardware platforms that are being delivered by server vendors, you can use virtualization software to maximize the utilization of these platforms by combining multiple systems onto that platform. With dual core and quad core processors, along with upwards of 32GB to 128GB of memory, you can optimize the servers and services you need to get the most of these hardware resources. This virtualization technology allows you to break free of the single application per server model. VMware case studies often report post-virtualization server utilization rates of 60% or better, with resulting reductions in required hardware of 7 to 1 or better.

Virtualization software also allows for the migration of legacy systems that the business still needs, to newer hardware. The legacy hardware may be out of warranty or it may take up too much space or it simply doesn't perform well. Migrating these servers to new hardware in a virtual environment allows for legacy software to live on with very little disruption to the environment.

Facilities and Physical infrastructure

Virtualization can also help your datacenter with reduction in footprint and physical infrastructure cost reduction. Data center retail space is at a prime. With raised floors, racks, cooling requirements and other facility considerations, saving space, power and cooling is a huge benefit of virtualization. With newer pushes for “greener” data centers, virtualization can help by decreasing the actual number of physical systems without impacting the business needs.

Rapid provisioning

Rapid provisioning of servers is another reason why businesses are making the move to virtualization. You can easily provision a new server to meet a business need almost immediately instead of waiting for new hardware to arrive and set it up. Along with this, you can baseline your standard images for your business so that you know when you do provision a server; it will be deployed as you need it with minimal reconfiguration.

Testing and Development

Virtualization is a great resource for testing and development. With the ability to rapidly provision servers, you can test patches, applications, and various environments. You can also use virtualization to prep an environment before deploying it into production

High Availability

Virtualization also allows for various aspects of high availability. Using VMware’s ESX server you can use native abilities such as NIC teaming, virtual machine clustering as well as SAN features like Dynamic Resource Scheduling (DRS), High Availability (HA), and VMotion. This allows IT administrators to reduce downtime and achieve higher levels of availability.

Disaster Recovery

Disaster recovery is made easier because of the encapsulation and hardware independence that virtualization provides. It also reduces the amount of physical equipment needed at a DR site to provide full functionality in the case of a failure.

VIRTUAL INFRASTRUCTURE 3

VMware Virtual Infrastructure 3 is the entire suite of applications and abilities that VMware provides to create a virtual environment with high availability. These include:

- VMware ESX Server 3.x
- VirtualCenter
- VMotion
- Distributed Resource Scheduler (DRS)
- High Availability (HA)
- Storage VMotion (Only available in 3.5/3i)
- Site Recovery Manager (SRM)
- VMware Consolidated Backup (VCB)
- VMware File System (VMFS)

VMware ESX Server

VMware ESX Server is enterprise infrastructure virtualization software for mission-critical environments. While VMware Server sits on top of the server's already-installed OS, VMware ESX Server replaces the host OS on the server, applying the server's full resources to the virtual machines. This maximizes the performance of the virtual machines running on the server. Figure 1 shows the VMware ESX Server architecture.

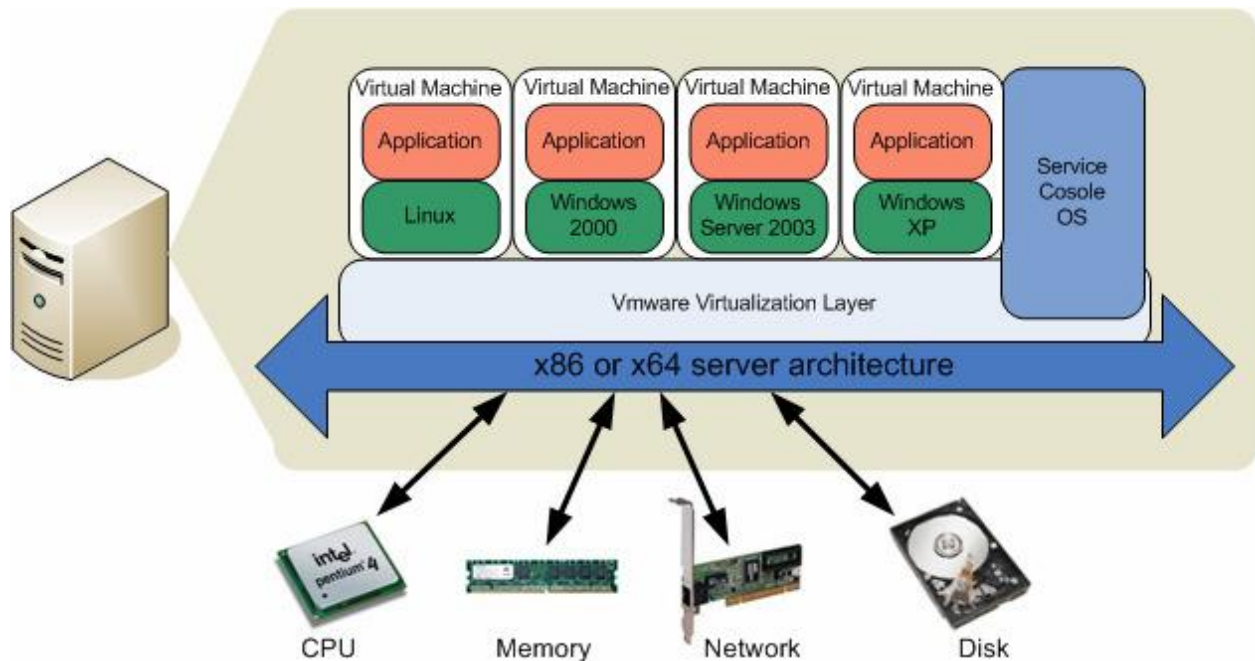


Figure 1: VMware ESX Server Architecture

VMware ESX Server is highly scalable and provides additional benefits in management, security, high availability, and performance beyond those in VMware Server. VMware ESX Server runs on a wide variety of x86-based servers.

VMware ESX 3i introduces a new type of hypervisor. It is essentially ESX 3.5 but because of its 32MB footprint, server vendors will be able to imbed it into the server itself allowing for diskless servers.

VMware VirtualCenter

VMware VirtualCenter provides a single control center and management center for VMware ESX Servers. VMware VirtualCenter allows you to:

- Manage hundreds of Servers
- Performance monitoring and resource optimization
- Use standard templates to provision new servers in minutes
- Apply access control to the virtual environment
- Utilize advanced features such as VMotion, HA, and DRS
- And more

The following advanced features require a SAN for centralized storage.

VMware VMotion

VMware VMotion is an optional feature enabled by VMware VirtualCenter. This technology allows for Virtual Machines to be moved from one ESX server to another. This can be done for load balancing or better resource utilization. This capability also makes it possible to perform zero-downtime maintenance by moving virtual machines to different ESX servers to allow for physical maintenance as well as patching and upgrades. Figure 2 shows the VMotion architecture.

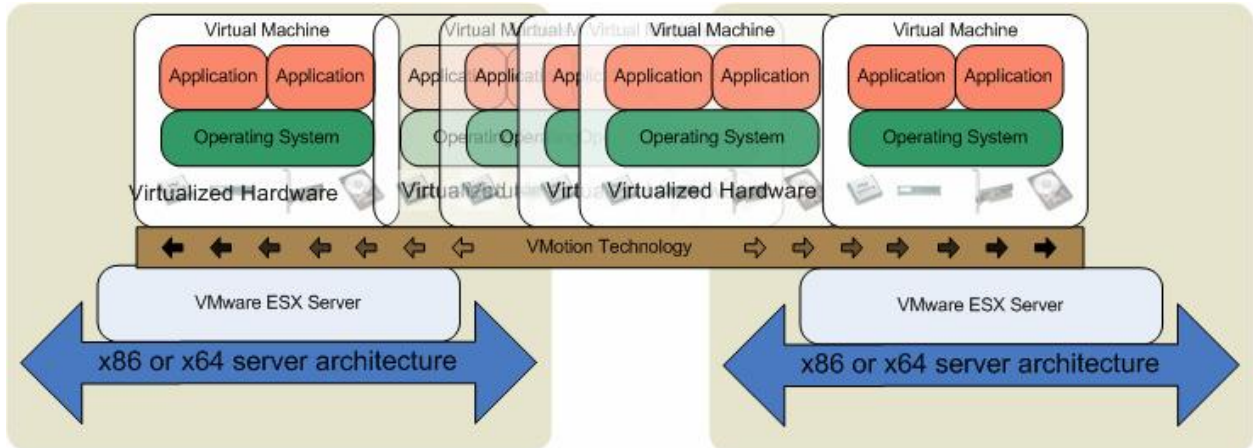


Figure 2: VMotion moving a virtual machine to a different server

VMware DRS

VMware DRS (Distributed Resource Scheduler) tracks resource use across ESX Servers and resource pools, and depending on how you have configured your virtual infrastructure, either recommends the migration of virtual machines to different hosts or uses VMotion to migrate the virtual machines automatically. Figure 3 shows the DRS interface screen.




Figure 3: DRS Settings

VMware HA

VMware HA (High Availability) allows you to organize virtual machines into failover groups. When an ESX Server fails, all of its virtual machines that were in failover groups are automatically restarted on different ESX Servers. Figure 4 shows the HA configuration menu.

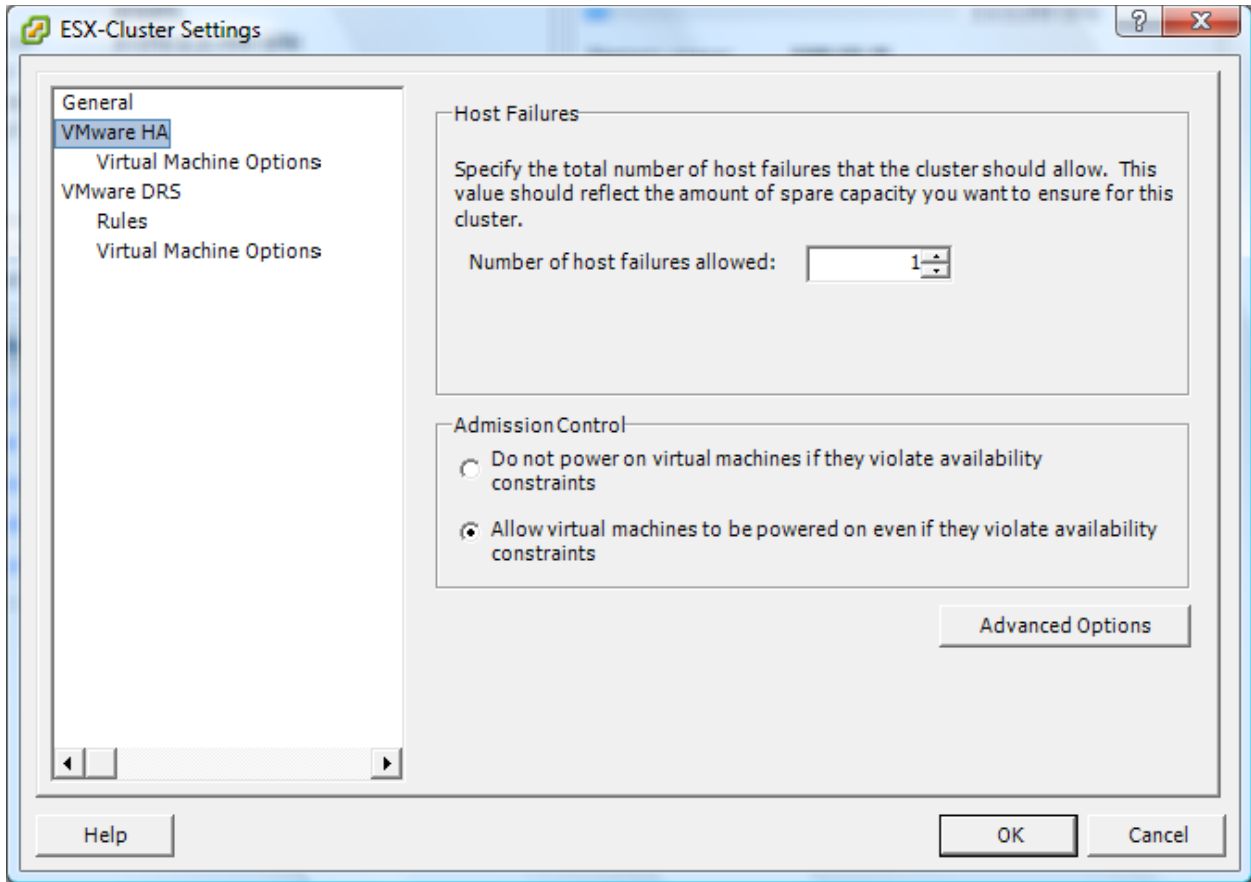


Figure 4: HA Configuration

VMware Storage VMotion (Only available in 3.5/3i)

VMware Storage VMotion is similar to VMotion except instead of migrating the Virtual Machine from one ESX host to another, instead it migrates the encapsulated data files from one datastore to another. This can be used to migrate Virtual Machines between different storage arrays, different tiers of storage or different volumes within the PS Series SAN.

Site Recovery Manager

Site Recovery Manager (SRM) is a new product that VMware announced that will help facilitate the protection of VMs in a disaster scenario. Working closely with storage vendors, SRM is a plug-in tool for VirtualCenter that will allow for the creation of protection groups and recovery plans to bring VMs online in a DR site.

VMware Consolidated Backup

VMware Consolidated Backup (VCB) provides a way to backup the VMs in an ESX environment. This is done by mounting them on a proxy server and using backup software to bypass the ESX server and backup the files. This reduces the workload on the actual ESX server.

VMware File System

VMware File System (VMFS) is a high performance cluster file system that ESX uses to host VMs.

EQUALLOGIC ISCSI ARRAYS FOR VIRTUALIZED STORAGE

EqualLogic, a VMware Technology Alliance Partner, produces intelligent, enterprise-class, virtualized iSCSI SANs that are compatible with all VMware products. EqualLogic PS Series SANs bring the same kind of virtualization and dynamic operation to storage that VMware brings to servers. The EqualLogic PS Series storage arrays are easy to deploy, operate, and manage.

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.

VMware + EqualLogic = An Adaptable Virtual Environment

An EqualLogic PS Series SAN is a perfect complement to VMware server virtualization products. Both products virtualize their part of the environment with the intention of reducing TCO through:

- Optimized asset utilization
- Reduced deployment time
- Reduced management complexity
- High availability
- High scalability
- Increased flexibility

Optimized Asset Utilization

In today's datacenter we see that workloads and utilization vary throughout the day and month. Instead of having idle servers sitting around, VMware allows you to optimize your physical assets to make the most use of available resources.

EqualLogic PS Series storage arrays perfectly complement this process by dynamically balancing loads between multiple storage arrays as the usage changes. Additionally you can configure thin-provisioned volumes to grow on demand only when additional storage is needed for those volumes. Thin-provisioning can increase the efficiency of how storage resources are utilized.

Reduced Deployment Time and Effort

Installing physical servers is a time and resource-intensive process, often taking hours or days (or weeks if the hardware must first be purchased). Deploying virtual servers, on the other hand, takes only minutes to hours. You can use predefined templates to provision multiple servers rapidly based on a company standard or build out new servers with just a few clicks.

EqualLogic PS Series arrays similarly reduce the time and effort required to deploy a SAN. Inside the EqualLogic Group Manager, you can easily create new volumes and assign them to your virtual infrastructure and immediately access those volumes.

Reduced Management Complexity

VMware ESX Server with VirtualCenter simplifies server management by reducing the number of physical servers required and providing a single-pane management view of the entire virtual datacenter. With VMware Distributed Resource Scheduler, VMware intelligently balances VM workloads across the entire virtual datacenter without need for human intervention.

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 SAN. The PS Series intelligently balances workloads across the available arrays without need for human intervention, automatically adapting to changes in workload caused by DRS or increased I/O.

High Availability

When virtual machines are stored on a SAN, VMware's High Availability (HA) will detect if an ESX host has failed and dynamically restart the Virtual Machines that were on them on another physical host server that can meet the VM's resource commitments.

PS Series arrays offer best in class reliability, with hot-swappable and fully redundant components, RAID 5/10/50 support and hot spares, along with support for clustering and MPIO.

High Scalability

Scaling of a non-virtualized environment requires significant planning and infrastructure work. In a virtualized environment, scaling is much simpler. You can add new hardware to your virtual infrastructure, and let VMware manage the migration of VMs to make the best use of the new physical resources.

EqualLogic PS Series storage arrays scale in much the same way. As you add more array members, the storage capacity scales along with the performance. Plus the SAN will dynamically balance the load to optimize storage resources.

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 to respond quickly and flexibly to enterprise demands and initiatives.

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

PLANNING CONSIDERATIONS

Boot Options

There are three boot options for ESX server. You can boot from local drives or boot from SAN. With the introduction of ESX 3i there is another option that server vendors will be providing to boot from flash card and customers can use with bootable USB media and other local storage devices.

1. Boot from SAN

In order to configure ESX to boot from SAN you must have an iSCSI HBA that has been approved on the hardware compatibility list (HCL) and has the capability to boot from SAN. Once you configure your HBA settings, you can create a boot volume on the PS Series SAN and point the HBA to that volume. When installing ESX you will choose to use that volume to boot to.

One of the benefits of boot from SAN is that the operating system of ESX resides on the SAN. This provides extra fault tolerance for your server environment. If the physical server goes down, you can point another server to that boot volume and bring up ESX much faster than a full reinstall.

By storing server boot images and production data in a PS Series SAN, you can utilize all of the SAN features such as snapshots and replication on the boot volume. The boot volumes can be replicated to a remote location which will also help in a disaster recovery scenario. These boot volumes can be brought online at the DR site and attached to a server and the environment will be back up and running with minimal configuration change.

Because the boot volume resides on the SAN, this also allows for the rapid deployment of diskless servers and diskless blades into an ESX environment.

2. **Boot from local disk**

Booting from local disk is the easiest and most familiar way to provision an ESX server. There are no configuration changes needed.

3. **3i Hypervisor**

With the introduction of ESX 3.5 and 3i, there is a new hypervisor. It only takes up 32MB of space so server vendors are releasing servers that have a flash card with 3i already built into the system. When the server is delivered, plug it in and power it up and ESX will be running after a few configuration entries. VMware will also be providing this hypervisor and it can be put onto something bootable such as a bootable USB flash drive. The HCL for ESX 3i is much more limited than the standard VI3.

Server and Storage Connectivity

There are multiple ways that you can attach a virtual infrastructure to an EqualLogic PS series SAN. ESX 3.x has two ways to connect to volumes on the SAN. You can use the ESX iSCSI software initiator or an approved HCL iSCSI NIC or HBA. In addition, VM's can utilize the native guest OS iSCSI software initiator to connect to an EqualLogic PS series SAN.

Typically the hardware iSCSI HBA has less CPU overhead because more processing is performed by the HBA card. The requirement is an approved HBA card and the physical resources needed for it such as space in the server, correct PCI interface type etc. As long as the HBA supports it, the server can be set up for boot from SAN.

Software iSCSI initiators will run over standard network interface cards which typically all newer servers have onboard. Because the software iSCSI initiator requires ESX to be up and running, it cannot be used for boot from SAN.

Volume Connectivity and Data Formats

There are three ways in which to access the volumes residing on a PS Series SAN. ESX can see a VMFS volume and a Raw Device Map (RDM) volume. It is also possible to utilize the guest operating system of the VM to access the volumes.

1. **VMFS**

The first and most common access method is a VMFS volume. This is a proprietary disk format that ESX uses to allow multiple VMs to reside on and multiple servers to access a single volume. VMFS has some built in functionality to allow for taking VMware software snapshots, file locking for multiple host access, as well as being able to place multiple VMs on a single volume. Each '*disk*' is encapsulated in a single .vmdk file. VMFS has a maximum size of 2TB but can use extends to increase the size of this.

2. **RDM**

The second alternative to access storage is a Raw Device Map (RDM). A RDM is a volume that is not formatted in VMFS but instead is presented to the VM and the VM formats it in a way that it sees fit (NTFS, ext3, GPT, etc). However, because ESX is actually handling this volume, the VM treats it as a local disk. There are two types of modes for a RDM: Virtual Compatibility Mode and Physical Compatibility Mode.

- a) In Virtual Compatibility Mode, ESX virtualizes much of the volume so that the guest OS VM sees it as a VMware device. This allows for VMware to take software snapshots. Snapshots are available and the LUN is removed from the list of available LUNs for RDM. A downside is that the volume size is encoded in the virtualization file. EqualLogic Volume resizes are not seen by the VM. In order to see new volume sizes you must power down the VM, remove the RDM and then re-add it to the VM to see the correct size.
- b) In Physical Compatibility Mode, VMware does not virtualize the volume and it allows for the device information and all SCSI commands to pass through to it. VMware cannot do snapshots on a Physical RDM. One of the downsides is that the LUN is not removed from the list of RDM LUNs so you must be diligent in LUN assignments.

3. Guest OS iSCSI Initiator

The third way to access volumes on the SAN is done by using the native Guest OS's iSCSI software initiator. Any OS that has iSCSI support will be able to use their native initiator to connect to volumes on the SAN. The benefit of this is that you can utilize the operating system's multipath iSCSI protocols (MPIO) as well as any native software designs. For example, for windows 2003 EqualLogic provides a MPIO device specific module (DSM) that can improve the performance of native windows MPIO when connected with the PS Series SAN. Another huge advantage of using guest OS iSCSI initiators is the ability to use native tools such as Windows VSS for application aware backups and snapshots. It is also possible to use the EqualLogic Auto-Snapshot Manager to take application consistent snapshots of the data residing on volumes accessed using the guest OS's iSCSI initiator.

Volume size considerations

There are no hard and fast rules for volume size considerations. Depending on the environment you may decide to do multiple 500GB volumes with multiple VMs or smaller volumes dedicated to single VMs. One thing to consider is the volume on which the VM resides is also used as the Guest OS's memory cache. So for example, if you have a 12GB windows C drive and have allocated 1GB memory, you need a little over 13GB minimum of space for the file system, the swap files, and miscellaneous VMware files.

With the addition of Storage VMotion with ESX 3.5/3i, you will be able to move VMs from one volume to another while the VM is up and running. If you are not using 3.5/3i, you can still migrate VMs from one datastore to another but the VM must be powered down first.

Storage Array RAID Levels

You should consider the workload in your virtual environment to determine which RAID level to use for the PS Series storage. In general, RAID 10 provides the best performance, at the expense of storage capacity. RAID 50 generally provides more usable storage, but has less performance than RAID 10 in random I/O situations and requires additional overhead in the case of a drive failure scenario. RAID 5 provides the most storage capacity at the expense of a little lower performance and availability.

Server and Storage Redundancy

You want to make sure your servers are enterprise class servers that are approved by VMware along with the normal raid, dual power, and other things you normally include in an enterprise deployment.

PS Series SANs offer best in class reliability, with hot-swappable and fully redundant components such as fans, controllers, and power supplies. SANs also offer data protection support such as RAID 5/10/50 and hot spares, along with support for clustering and MPIO.

With VMware HA enabled, if there is a failure of an ESX server, all of the VMs that were on that server will be moved and restarted on other ESX servers. HA clustering constantly monitors the health of each ESX server. When an ESX server goes down, VirtualCenter is responsible for bringing the VMs that the failed ESX server was controlling back online on other ESX servers. Additionally, you can use native guest OS clustering software to provide an even higher level of redundancy at the VM level. By clustering VMs across different ESX servers, even if an ESX server fails, the other node of the cluster will be up and running and in the meantime, HA will online the VM that failed onto another ESX server.

MPIO

Multipath I/O is a concept to provide multiple paths to the data for load balancing as well as fault tolerance. This can be achieved in a few different ways but usually involves configuring ESX as well as the physical infrastructure to support more than one path. Multipath I/O can be used to utilize all the network and storage resources evenly. The added benefit is that now there is an inherent fault tolerance built in, if a path fails all access to the data is not cut off.

1. MPIO with ESX 3.x iSCSI HBAs

ESX supports hardware iSCSI HBA MPIO in active/standby mode. If the active path to the data fails the standby path will come up and data communication will continue with no interruption of service. It is recommended that each HBA is attached to a different physical switch so there is no single point of failure to the data. By configuring multiple volumes on different HBAs, you can achieve traffic on all available cards. Figure 5 shows one example of MPIO using HBAs.

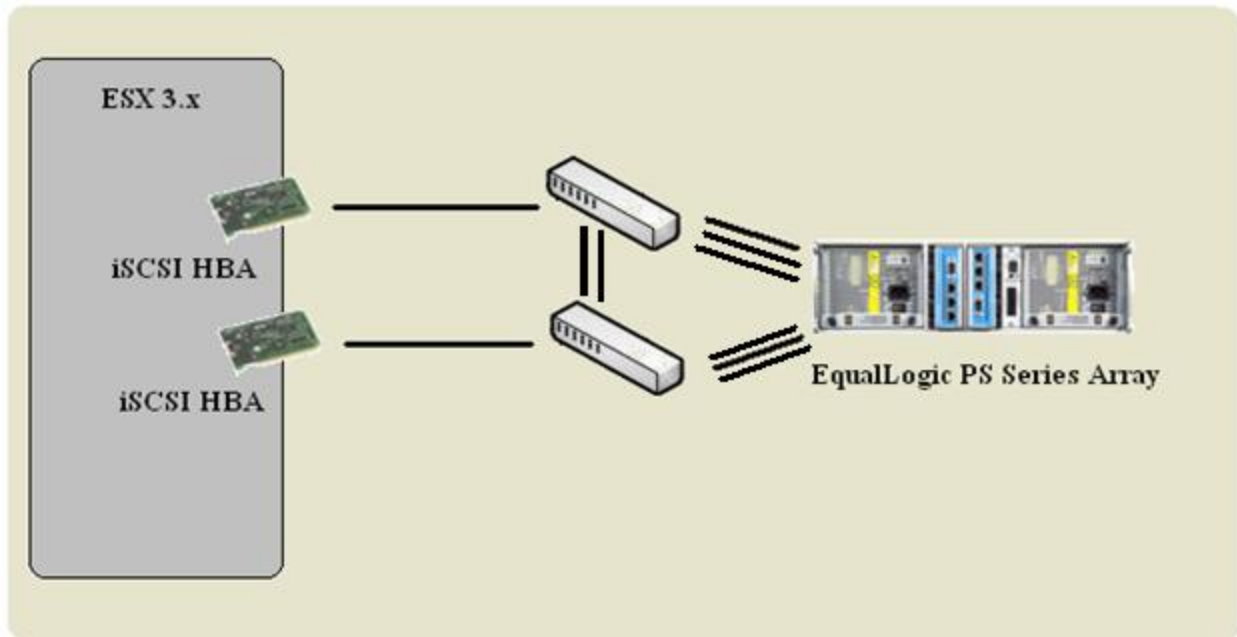


Figure 5: MPIO with iSCSI HBAs

2. MPIO with ESX 3.x iSCSI Software Initiator

ESX native software iSCSI initiator supports MPIO by using the network NIC teaming inside the networking properties. This can be set up a few different ways such as route based on MAC source hash or based on IP hash. You can increase the number of network interfaces to increase the number of paths to the SAN. Each connection however will still only utilize one physical port so the load balancing comes from the entire system not an individual volume or VM. Again, having different NICs on different switches will prevent a single point of failure. Figure 6 shows an example of this.

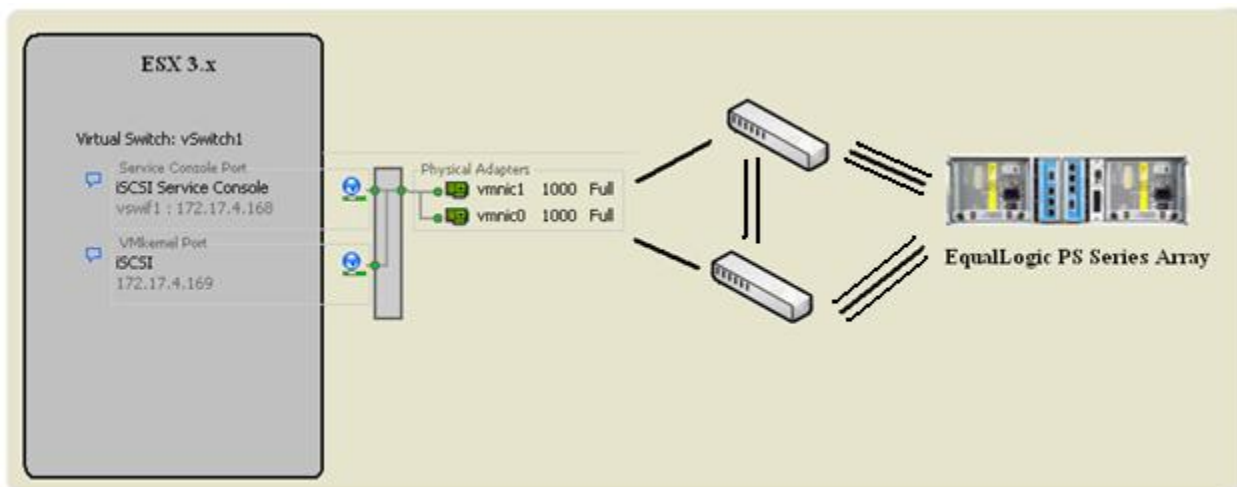


Figure 6: MPIO with ESX 3.x iSCSI Software Initiator

3. MPIO with ESX 3.x using Guest OS Software Initiator

Some operating systems have their own native iSCSI software initiators that can be used to achieve multipath I/O and redundancy to the SAN. To establish this you create multiple virtual switches using the

physical vmnic ports and assign those to the VM. Then inside the guest OS you configure the multipath and iSCSI initiator to talk to the SAN. For example, in Windows 2003 you can assign 2 virtual switches to the VM along with any production network traffic. Each of those NIC ports communicate with the storage. The guest OS will see each of these virtual switches as another network interface. You can then install Microsoft MPIO and the iSCSI software initiator and attach to the storage from the VM. Figure 7 shows one way to configure your guests to see multiple paths.




Figure 7: MPIO using ESX Networking for Guest OS iSCSI

MPIO with ESX 3.x Networking

Just as you want to enable MPIO for storage access, the same concepts can be applied to the data network setup as well. Using multiple virtual switches you can use the guest OS NIC teaming software or you can add multiple NICs at the ESX level to an individual virtual switch to provide fault tolerance and load balancing.

Data Protection

Moving to a virtual infrastructure does not necessarily change the way you protect your data. Backups, replication, fault tolerance, DR, and other means of protecting your data still apply.

You need a plan for backing up ESX Servers, boot volumes, virtual machines and their data. Standard snapshots can be used to back up an ESX Server if it resides on the SAN as well as virtual machine data. Or VMware Consolidated Backup can be used in conjunction with traditional backup software to provide quiesced backups of virtual machines stored on the SAN. VMware Consolidated backup is a serverless backup utility that mounts the volumes of an ESX server onto a Windows proxy server for backup thus limiting the amount of overhead that the ESX server sees during the backup process.

If your standard backup policy is to install agents on the servers and back them up using a remote network backup server you can continue to use these techniques in a virtual environment. You can use some of the SAN features such as snapshots and replication along with classic backup methods to provide multiple layers of data protection.

Replication

Because of the encapsulation that VMware provides, replication to a DR site becomes even easier. With the volumes being replicated, it is easy to break the replication on the DR site, bring the volume online to an ESX server sitting in the DR environment, browse the datastore and bring VMs back online. This entire process takes very little time and decreases the time to recovery in the event of a disaster.

With ESX 3.5/3i, VMware is introducing a new product named VMware Site Recovery Manager. This is a plug-in that storage vendors will use to help automate the DR recovery process. This includes the replication of certain datastores that contain VMs as well as the breaking of the replication and bringing VMs online in the DR site.

Snapshots

There are two types of snapshots that can be used in an ESX and PS Series SAN environment. ESX software-based snapshots and PS Series SAN-based snapshots.

When taking a VMware snapshot, ESX puts the VM into a snapshot state so that it can track changes. It then writes these changes to the volume so that you can roll back in the case of a failure or if you need to go back to a point-in-time. Depending on how many software snapshots are being run this can lead to running out of space as well as performance impacts on the ESX host server.

When using the PS Series SAN based snapshot technology, you utilize the performance and intelligence of the PS Series SAN to take point-in-time copies of the entire volume and all VM's running on that volume. This option uses snapshot reserve space on the SAN rather than space inside the VMFS volume. This also provides the added benefit of being able to connect that snapshot to the ESX server, browse the datastore and bring VMs online either in a test environment or as a new VM.

Templates and Rapid Provisioning

Templates allow for rapid provisioning of new virtual servers. You can have a gold master image and deploy new servers with all the patches and software you need. Depending on the OS you may have to run tools like sysprep or new sid. This allows for rapid deployment of new servers as well as creating test environments.

DEPLOYMENT

VirtualCenter

The process of building a virtual infrastructure with VMware ESX begins with VirtualCenter. VirtualCenter will be the central management of all of the ESX hosts, the VMs, and configurations.

VMware recommends that VirtualCenter reside on its own server. If you are using multiple ESX Servers and managing them with VirtualCenter, you will need to install and configure a SQL Server or Oracle database to manage the VirtualCenter server information database. Follow the procedures in the *Preparing the VirtualCenter Server Database* section of the VMware Virtual Infrastructure *Installation and Upgrade Guide* to prepare the database.

Note: For large installations, you can also install an optional VMware License Server.

Install the VirtualCenter software on the designated server using the VirtualCenter CD. Consult the *Installing the VirtualCenter Server* section of the *Installation and Upgrade Guide for VMware Virtual Infrastructure* if you need more information.

EqualLogic PS Series SAN

Configure the EqualLogic PS Series SAN. Consult the latest PS Series Quickstart Guide for more information.

Deploy ESX 3.x

Depending on the boot method, install ESX 3.x into the environment you have chosen. See Appendix A for one example of an ESX install.

Once ESX is up and running, use VirtualCenter to add it to the inventory as well as configure the license type. From VirtualCenter you can configure the networking as well as the iSCSI connections.

Connect

After your ESX environment is up and running, use the connection methods discussed in this report to connect to the volumes on the SAN. Appendix A has one example of using the ESX software iSCSI initiator and Appendix B has an example of Multipath I/O with an iSCSI HBA.

Provision

Once you have volumes attached to your ESX environment you can start with provisioning new Virtual Machines. This can be done in a number of ways. Using VMware converter tool you can take existing physical servers and convert them to virtual servers on the ESX server. You can also create new Virtual Machines from VirtualCenter. Once you have created the Virtual Machine you follow all the standard steps for installing the operating system and configuring the VM to your needs. You can also convert VMs that were created in VMware Server into the ESX environment.

Utilizing VirtualCenter existing VMs can be converted to templates which will allow for more rapid server deployment.


CONCLUSION

Virtual infrastructures built with VMware ESX Server and PS Series SANs can help lower cost and increase functionality. Combined, the technologies create a dynamic, scalable and highly available environment that works hand in hand to meet your current and future business needs.


APPENDIX A: ACCESSING A PS SERIES SAN USING THE ESX SOFTWARE ISCSI INITIATOR

One of the ways to connect ESX to a PS Series SAN is through the ESX native iSCSI software initiator. The following section describes how to configure ESX, configure the SAN, connect and access the volume.

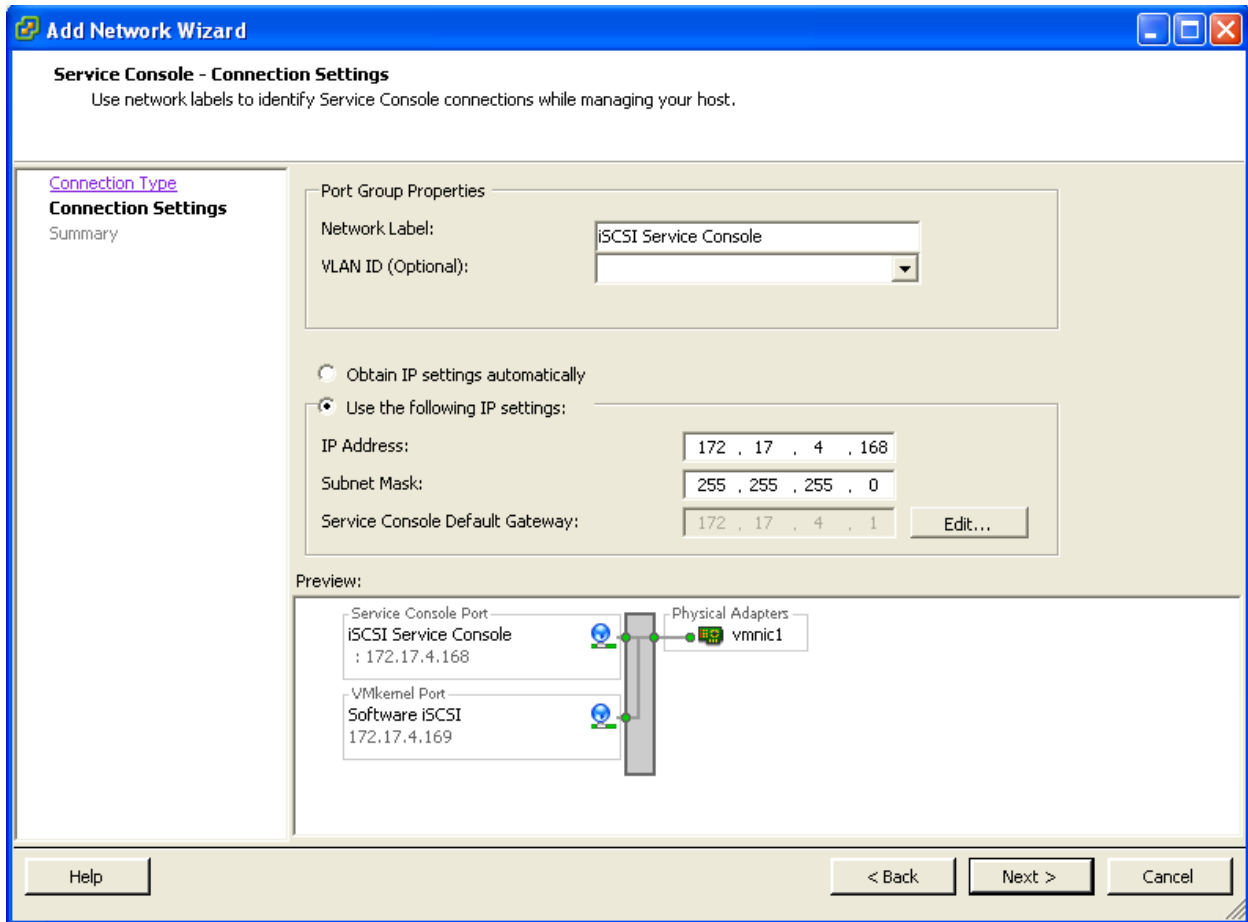
The first thing we need to do is configure the ESX iSCSI software initiator. Open up the **Configuration** tab and select **Add Networking**. Then add a new VMkernel connection.



Choose as many physical NICs as you want for your connection to the SAN. When this is done configure the IP address for the VMkernel.

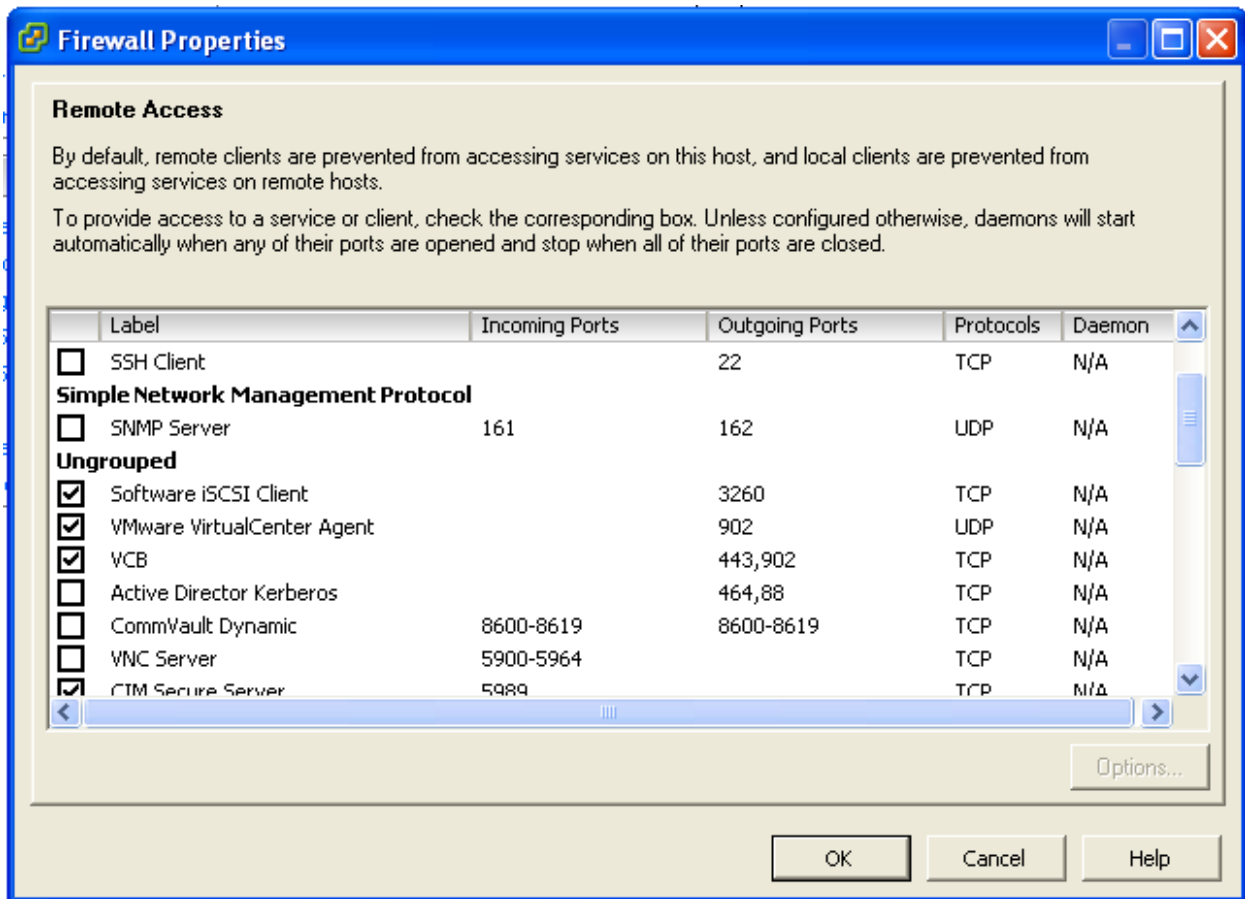


After this we need to add an iSCSI Service Console to that new network. Select **Properties** of that new Network and choose the **Add** button. This time select a Service Console. Configure the name and IP Address of the iSCSI service console. Remember that this is only for iSCSI connectivity and not to be used in place of the default service console.



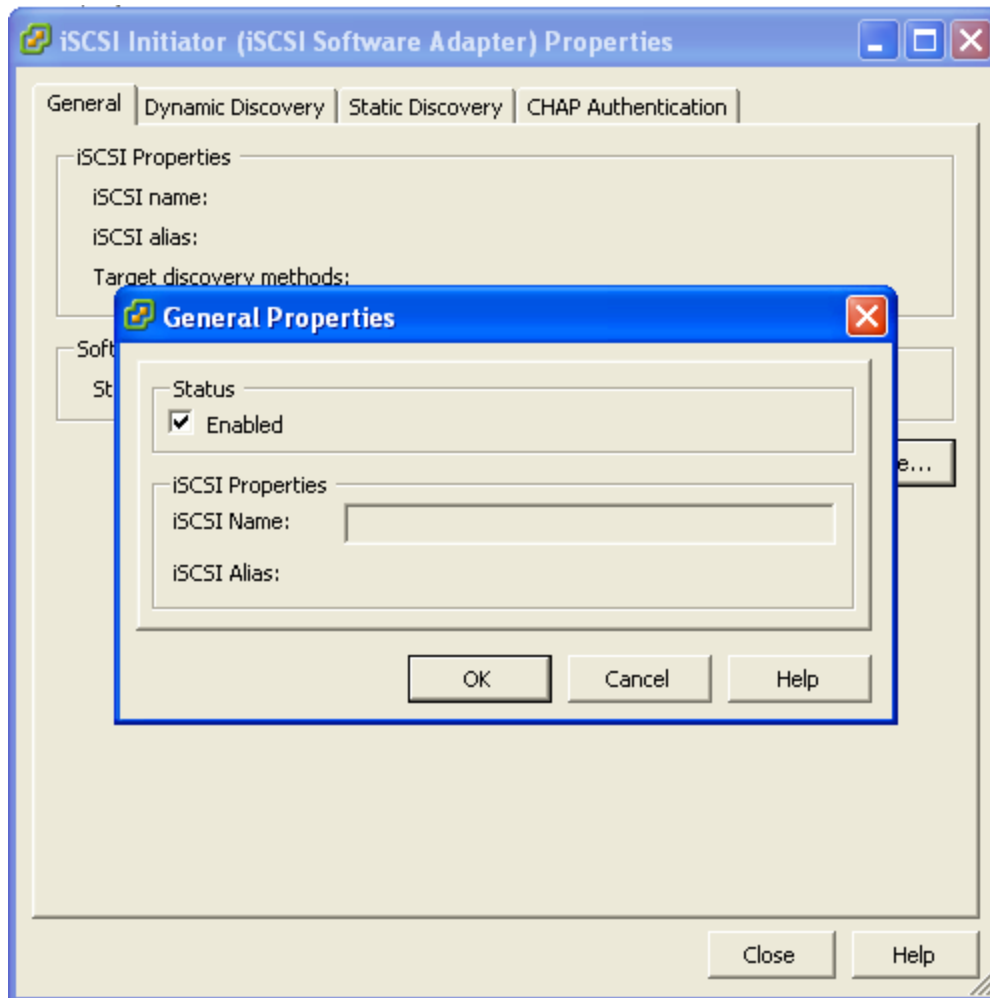
When this is finished you will have a new Network with however many physical adapters you chose, a VMkernel with an IP address dedicated to iSCSI traffic and an iSCSI service console.

The next thing to do is to configure the VMware ESX Firewall to allow iSCSI traffic out. Click the **Configuration** tab, and then **Security Profile**. Click the **Properties** next to the firewall settings. Find the software iSCSI Client and make sure it is checked.

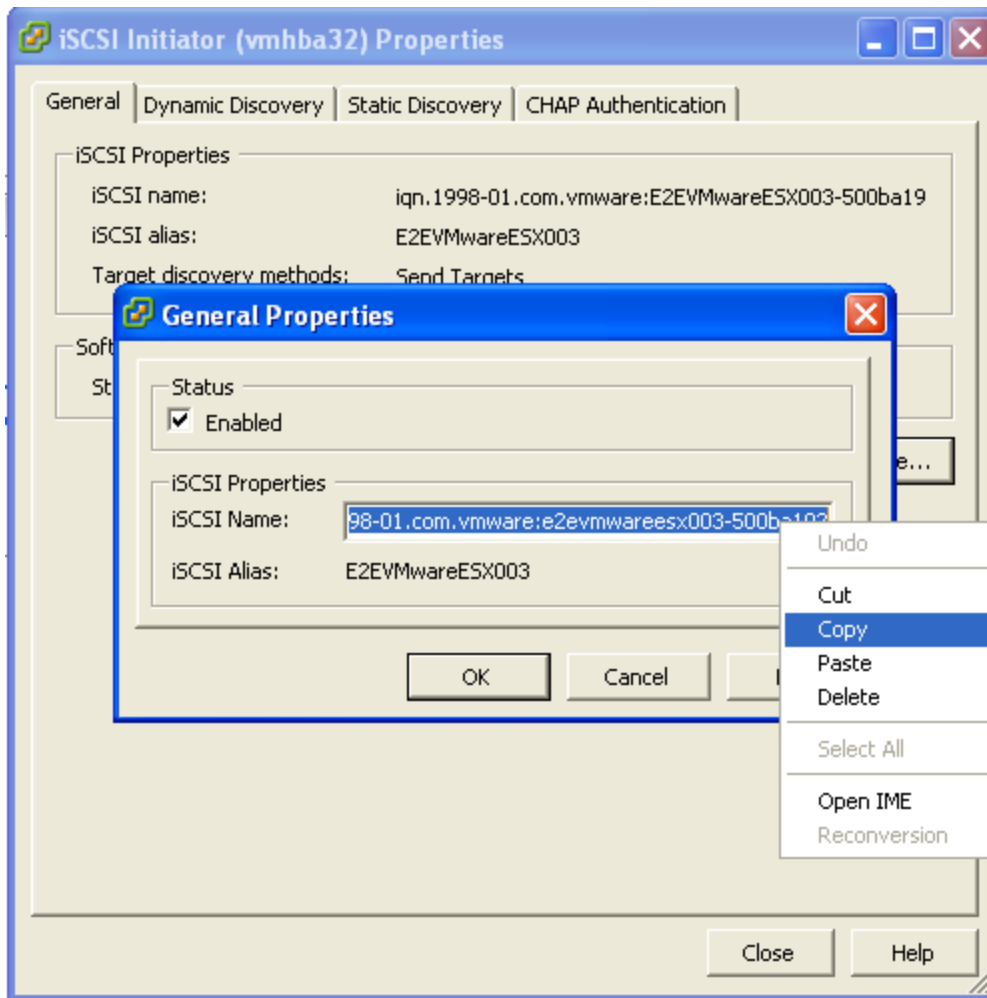


You may need to reboot your ESX host after completing this. You will see a message if you need to. After that click the **Configuration** tab and select **Storage Adapters**. In your list of storage adapters you will see the iSCSI Software Adapter. Highlight this and click **Properties**.


On this first screen click the **Configure** button and put a checkmark in the Enabled field.



Once this is completed you can click **Configure** again and copy the iqn name of the software iSCSI connection.



Using this information you can add this server to the access control list of the volume on the PS Series SAN. Open up group manager, select the volume you are connecting to and modify the Access settings.



When this is done, on the ESX server click the **Configuration** tab then click **Storage Adapters**. Select the iSCSI Software Adapter and hit the **Rescan** button. When the rescan is finished you should see the volume in the target list.

Storage Adapters

Rescan...

Device	Type	SAN Identifier
vmhba1	iSCSI	iqn.2000-04.com.qlogic:ql...
vmhba2	iSCSI	iqn.2000-04.com.qlogic:ql...
Smart Array P400i		
vmhba0	Block SCSI	
iSCSI Software Adapter		
vmhba32	iSCSI	iqn.1998-01.com.vmware:...

Details

vmhba32 [Properties...](#)

Model: iSCSI Software Adapter IP Address:

iSCSI Name: iqn.1998-01.com.vmware:e2evmwareesx003-500ba192 Discovery Methods: Send Targets

iSCSI Alias: E2EVMwareESX003 Targets: 1

SCSI Target 1

iSCSI Name: iqn.2001-05.com.equallogic:0-8a0906-98bbf1102-8eb598bd966478e7-esxvolume0


iSCSI Alias: ESXVolume0

Target LUNs: 1 [Hide LUNs](#)

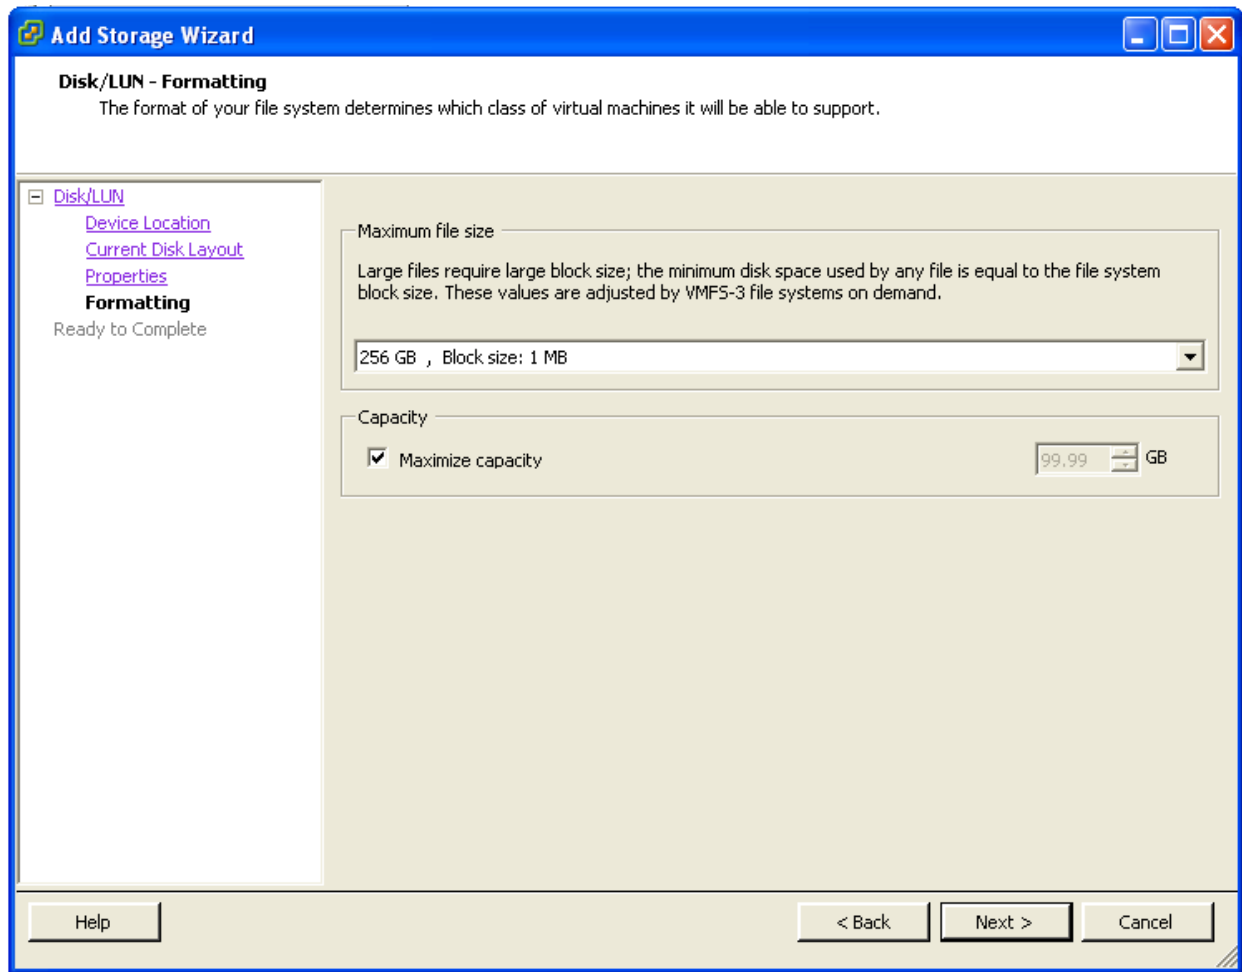
Path	Canonical Path	Type	Capacity	LUN ID
vmhba32:1:0	vmhba32:1:0	disk	100.00 GB	0

Once the volume shows up in the target list the last thing we have to do is format it so that ESX can use it for VMs. The next step shows how to format the volume for VMFS.

Click on the **Configuration** tab and then **Storage**. Click **Add Storage**. Choose to add a Disk/LUN and you should now see that volume available. Select the volume and give the Datastore a name. It is recommended to assign the same name as the EqualLogic Volume so that you can easily tell which datastore relates to which volume.




After this you are shown a screen in which to format the volume. One thing to keep in mind is that the block size determines the maximum size of an individual .vmdk file regardless of the size of the volume. If you choose the default block size of 1MB, your biggest .vmdk file can only be 256GB. So if you are planning on a 500GB data volume on vmfs you must change the size of the block size accordingly.



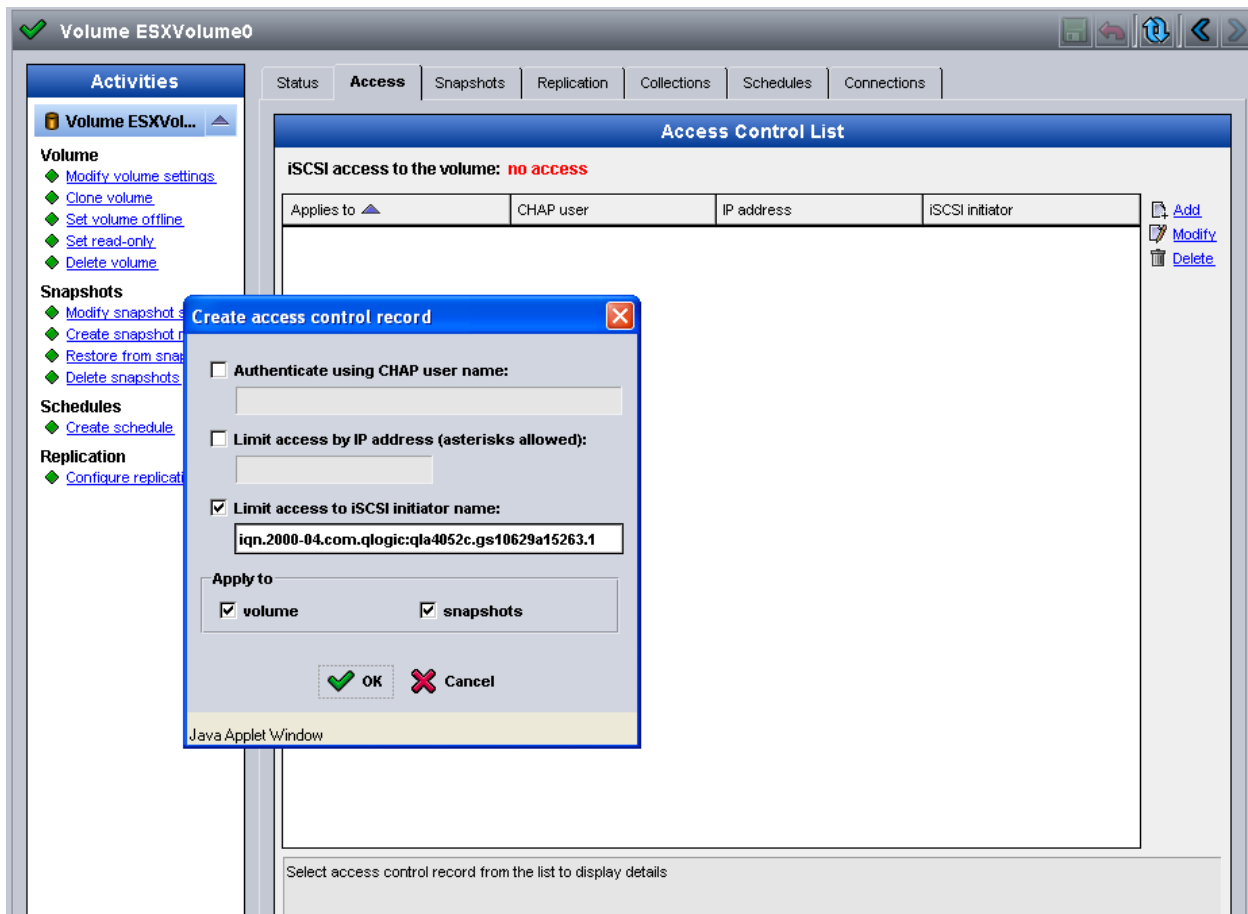
When this is done it will show up in the available Storage pool and allow you to begin using it for VM storage.

APPENDIX B: ACCESSING A PS SERIES SAN USING AN HBA HARDWARE INITIATOR


This section will explain the process of connecting your ESX host to a PS Series SAN using approved hardware HBAs. You must first configure an IP address for each HBA port you are using. To do this, in Virtual Center Client click on the ESX host, then click the **Configuration** tab and then click **Storage Adapters**. Select the storage adapter you want to configure and click **Properties**. Click **Configure** and then you can configure the IP address of the HBA.



From here you can also copy the iSCSI iqname of the HBA initiator to add to the volume's access control list. Open up Group Manager and go to the Access tab of the volume. Here you can add the iqname of the HBA to the access tab.



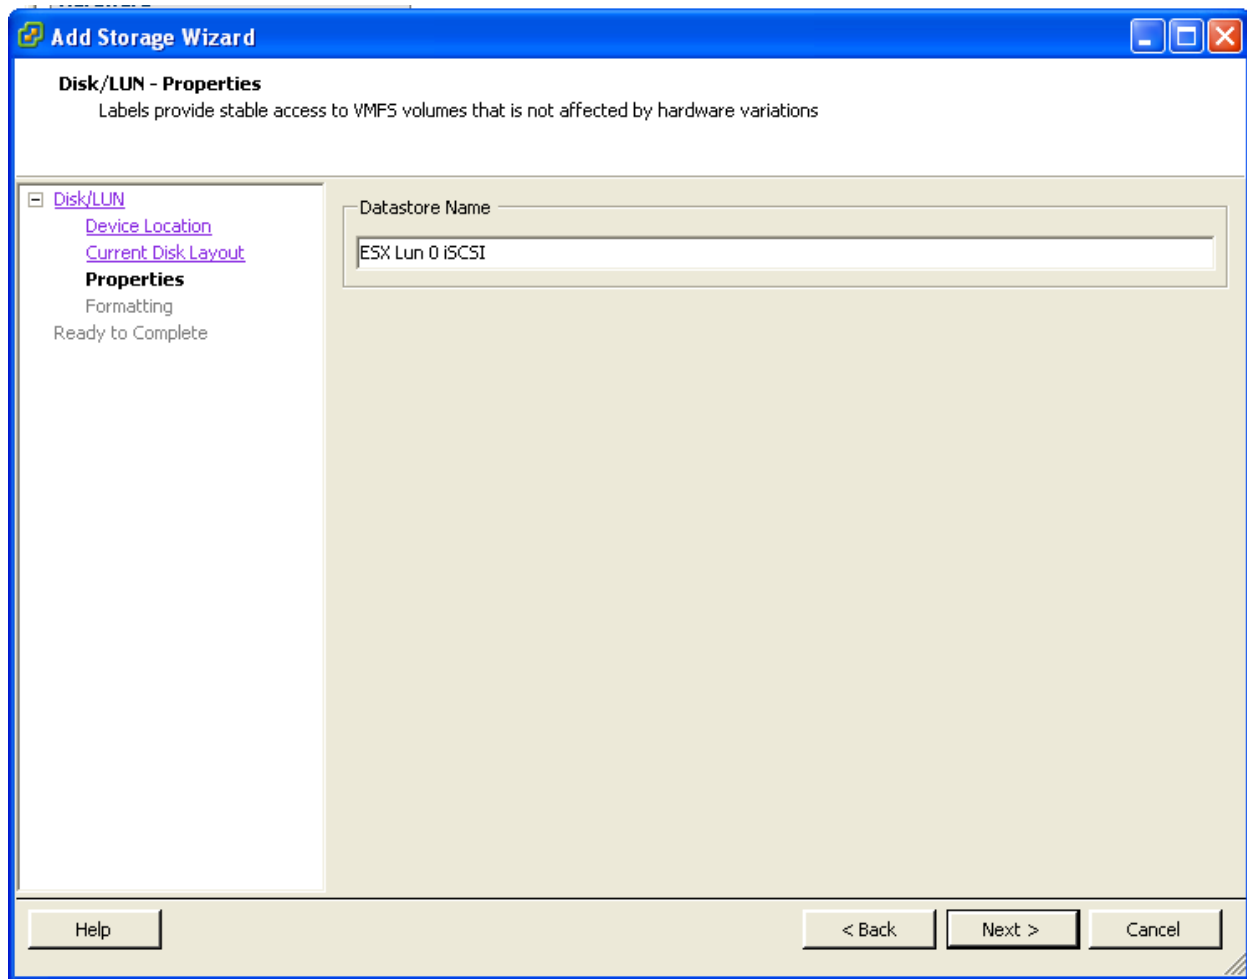
Do this for each HBA you have. When this is done go to the **Configuration** tab in ESX. Click on **Storage Adapters** and select the adapter and click **Properties**. Click the **Dynamic Discovery** tab and type in the IP address of the PS Series group.




When that is done, do the same for each HBA in the system. After that click on one of the HBAs and click **Rescan**. You should see the volume you created show up in the Target list.

Once the volume shows up in the target list the last thing we have to do is format it so that ESX can use it for VMs. The next step shows how to format the volume for VMFS.

Click on the **Configuration** tab and then **Storage**. Click **Add Storage**. Choose to add a Disk/LUN and you should now see that volume available. Select the volume and give the Datastore a name. It is recommended to assign the same name as the EqualLogic Volume so that you can easily tell which datastore relates to which volume.



After this you are shown a screen in which to format the volume. One thing to keep in mind is that the block size determines the maximum size of an individual .vmdk file regardless of the size of the volume. If you choose the default block size of 1MB, your biggest .vmdk file can only be 256GB. So if you are planning on a 500GB data volume on vmfs you must change the size of the block size accordingly.



When this is done it will show up in the available Storage pool and allow you to begin using it for VM storage.

EQUALLOGIC DOCUMENTATION AND CUSTOMER SUPPORT

Visit the EqualLogic Customer Service website, where you can download the latest documentation and firmware. You can also view FAQs, the Knowledge Base, and Technical Reports and submit a service request.

EqualLogic PS Series storage array documentation includes the following:

- *Release Notes*. Provides the latest information about PS Series storage arrays and groups.
- *Quick Start*. Describes how to set up the hardware and start using a PS Series storage array.
- *Group Administration*. Describes how to use the Group Manager GUI to manage a PS Series group. This manual provides comprehensive information about product concepts and procedures.
- *CLI Reference*. Describes how to use the Group Manager command line interface to manage a group and individual arrays.
- *Hardware Maintenance*. Provides information on maintaining the PS Series storage array hardware.

You can also find the latest version of the VMware FAQ online. This document is specific to VMware and helps answer questions and configuration issues specific to VMware in an EqualLogic PS Series SAN.

To access the Customer Support website, go to <https://www.equallogic.com/support/> and click LOGIN to log in to your support account. If you do not have an account, you can request one on this page.

If the issue is urgent, please call us at 1-877-887-7337 (toll free US & Canada) or 919-767-5729 to speak with a member of the customer support team.

If you have any comments or suggestions related to this technical report, please send them to techreports@equallogic.com