Deploying iSCSI Storage Solutions on the Microsoft® Windows Server™ Platform

Dennis Martin

Microsoft MVP for Windows Server Storage
President, Demartek

Published: June 2007

Abstract

In an effort to provide an improved experience for the growing iSCSI storage market, Microsoft and its storage partners have created various iSCSI target solutions for the Microsoft Windows Server environment. These solutions include iSCSI targets based on Microsoft's iSCSI target software and storage partner hardware, and iSCSI targets based on storage partner iSCSI target software and hardware. This paper provides an update on the state of iSCSI storage technology, specifics on several Microsoft storage partner solutions, including the deployment of each solution for specific solutions such as Microsoft Exchange, Microsoft SQL Server, Microsoft SharePoint Server and Microsoft Cluster Server (MSCS).

This report is designed for managers of IT departments and system administrators who are exploring the possible benefits of iSCSI storage solutions or who are looking for actual deployment examples of iSCSI storage solutions.

Contents

Introduction

Internet SCSI (iSCSI) is an industry standard developed to enable transmission of SCSI block storage commands and data over an existing IP network by using the TCP/IP protocol. The encapsulated SCSI commands and data can be transmitted over a local area network (LAN) or a wide area network (WAN). As with traditional SCSI, an iSCSI storage solution requires at least one “initiator” residing on the application server and at least one “target” residing on the storage.

This report provides background on iSCSI technology and information on the current state of iSCSI storage solutions for the Microsoft Windows environment, focusing on the deployment of the iSCSI target solutions from some Microsoft storage partners. Some of these storage solutions are based on the Microsoft iSCSI software target and run on a Microsoft Windows-based platform. Some of these storage solutions are based on the storage partner’s own technology and run on a non-Microsoft platform. All these storage solutions provide storage for hosts that use the free-of-charge Microsoft iSCSI initiator.

Storage Solutions Deployed

The following iSCSI storage solutions were deployed for this report.

- Dell™ PowerVault™ NX1950 Networked Storage Solution
- EqualLogic® PS3800XV
- HDS TagmaStore™ AMS1000
- HP StorageWorks 1200 All-in-One Storage System
- LeftHand Networks® SAN/iQ®

The Dell solution is based on Microsoft Windows Unified Data Storage Server 2003. The HP solution is based on Microsoft Windows Storage Server 2003 R2 and the Microsoft Software iSCSI Target application pack. The EqualLogic, HDS and LeftHand Networks solutions are based on their own respective technology.

These iSCSI target storage solutions provide a variety of advanced storage features including hardware RAID, Multi-path I/O (MPIO), snapshot copy, replication, remote copy and others. Some include integration with Microsoft Volume Shadow Copy (VSS) and provide Microsoft Virtual Disk Service (VDS) providers.

Basic Storage Architectures

The two basic forms of storage for host computers are direct-attached storage (DAS) and networked storage. DAS is storage that is directly attached to a host computer and is generally privately owned by that computer. Networked storage is storage that is connected to a host computer via some sort of network, such as an Ethernet network or Fibre Channel network, and can take several forms, including variations of Storage Area Networks (SAN) and Network Attached Storage (NAS).

There are two basic forms of networked storage: the Storage Area Network (SAN) and Network-Attached Storage (NAS) and they are generally distinguished by their Input/Output (I/O) characteristics. SAN is generally used for applications that require “block” I/O access. NAS is generally used for applications that require “file” I/O access. An application that uses block I/O is any application that reads or writes its data blocks directly to the storage device or subsystem, such as a databases, email servers, or file systems themselves (such as NTFS). SAN storage devices appear to the applications the same way that DAS devices do, allowing applications to use what appear to be local storage devices. An application that uses file I/O is one that makes its read and write requests in the form of files, such as a network client reading and writing files from a file server. NAS devices typically appear as one or more network file shares to the applications and users. NAS devices are actually host servers themselves that internally use a DAS or SAN I/O connection, but share a “file-system” type of view of their storage resources to other hosts on the network.

The primary reasons for using any form of networked storage are to overcome the disadvantages of the DAS storage model. The various implementations of networked storage can allow the storage to be located potentially many miles from the host CPU requiring the storage and can scale to hundreds, thousands or even millions of storage devices. In addition, the networked storage model allows storage to be placed into a “pool” that is not necessarily owned by any one application client or server but can be shared among many applications or servers.

Direct-Attached Storage (DAS)

DAS is probably the most well-known form of computer storage. In a DAS implementation, the host computer has a private and usually exclusive channel between the host CPU and the storage device or devices, so that the host “owns” the storage. In this context, DAS storage has also been called “server-centric” or “silo” storage. Everything from personal computers to mainframe computers have used this implementation. Over the years, various interfaces have been used for this purpose, including IDE/ATA and SCSI. The advantage of DAS is that it is relatively simple to understand and implement. The disadvantages are that there is limitation to the number of devices that can be connected on the same interface, a relatively short distance between the host CPU and the storage device(s) due to cable length restrictions, and when larger storage devices are required the data must often be moved from the smaller device to the larger device, potentially consuming a large amount of time. In addition, many DAS storage architectures require that the host computer be taken offline when adding or removing storage devices. The DAS model doesn’t scale to large or distant environments very well. A server in this model doesn’t directly share its storage resources.

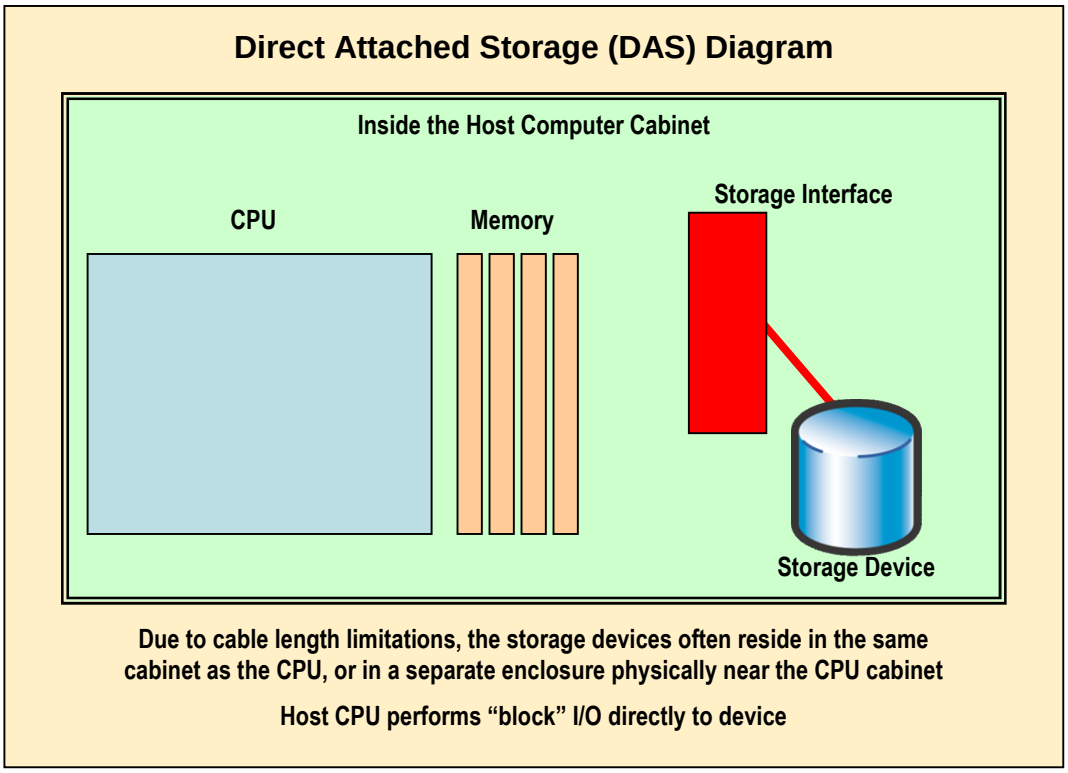

Network-Attached Storage (NAS)

NAS devices share their storage resources with other clients on the network, in the form of file "shares." The clients read and write files on the NAS server using either SMB/CIFS or NFS file protocols. The NAS device has its own storage and internally uses block I/O to store the data in its own internal format. NAS devices typically can have many file "shares" and can potentially be a great physical distance from the network clients.

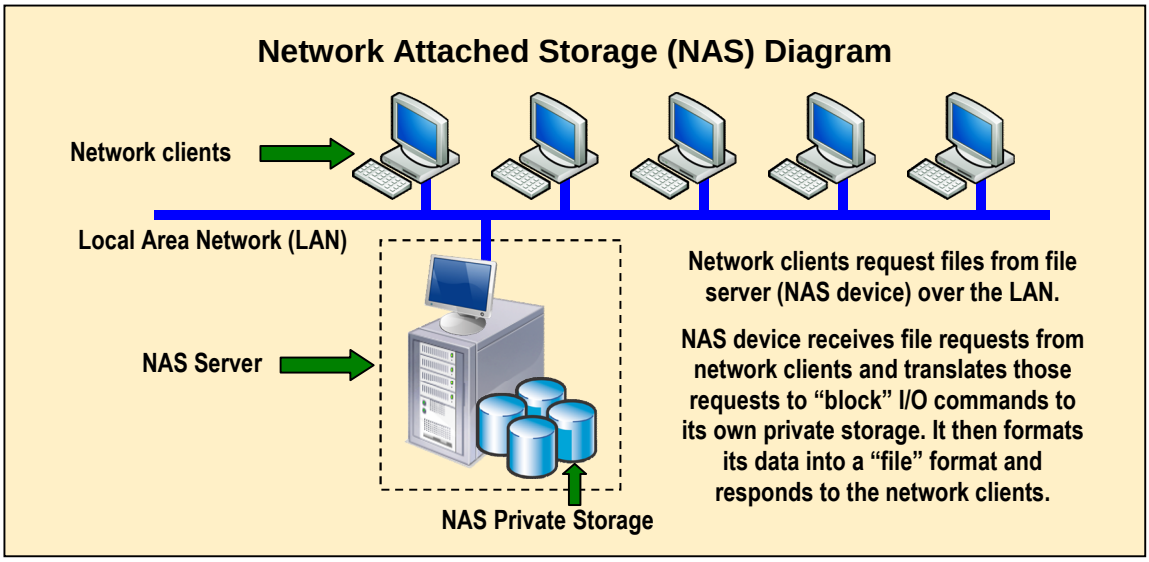

Storage Area Network (SAN)

SAN architecture, using “block” I/O, can be implemented over an Ethernet network or a Fibre Channel (FC) network, each having its own strengths and weaknesses.

Ethernet networks are ubiquitous, relatively inexpensive, and they offer a wide variety of choices of built-in, as well as peripheral network hardware solutions. The least expensive solutions use copper cabling and connections. Ethernet networks using TCP/IP protocol typically manage network traffic with software and frequent interrupts to the host computer.

By comparison, Fibre Channel networks tend to be more expensive, because they use dedicated fiber-optic technology, they are physically separate from the local Ethernet network, and they require specialized expertise. FC host bus adapters (HBAs), managed switches, and optical networks are dedicated to block level storage I/O. FC networks manage traffic using highly efficient hardware processing that offloads functionality from the host CPU.

SAN architectures use block SCSI protocol for sending and receiving storage data over their respective networks. Fibre Channel (FC) SANs implement the SCSI protocol within the FC frames. Internet SCSI (iSCSI) SANs implement the same SCSI protocol within TCP/IP packets. Because both technologies allow applications to access storage using the same SCSI command protocol, it is possible to use both technologies in the same enterprise, or move from one to the other.

For the larger enterprises that have implemented SAN technology, most have implemented Fibre Channel technology. These enterprises typically demand proven technology, have the need for high bandwidth storage solutions, have the budgets to pay for more expensive hardware to meet their performance and reliability needs, and typically have full-time staff dedicated to storage management.

Some organizations have implemented only Fibre Channel SANs into production use. Some organizations have implemented only iSCSI SANs into production use. Some organizations have chosen to implement Fibre Channel SANs and iSCSI SANs. An iSCSI connection can be used to bridge a server into a Fibre Channel SAN. In many cases, iSCSI SANs can be deployed more quickly than Fibre Channel SANs.

Many disk storage solutions available today offer both Fiber Channel and iSCSI interfaces into the same disk subsystem. Many iSCSI-only solutions use the same high-end components in the disk subsystem as the Fibre Channel-only solutions. These design and implementation factors reduce or eliminate concerns about performance and reliability of iSCSI storage solutions.

IP-based SAN technology such as iSCSI has not yet been deployed as widely as Fibre Channel SAN technology, in part, because it is a newer technology. But iSCSI SAN technology has been proven to work reliably, provide excellent performance, and is a cost-effective choice for many storage environments.

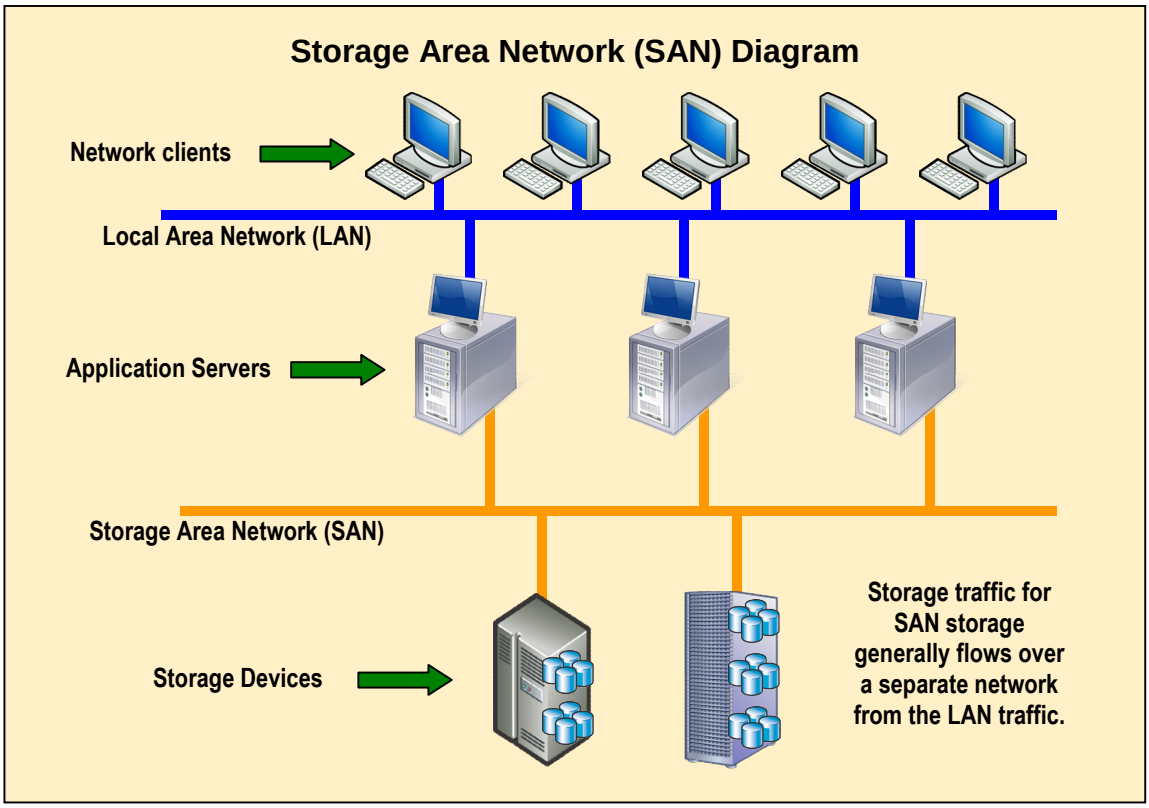

A typical SAN environment consists of one or more application servers, one or more storage devices, a switch between the servers and the storage, the appropriate external hardware interface in the application server, and appropriate cabling.

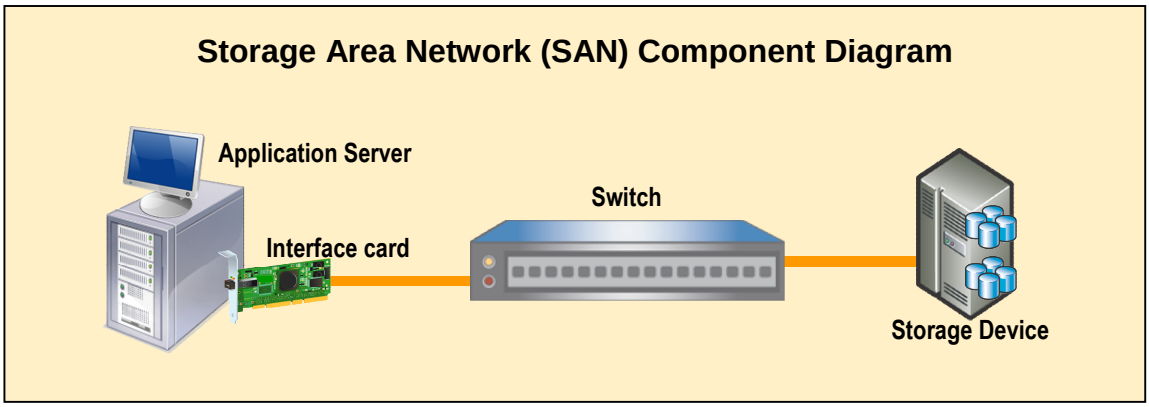

The components of an iSCSI SAN are equivalent in concept to those found in a Fibre Channel SAN, but typically are less expensive. However, there is a reason for the difference in price of the hardware components. Fibre Channel switches and host bus adapters almost always use managed switches and fiber optic technology including cables, connectors, transceivers, etc., which are more expensive than the typical copper-based technology found in Category 5, 5e, or 6 Ethernet cables, standard NIC cards and unmanaged Ethernet switches. The price differences

begin to diminish when comparing fully managed fiber-optic Ethernet and Fibre Channel environments. The line rates found in the Fibre Channel environments are 1, 2, and 4 Gbps. The line rates found in typical Ethernet environments are 1 Gbps or less.

When comparing costs of Fibre Channel and iSCSI storage solutions, it is important to include advanced storage and management functions in the calculations. Many of the iSCSI solutions have done a good job of simplifying common functions such as LUN provisioning and provide advanced storage and management functions in the solution for less cost than other types of solutions.

Unified Storage

Unified Storage is the concept of combining the technologies used in Fibre Channel SANs and/or iSCSI SANs with Network-Attached Storage (NAS) into a single, integrated storage solution. The solutions provide both block and file access to the shared storage environment. This type of storage solution can provide simplified management by combining the management of all storage regardless of the transport or “plumbing” into a single management console. Unified storage solutions often also provide advanced storage capabilities including replication functions, storage resource management, clustering and more.

10-Gigabit Ethernet Technology

10-Gigabit Ethernet (10 GbE) technology promises to deliver increased speed and a unified approach to networking, clustering and storage applications. Current deployments of 10-Gigabit Ethernet are found in server clustering and network-trunking applications, and many 10-Gigabit Ethernet switches available today also support 1-Gigabit connections into the same switch. Fiber-optic technology is used for these applications, as copper-based technology for 10-Gigabit Ethernet is still relatively new. The price ratio between 10-Gigabit Ethernet technology and 1-Gigabit Ethernet technology is dropping, but currently 10-Gigabit Ethernet technology remains priced too high for many organizations.

The PCI-Express (PCIe) bus has the bandwidth to handle the new high-speed interconnect technologies, including 10-Gigabit Ethernet technology. As PCI-Express becomes more common in servers and desktop computers, and prices continue to drop for high-speed offload adapters and other devices, 10-Gigabit Ethernet technology may begin to become cost-effective for more common use in iSCSI SANs over the next 12 – 24 months.

iSCSI Technology

Internet SCSI (iSCSI) is the combining of the SCSI command protocol that storage devices use with a TCP/IP network as the transport mechanism in order to provide “block” storage connectivity over an existing network. The iSCSI technology is implemented as a layer above the TCP layer in the TCP/IP protocol stack. Using the SCSI nomenclature of initiator and target, a host server would be the iSCSI initiator and a logical storage device or subsystem would be the iSCSI target. The iSCSI target can be implemented in hardware or software.

Once the connection is established between the iSCSI initiator and the iSCSI target, the operating system on the iSCSI initiator sees the storage as a local storage device that can be formatted, read and written in the usual manner. For example, Windows “Disk Management” sees disk volumes the same way whether they are connected via iSCSI or connected via more traditional means. Some restrictions on the iSCSI initiator are listed on page 19.

The iSCSI hardware and software components comprise an iSCSI storage area network (SAN). Hardware components required to implement iSCSI include a network adapter, a network switch, and an iSCSI target storage device or subsystem. The network adapter used for iSCSI traffic can be the same adapter used for traditional network traffic, but in many cases one or more separate adapters are used for iSCSI traffic. Software components required to implement iSCSI include iSCSI initiator software.

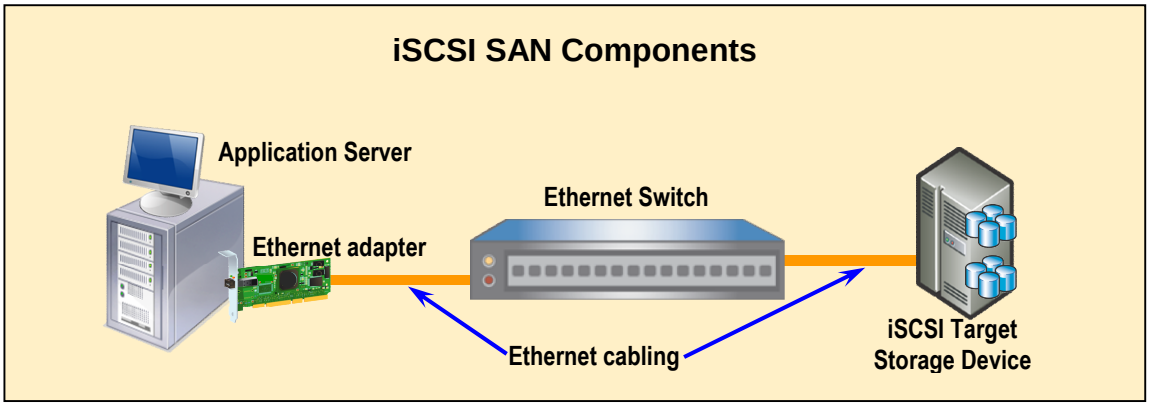

Because Ethernet infrastructure already exists in many environments, adding iSCSI components to an existing Ethernet infrastructure can be relatively inexpensive. Typically the management expertise and Ethernet network infrastructure (network switches and cabling) is already in place in most organizations.

Best practices for iSCSI SANs are to separate iSCSI storage traffic normal LAN traffic through the use of virtual LAN technology or by deploying iSCSI traffic to a physically separate network. This is because adding iSCSI traffic to existing LAN traffic may cause degraded overall network performance due to the different nature of iSCSI storage traffic.

Hardware requirements for iSCSI include a Gigabit Ethernet adapter in the host server connected to a Gigabit Ethernet switch. The iSCSI target storage device must also be connected to the Gigabit Ethernet switch.

As many environments have already moved to a Gigabit Ethernet infrastructure for their basic networking needs, iSCSI can be added for no additional hardware costs. For those who do not have a Gigabit Ethernet infrastructure, it can be created relatively inexpensively. Low-end, unmanaged, five-port Gigabit Ethernet switches are available for less than $100 today. Low-end, unmanaged, eight-port Gigabit Ethernet switches are available today for less than $200. Some unmanaged 16-port and 24-port Gigabit Ethernet switches are available today for less than $300. Category 5E or category 6 cabling is recommended for Gigabit Ethernet networks, which is present in many environments. Most of the newer servers have onboard Gigabit Ethernet network interfaces. Simple Gigabit Ethernet Network Interface Cards (NICs) are widely available for less than $50 and specialized, dual-port, server Gigabit Ethernet NICs with advanced networking features are available for less than $200.

Best practice for iSCSI is to use enterprise class, managed switches that support jumbo frames. Enterprise switches are generally designed to be used in higher-traffic networks and are better choices than low-cost switches for iSCSI traffic.

Some switches are designed with “oversubscription,” which is a design that takes advantage of the fact that average utilization of Ethernet links tends to be significantly less than full-bandwidth. As such, these switches cannot run at full-bandwidth on all ports at the same time. Care must be taken when deploying iSCSI traffic on this type of switch so that the switch is not oversubscribed.

Because iSCSI runs over standard Ethernet networks, there is virtually no distance limitation in the basic technology. Wide Area Networks (WANs) can be used to implement iSCSI technology, and iSCSI technology can be used to bring a remote or “stranded” server into an existing storage infrastructure. However, many applications and operating systems do not have a high tolerance for latency, so response time on storage devices should be considered. It is possible that a very large distance (thousands of miles) may generate a response time that is unacceptable. Some iSCSI storage providers recommend a network latency of less than 5 milliseconds, resulting in a distance of approximately 100 kilometers.

Initiators

Microsoft provides a free iSCSI software initiator, which can be downloaded from the Microsoft web site. There are other iSCSI initiators available, but this report will only discuss the Microsoft iSCSI software initiator. The focus of this report is primarily on iSCSI targets.

The iSCSI initiator works in combination with the network adapter. There are four basic combinations of iSCSI initiators and network adapters available today. Varying degrees of processing can be offloaded to the adapter hardware, depending on the type of network adapter deployed. These combinations of iSCSI initiator and network adapter are:

- Software iSCSI initiator with standard network card
- Software iSCSI initiator with advanced network card that supports Receive-side Scaling
- Software iSCSI initiator with network card that includes a TCP/IP Offload Engine (TOE)
- Hardware iSCSI host bus adapter (HBA) that provides offloaded TCP/IP and iSCSI processing.

The performance aspects of these types of iSCSI initiators are discussed below in the iSCSI performance section. The hardware iSCSI HBAs were not tested for this report but they do help to optimize CPU utilization in servers.

Targets

There are a variety of iSCSI target solutions available. One way to organize iSCSI target solutions is by the underlying technology. Some of these solutions are based on the Microsoft iSCSI target software and run on a Microsoft Windows Server platform. Other solutions run on a non-Microsoft platform. Some specific examples are discussed in further detail in subsequent sections of this document. Target solutions are available in a wide variety of storage capacities, performance levels, and prices.

Multi-Path I/O

Microsoft MPIO is supported with iSCSI Storage Area Networks as well as Fibre Channel and Serial Attached SCSI (SAS) storage. Microsoft includes a Microsoft iSCSI Device Specific Module (DSM) with the Microsoft iSCSI Software Initiator which supports many arrays and allows the creation of multiple paths for failover and load balancing. Storage array vendors can also license the Microsoft MPIO DDK from Microsoft and implement their own DSMs specific allowing their storage to interface to the Microsoft MPIO core driver stack. The Microsoft iSCSI initiator can be installed with Microsoft MPIO, the same MPIO that is available for other types of storage. Multi-path I/O provides the benefits of fail-over if a path fails and load balancing across multiple active paths to increase throughput.

It is important to note that when using multi-path I/O for iSCSI storage solutions, both the iSCSI initiator and the iSCSI target need to support MPIO. Each network adapter and its associated ports in the iSCSI initiator and iSCSI target should have identical features to insure consistent performance. The iSCSI DSM implements several load balance policies designed for different link performance metrics.

Management of iSCSI

The iSCSI solutions discussed in this report are managed using the Microsoft iSCSI initiator. The storage volumes can be managed using standard Windows tools such as “Disk Management”. In addition, most iSCSI target storage solutions provide Microsoft VSS and VDS hardware providers and can be managed with Microsoft Storage Manager for SANs (SMfS), which is available in Windows Server 2003 R2.

Storage Performance and iSCSI

General Performance Comments

One of the concerns about iSCSI is the overall performance of the solution, including the load on the host CPU, the iSCSI target performance, and the Ethernet network performance, especially during periods of heavy I/O. Although this report is not intended to be an exhaustive performance benchmark, some of these performance issues will be discussed.

iSCSI solutions are a blend of traditional network and traditional storage technologies, and most of the iSCSI storage solutions are pre-configured to provide good overall network and storage performance. Administrators may choose to fine-tune various advanced network and storage settings for additional performance or configuration purposes.

The implementations discussed for the various iSCSI target solutions were intentionally disparate to illustrate the variety of ways in which iSCSI targets can be deployed and the configurations discussed for individual products were not necessarily optimized for performance. As a result, the performance of the iSCSI target solutions varied widely, due to the variety of designs and components used. These storage solutions used a variety of storage devices, including SATA disk drives, SCSI (parallel) and SAS disk drives. The disk drives spun at various RPM including 7200, 10K and 15K RPM. Each storage array had a different number of disk drives in the array. Different RAID stripe sizes were used with different arrays. Various disk subsystem caching designs were used, not all of which have been publicly disclosed.

Basic I/O tests were performed with IOMeter, an open-source I/O load generator. The same group of block sizes and I/O patterns was tested with each iSCSI target solution; however the queue depth was varied as an additional data point. Some of the iSCSI target solutions supported multi-path I/O, and where possible, multiple paths were used. The deployment scenarios outlined below include up to 2 sessions. Although the purpose of these tests and this report is not to be a head-to-head performance comparison of the iSCSI target solutions, performance was measured in order to provide some general reference points for the expected performance range of iSCSI target solutions. Some interesting reference points comparing various types of network adapters in the host servers (iSCSI initiators) were also made. The IOMeter test results for each iSCSI target solution are included in their respective sections. Readers should take notice that IOMeter testing is by no mean a substitute for workload testing and modeling. In addition, tools from Microsoft such as LoadSim for Microsoft Exchange and SQLIO and SQLIOSim for SQL Server can be used to test how an iSCSI initiator and target respond for those particular applications.

IMPORTANT NOTE: the iSCSI targets presented in this white paper are different in class, price, and disk I/O characteristics, so head-to-head comparison of the iSCSI targets in the context of this report is not possible. In addition, the tests were run with different parameters to emphasize that this report is not a benchmark report.

Improving iSCSI Storage Performance

Performance improvements for iSCSI solutions can be determined by measuring either the increase in absolute network throughput or the reduction in system resources such as CPU utilization. Benefits may vary depending on the applications. Application performance improvements may depend on the network packet size and/or storage block size in use.

There are several areas that can be adjusted to improve iSCSI initiator performance on Microsoft Windows host platforms. It should be noted that several of these items listed below will improve general network performance as well as iSCSI initiator storage performance.

- Network Infrastructure Settings
- Microsoft Scalable Networking Pack
- Receive-side Scaling
- TCP Offload adapters
- Full iSCSI Host Bus Adapters (HBA)

Network Infrastructure Settings

Many network cards have various feature settings that can improve performance. Not all of these features are available on all network adapters. Jumbo Frames, TCP Checksum Offload and Large Send Offload can be enabled to improve performance. Windows Server 2003 is the first Windows platform that supports network adapters that include hardware TCP Checksum Offload and Large Send Offload features.

In the case of Microsoft Windows Server-based iSCSI target solutions, the network interface adapter settings should be examined on both the iSCSI initiator and the iSCSI target solution. It may be possible to have one side of the iSCSI communication highly optimized and the other side not optimized, resulting in reduced performance. The network features discussed in this section should be examined on the iSCSI initiator and, where possible, the iSCSI target. Implementations and impacts of these features on the iSCSI target may vary.

Network switches should have Jumbo Frames enabled. Flow control may also need to be enabled in the switch and network adapters if there is heavy network congestion. Enterprise switches are generally designed to be used in higher-traffic networks and are better choices than low-cost switches for iSCSI traffic.

Microsoft Scalable Networking Pack

With the gaining popularity of multi-core and multi-processor systems, deployment of the Microsoft Scalable Networking Pack with advanced, server-class network adapters is highly recommended.

Microsoft makes the Scalable Networking Pack (SNP) available as a free download for Microsoft Windows 2003 Server (32-bit and 64-bit) and for Windows XP 64-bit platforms. It is also an integrated component within Windows Server 2003 R2 Service Pack 2. This package provides new and improved network acceleration and compatibility with hardware-based offload technologies. Three technologies included in the Scalable Networking Pack help optimize server performance when processing network traffic. Because iSCSI uses the network, it can take advantage of these technologies. These technologies are Receive-side Scaling, TCP Offload, and NetDMA. NetDMA was not tested for this report.

Receive-side Scaling

Receive-side scaling is especially important in multi-core and multi-processor systems because of the architecture of the NDIS 5.1 miniport drivers. Without the SNP and Receive-side Scaling, multi-processor and multi-core Windows 2003 Server systems route all incoming network traffic interrupts to exactly one processor core, resulting in limited scalability, regardless of the number of processors or processor cores in the system. With SNP and Receive-side Scaling and the NDIS 5.2 miniport driver, incoming network traffic interrupts are distributed among the processors and processor cores on the computer. Receive-side Scaling-capable network

adapters are now available, and are required to take advantage of this feature. Support for this feature is currently found in some, but not all server-class network adapters.

The Scalable Networking Pack monitors network adapters for Receive-side Scaling capabilities. If a network adapter supports Receive-side Scaling, the Scalable Networking Pack uses this capability across all TCP connections, including connections that are offloaded through TCP Offload.

TCP Offload Adapters

TCP Chimney is the Microsoft Windows Server term for offloading the TCP protocol stack into network interface adapters. Network adapters that support this feature are also known as TCP/IP Offload Engines (TOE). TCP Chimney is an operating system interface to advanced Ethernet network adapters that can completely manage TCP data transfer, including acknowledgement processing and TCP segmentation and reassembly.

Full iSCSI Host Bus Adapters (HBA)

Another approach to use for offloading CPU processing cycles is to combine the iSCSI initiator and the full TCP processing onto one adapter card and perform all these functions in hardware. This work performed for this report used the Microsoft iSCSI software initiator for all examples, so iSCSI HBAs were not used, but many models are supported on Windows Servers.

Performance Result Summary by Initiator Network Adapter Type

Although this report is not a full performance benchmark, several performance measurements were taken using various network adapters with the same I/O workloads.

Ethernet network adapters are one important component of an iSCSI storage solution. **It should be noted that best practices recommend that a true server-class network adapter should be used for iSCSI storage applications.** The low-cost network adapter listed below that was used in these tests is not a true server-class network adapter, but was used only as a point of reference. This truly shows the importance of server-class network adapters in iSCSI deployments.

Three different types of Gigabit Ethernet network adapters were used for these tests. Two low-cost network adapters were deployed, each with one port. The advanced network adapter and the TCP Offload adapter are dual-port, server-class network adapters. The low-cost network adapters used in this report are available for the least cost, **but are not recommended for iSCSI storage solutions**. The advanced network adapters are available for a mid-range price. The TCP Offload adapters are, by comparison, more expensive. The three types of network adapters used in the Demartek lab were:

- Low-cost network adapter: *NetGear® GA-311*
- Advanced network adapter supporting Receive-side Scaling: *Intel® Pro/1000 PT*
- TCP Offload network adapter: *Alacritech® SEN2002XT*

The CPU usage on the dual-core, single processor server using the low-cost network adapter, without the Scalable Networking Pack was significantly skewed toward the first core with very little activity on the second core, especially during read operations. When SNP was installed and the advanced network adapter and the TCP Offload adapter were each used, the dual-core processor server exhibited a lower and more evenly balanced CPU utilization. The following Task Manager snapshots highlight the differences for light to moderate workloads using a mid-range iSCSI target solution.

Dual-core server with
low-cost network adapter

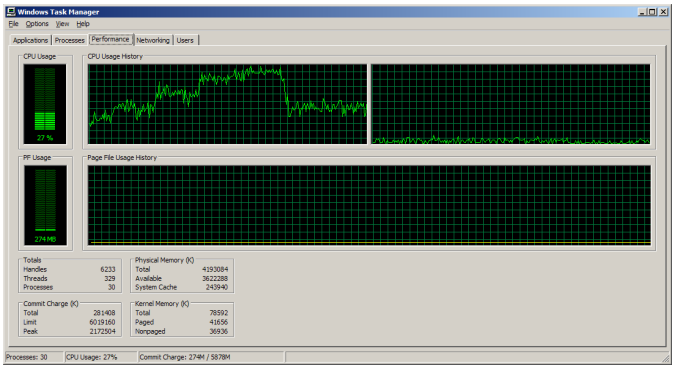

Dual-core server with
SNP and Receive-side Scaling
server-class network adapter

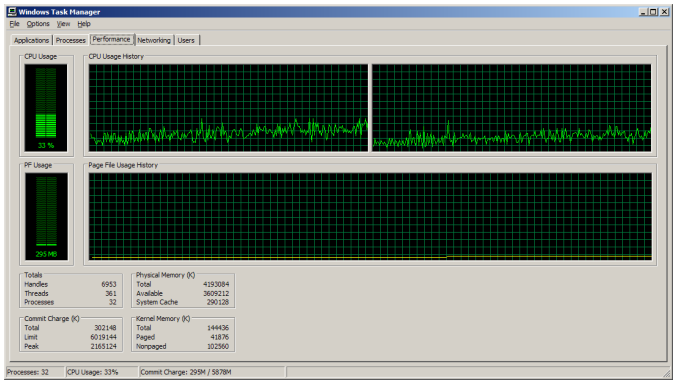

Dual-core server with
SNP and TCP Offload
server-class network adapter

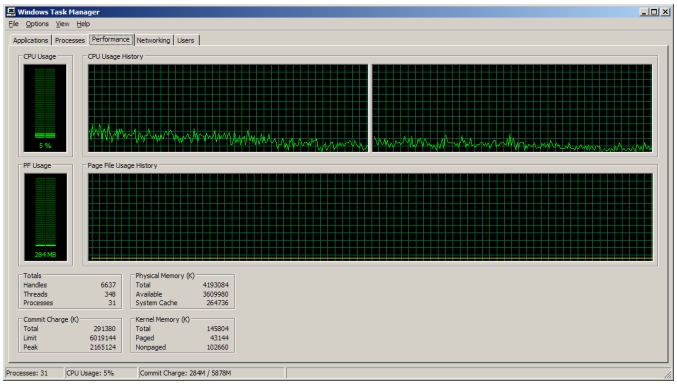

The differences between server-class network adapters and low-cost network adapters became obvious during our tests. We found that under heavy workloads with a high-performance iSCSI target solution, the low-cost network adapter configuration became unacceptably slow, fully utilizing the processor and locking out other processes on the server including mouse and keyboard controls. The same workloads using the advanced network adapter and the TCP Offload adapter, which are server-class network adapters, completed the workloads in the expected time and did not lock out other processes.

The charts below show a representative sample of percentage of CPU utilization for the three types of network adapters. Two paths, using MPIO, were used for this sample.

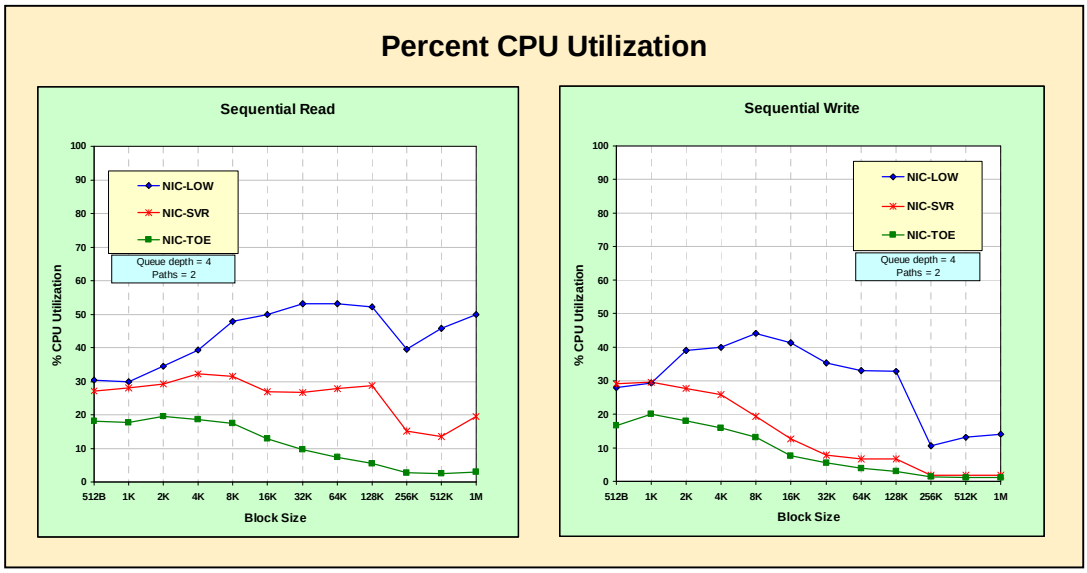

Network adapter legend:

- NIC-LOW: low cost network adapter
- NIC-SVR: advanced server-class network adapter supporting Receive-side Scaling
- NIC-TOE: TCP Offload adapter supporting TCP/IP Offload Engine

Performance will vary depending on many factors, including number of processors and processor cores in the application server, amount of memory in the application server, network adapter type, specific network adapter features that are enabled, and the iSCSI target storage system characteristics.

Deployment Examples

Deployment Environment

The deployment took place at a lab at Microsoft headquarters and at a lab at Demartek headquarters. The equipment at each location was similar.

Servers

Two host servers were used at each location. Windows 2003 Server R2 Enterprise x64 Edition was installed on all the application servers.

	Microsoft		Demartek	
	Server A	**Server B**	**Server A**	**Server B**
Processor	Single Xeon 2.8 GHz	Dual Xeon 3.0 GHz	Pentium D 3.4 GHz	Pentium D 3.4 GHz
Cores	1	1 each	2	2
Memory	1 GB	1 GB	4 GB	4 GB

Network Adapters

The network adapters used in the Microsoft lab servers were standard server-class network adapters. At the Demartek lab, three different types of network adapters were used, including the standard network adapters, advanced server-class network adapters and TCP Offload server-class network adapters.

Network Switches

Both labs used Gigabit switches. The Microsoft lab used a NetGear GS108 unmanaged, 8-port switch. The Demartek lab used a NetGear GS724T smart-switch with 24 ports. Jumbo frames were enabled on the GS724T switch.

Deployment Processes

The deployment processes for this report were broken into several main steps.

- Application Host configuration steps
- Storage Solution configuration steps
- Run applications to use the iSCSI storage
- Additional host configuration steps

In addition, special application scenarios were performed, including:

- Configure Microsoft Exchange 2003 to use iSCSI storage
- Configure Microsoft SQL Server 2005 to use iSCSI storage
- Configure Microsoft Cluster Server to use iSCSI storage

Application Host Configuration Steps

The initial steps used to configure the application host servers were:

- Install fresh copy of Windows Server 2003 R2 Enterprise x64 Edition

- Install Microsoft iSCSI initiator
- Configure Microsoft iSCSI initiator

Storage Solution Configuration Steps

The iSCSI storage solution configuration was separated into the steps necessary on the iSCSI target (storage solution itself) and the steps needed on the iSCSI initiator (application server). The general steps for configuring the iSCSI storage are outlined below, including some optional steps. Each specific iSCSI storage solution configuration follows this general outline, but the exact steps required varied slightly for each iSCSI storage solution. A separate section of this report is devoted to each iSCSI storage solution configuration, including screen shots taken during the deployment.

Target Configuration Steps

1. Configure network settings for iSCSI target device
2. Launch management console
3. Create LUNs on disk array
4. Make LUNs ready for use (formatting, etc.)
5. Create iSCSI Targets
6. Optional – Configure multi-path I/O for iSCSI Targets
7. Optional – Configure security for iSCSI Targets (CHAP, etc.)
8. Make iSCSI Targets ready for use for iSCSI Initiators (virtual disks, etc.)

Initiator Configuration Steps

1. Optional – Configure multi-path I/O from application host (iSCSI initiator)

Run Applications to Use the iSCSI Storage

Applications were used to read and write to the iSCSI storage, including basic Windows management functions such as Disk Manager.

In addition, IOMeter was used to read and write various block sizes to the iSCSI storage. IOMeter is an open-source I/O load generator and performance analysis tool. IOMeter is available from Source Forge at: http://sourceforge.net/projects/iometer/.

Although this report is not a performance benchmark for iSCSI storage solutions, some performance data are included to provide a general idea of the type of performance that can be expected with some iSCSI storage solutions. Two different lab locations and two different sets of servers were used for the deployments described in this document.

Application Host Configuration Steps

The following steps were used to initially configure each of the application hosts. These hosts became the iSCSI initiators.

Install Microsoft Windows 2003 Server R2 Enterprise x64 Edition

Each host was configured by installing Microsoft Windows 2003 Server R2 Enterprise x64 Edition. This particular version was installed so that all possible memory and clustering options could be explored. All the required security updates were installed after installing the base operating system.

The Microsoft iSCSI initiator can be installed on the following Microsoft Windows platform families

- Microsoft Windows Server 2000 SP4 or later
- Microsoft Windows XP Professional (32-bit or 64-bit)
- Microsoft Windows 2003 Server (32-bit or 64-bit) – Service Pack 1 or higher is recommended

Beginning with Microsoft Windows Vista and the next release of Windows Server, the Microsoft iSCSI initiator is pre-installed.

Install Microsoft iSCSI Initiator

The Microsoft iSCSI initiator is a free download from the Microsoft web site. There are versions for x86 (32-bit), x64 (AMD64 and Intel EM64T) and IA64 (Intel Itanium) processors. Version 2.03 of the iSCSI initiator was used on all hosts for this report.

The iSCSI initiator installation process is the same relatively simple process for all the application hosts, and is wizard-based.

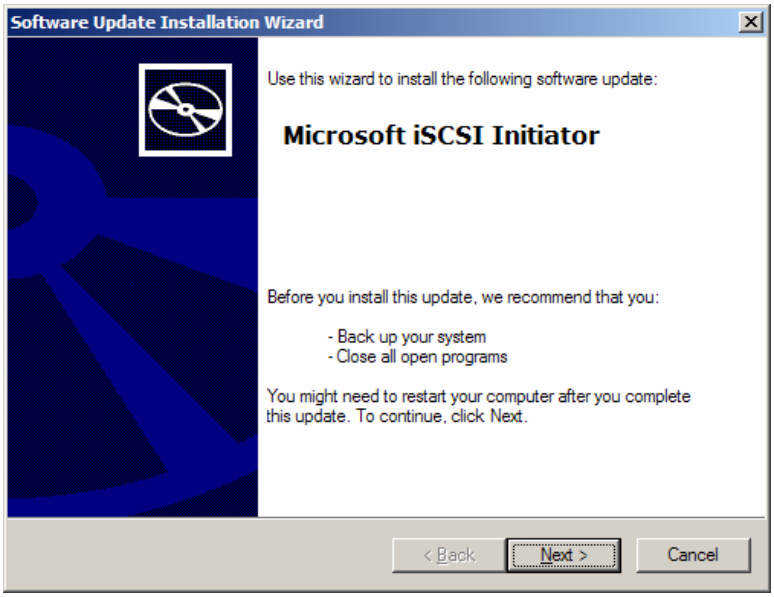

By default, the “Initiator Service” and “Software Initiator” features are checked. By default, the Microsoft MPIO multi-pathing feature is not checked. All the installations for this report used the MPIO feature and so this item was checked during installation.

If MPIO is not selected at installation, but desired later, the installer must be run again and the MPIO option selected. Beginning with Microsoft Windows Vista and the next release of Windows Server, MPIO is a feature that can be selected without re-installing.

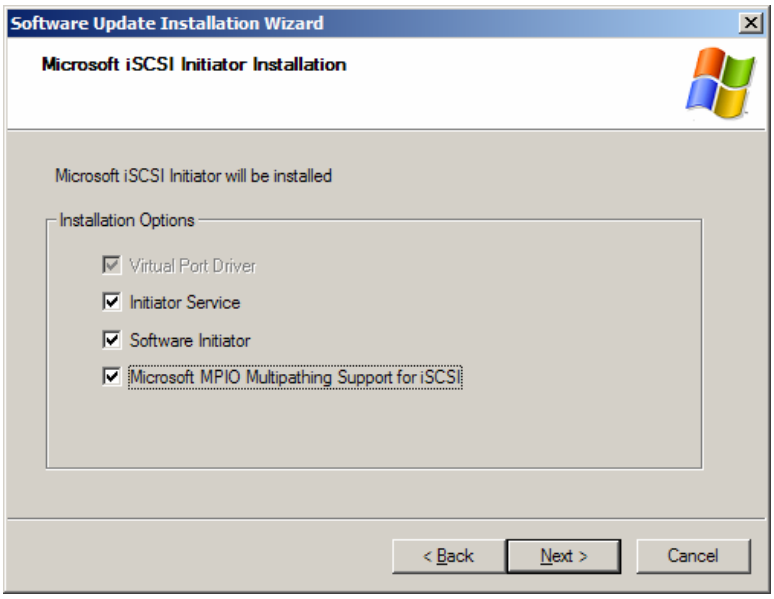

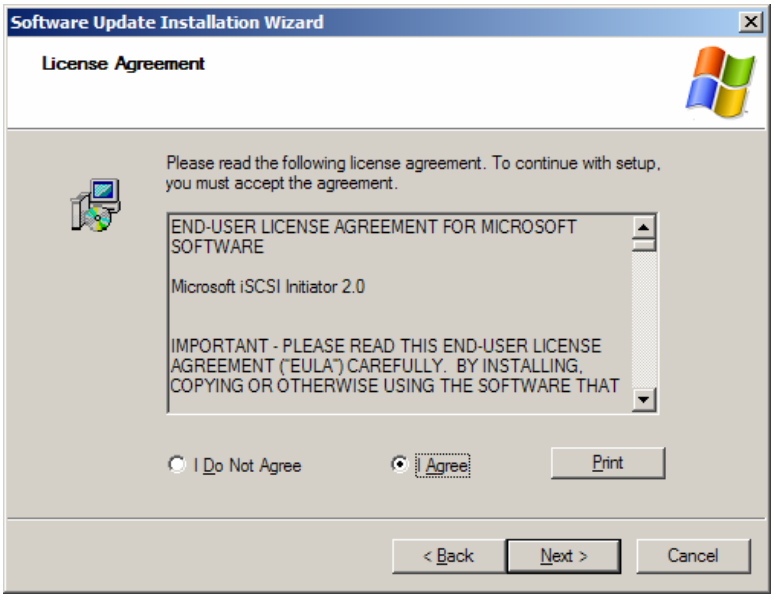

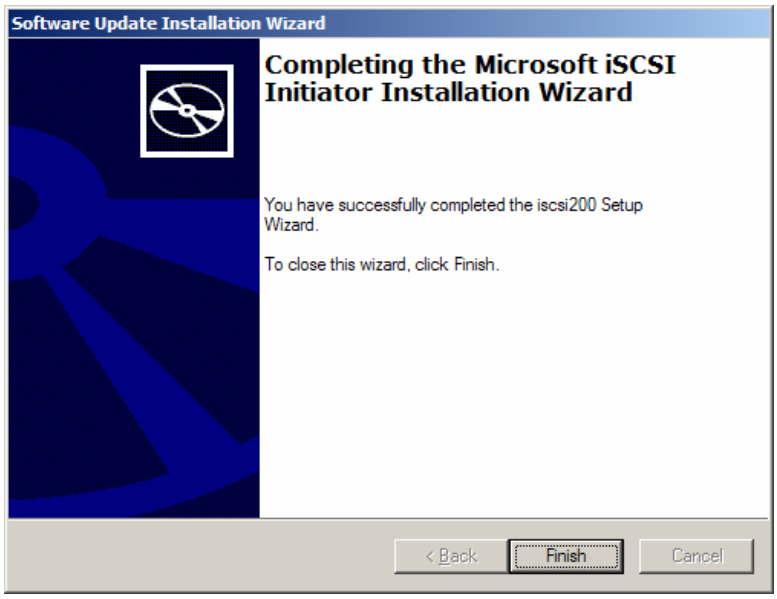

A command-line utility, “iSCSIcli” is also installed that can be used to configure connections to iSCSI targets from the Windows Server host.

The release notes and user guide are installed onto the local host when the iSCSI initiator is installed. A few items from the release notes are worth mentioning here. Some of the restrictions listed here may change in future releases.

- Dynamic disks on an iSCSI session are not supported. With Windows Vista and Windows Server Codename “Longhorn”, Dynamic Volumes are supported; however for best performance and redundancy, it is recommended that customers use basic disks or volumes in Windows combined with hardware RAID available in the storage arrays.
- Note that the default iSCSI node name is generated from the Windows computer name. If the Windows computer name contains a character that would be invalid within an iSCSI node name, such as '_', then the Microsoft iSCSI Initiator service will convert the invalid character to '-'.
- Both initiator and target CHAP secrets should be greater than or equal to 12 bytes, and less than or equal to 16 bytes if IPSec is not being used. It should be greater than 1 byte and less than or equal to 16 bytes if IPSec is being used.
- The checked and retail versions of the Microsoft iSCSI Software Initiator will only install on retail version builds of Windows. There is no package that installs on checked builds of Windows.

Configure Microsoft iSCSI Initiator

After installation, the Microsoft iSCSI initiator is used to manage the iSCSI environment.

General Tab

The general tab shows the initiator node name, which is the iSCSI Qualified Name (IQN).

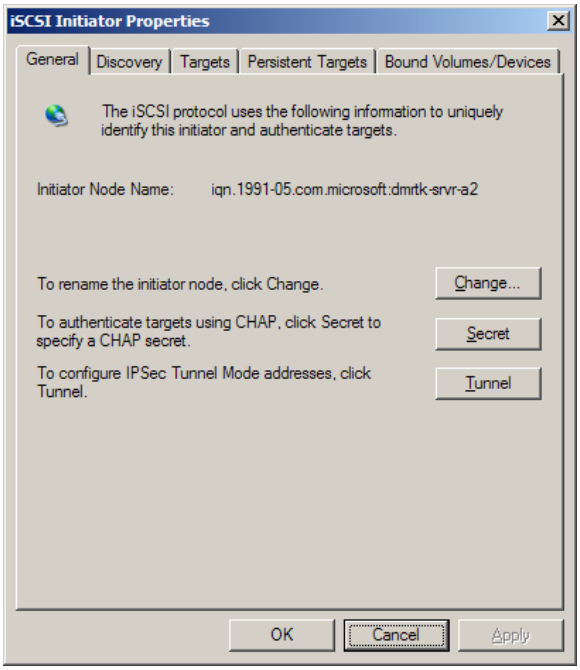

Discovery Tab

The discovery tab provides the list of discovered iSCSI target portals available to this initiator. The target portal is the primary IP address of the iSCSI target solution. Some target solutions use a virtual IP address and some iSCSI target solutions provide the first actual IP address of the solution. If there are no target portals listed, they can be added by IP address or DNS name. In this case two iSCSI target portals have already been discovered. If there is an iSNS server available, it can provide all the iSCSI discovery information.

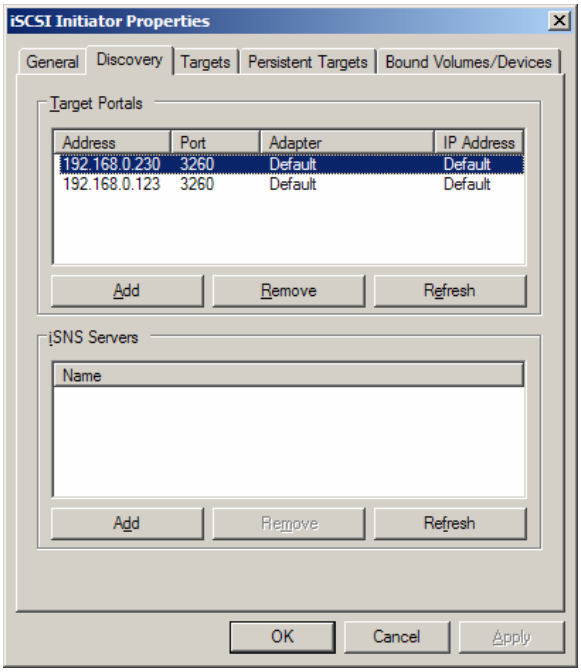

Targets Tab

The targets tab provides the list of individual targets available to the iSCSI initiator. In this case, three targets are available to the iSCSI initiator.

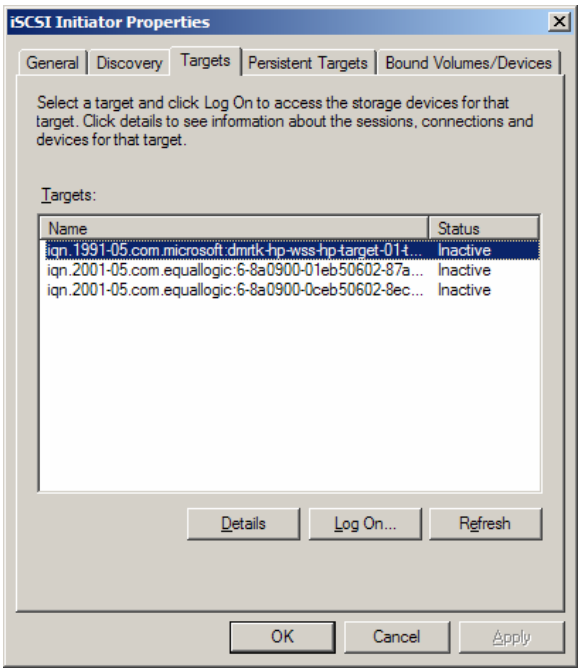

To gain access to the target, the initiator must “Log On” to the target. If there is only one path to the target, there is only one step needed for log on.

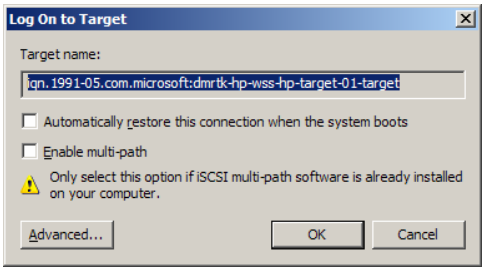

If there are multiple-paths to the target, then each path must be described to the iSCSI initiator. This is done by enabling multi-path and clicking on the “Advanced” tab. The Log on process must be repeated for each separate path.

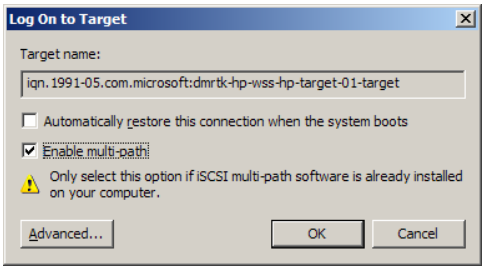

The advanced tab provides a drop-down box for all the possible source (initiator) IP addresses and a separate drop-down box for all possible target portal addresses. Some target solutions provide one virtual IP address for multi-path operations. In these cases, the target solution

manages the actual paths and IP addresses internally. Other target solutions expose each available IP address that can be used for multi-path operations.

The administrator must select each desired combination of source IP address and target IP address.

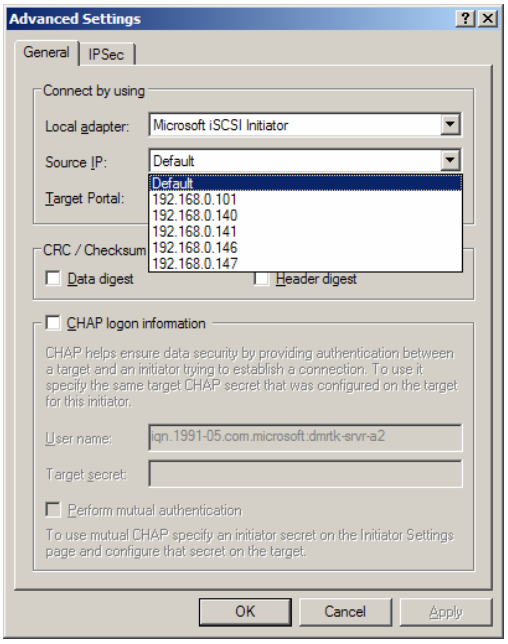

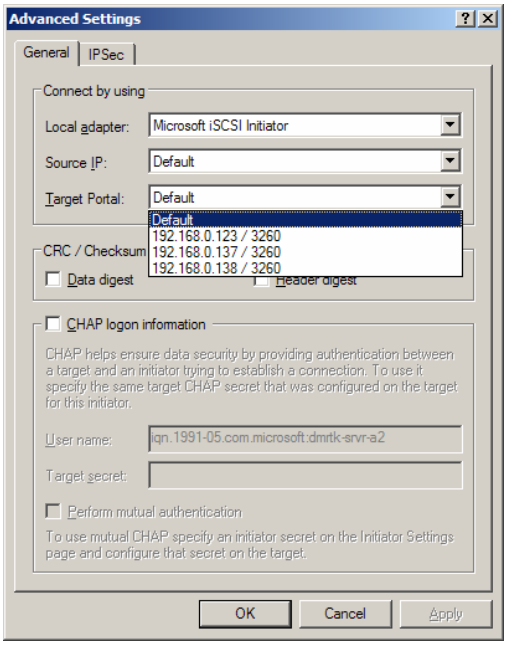

After the IP address or addresses have been selected, the initiator is connected to the target and the log on process is complete.

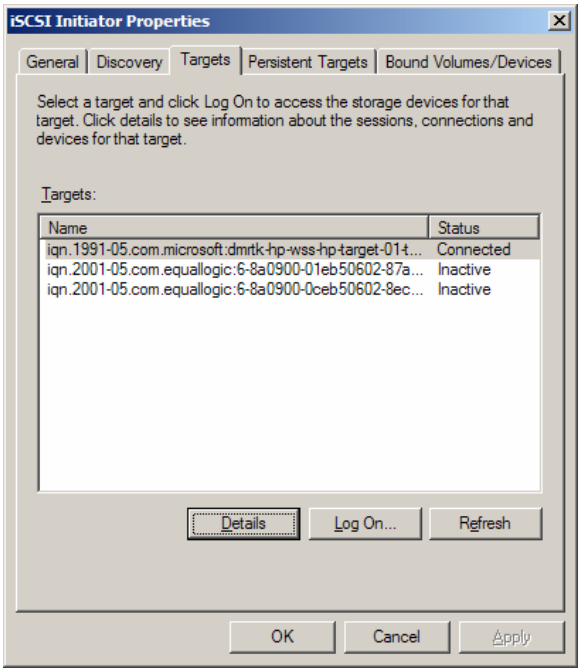

Path Load-balancing and fail-over are configured here. During the target logon process, multiple paths can be configured by selecting different combinations of source and target IP addresses. After all the paths have been selected, the desired load-balancing or fail-over behavior can be configured. This behavior must be assigned for the session and the individual targets.

In this example, three paths have been specified for this target as shown by selecting “Details” from the targets window. For even load-balancing, select “Round Robin” for the session connections load balance policy.

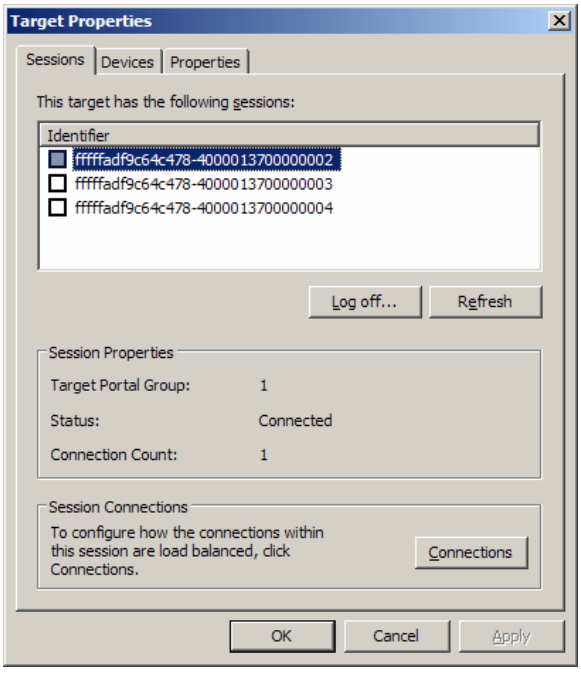

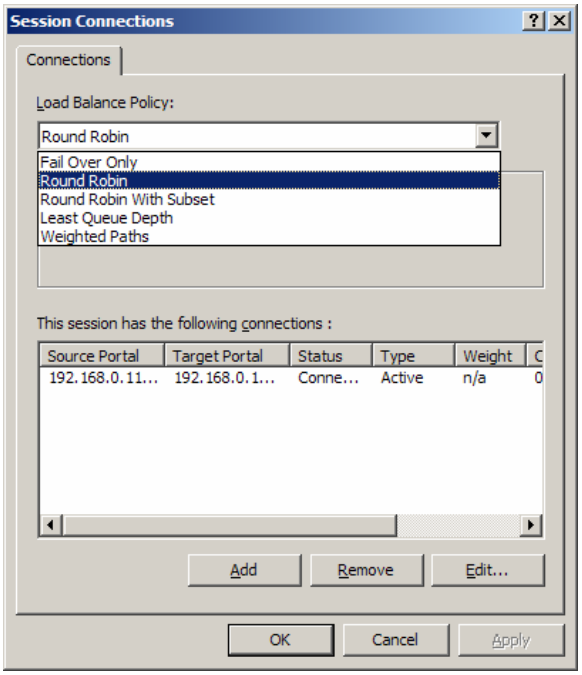

For the target multi-pathing configuration, select the “Devices” tab from the targets window, then select “Advanced” to get to the MPIO selection tab. Select the load balance policy for each device.

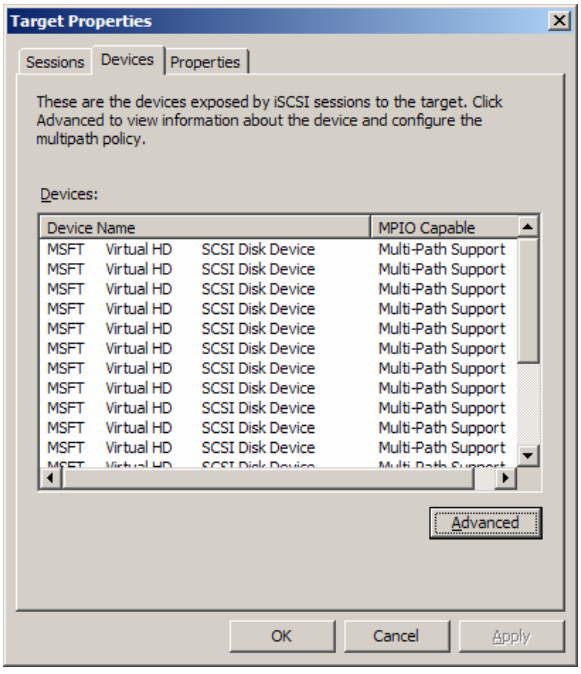

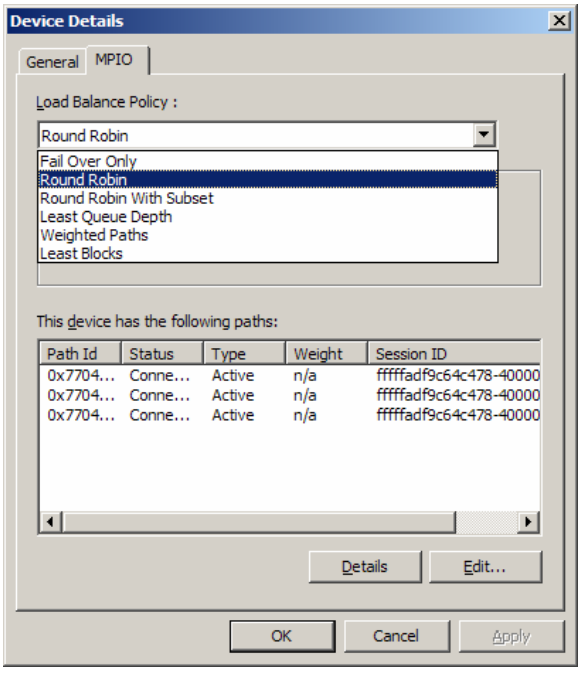

Persistent Targets Tab

Targets can be configured to be persistent, which means that the connection to the target is automatically restored when the system reboots. If the targets are configured to be persistent, they appear in this dialog box.

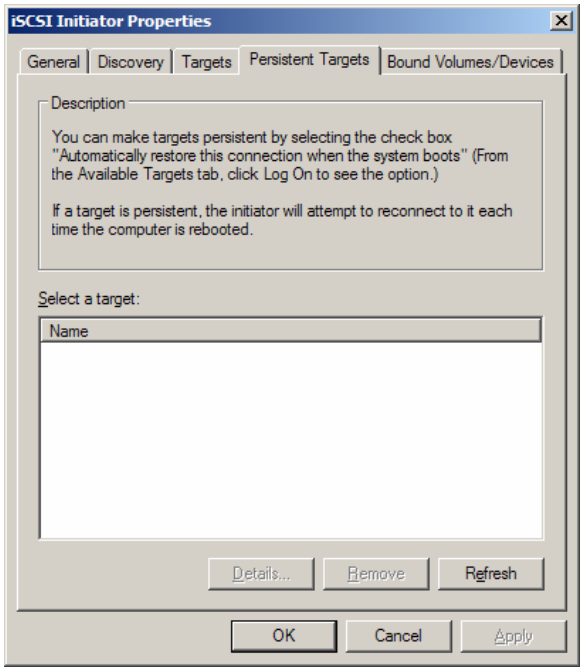

Bound Volumes Tab

If a host service or application depends on the availability of an iSCSI volume, it should be "bound" so that the iSCSI service includes each "bound" volume as part of its initialization.

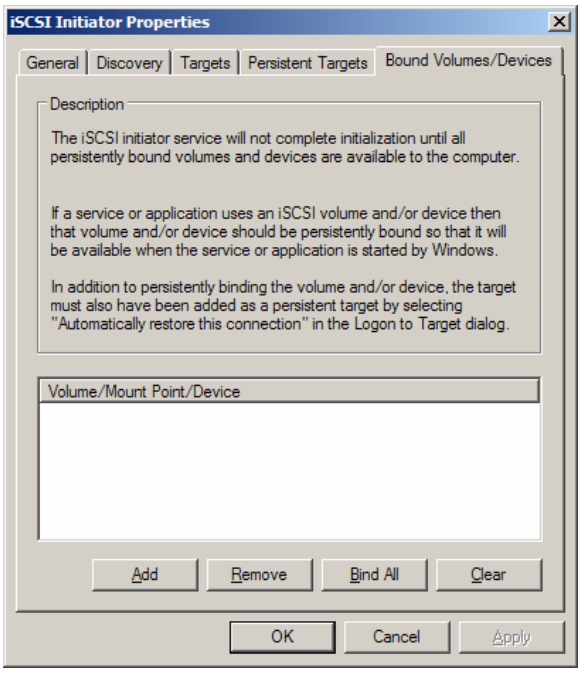

Security for iSCSI

Security for iSCSI includes some security features in the iSCSI layer itself, separate from any security layers that may be present in the lower TCP, IP, and Ethernet layers. The iSCSI security features can be enabled or disabled, as desired.

Each environment will need to address the issue of running storage traffic over the same network as the “public” LAN. Many will address this by running iSCSI storage traffic over a separate network or VLAN, which is the ***recommended best practice*** from Microsoft for applications using iSCSI storage. The items listed below are features of iSCSI which can provide increased security even if the iSCSI traffic is on a separate network.

The Microsoft iSCSI initiator uses Challenge Handshake Authentication Protocol (CHAP) to verify the identity of iSCSI host systems that are attempting to access storage targets. Using CHAP, the iSCSI initiator and iSCSI target share a predefined secret. The initiator combines the secret with other information into a value and calculates a one-way hash using MD5. The hash value is transmitted to the target. The target computes a one-way hash of its shared secret and other information. If the hash values match, the initiator is authenticated. The other information includes an ID value that is increased with each CHAP dialog to protect against replay attacks. Mutual CHAP is supported.

CHAP is generally regarded as more secure than PAP. More information is available on CHAP and PAP in RFC1334.

IPSec is also available for iSCSI. If IPSec is enabled, all IP packets sent during data transfers are encrypted and authenticated. A common key is set on all IP portals, allowing all peers to authenticate each other and negotiate packet encryption.

The Microsoft iSCSI initiator can be configured with the CHAP secret by clicking the “Secret” button from the “General” tab of the iSCSI initiator.

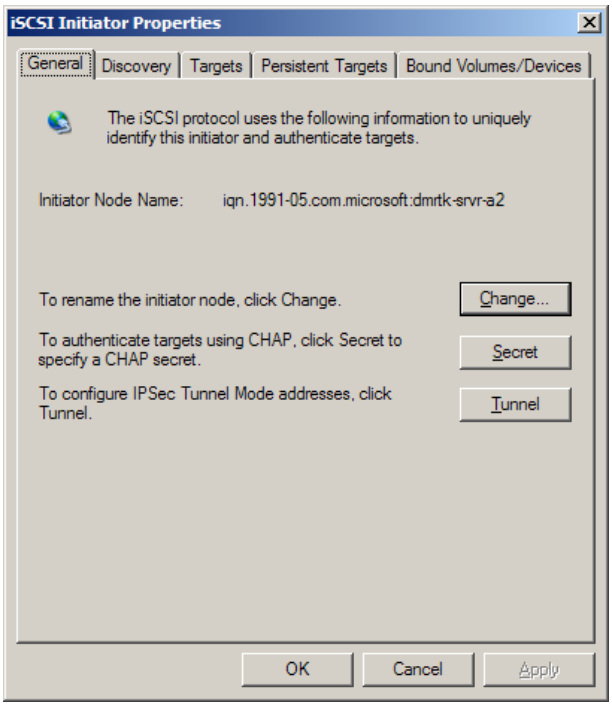

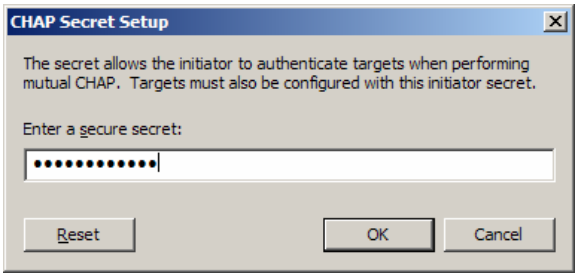

Microsoft Application Deployments for iSCSI

Microsoft Cluster Server

Microsoft Cluster Server (MSCS) supports the use of iSCSI connected disks as shared storage. Before creating the cluster, the storage that will be shared among the cluster nodes, including the quorum disk must be available. The volumes that will be shared by the cluster must be created on the iSCSI target and made available to the first node (iSCSI initiator) of the cluster. After the cluster is created, the iSCSI target should be configured to allow the other nodes of the cluster to access the same volumes as the first node.

Pre-Cluster Network Preparation

Microsoft clusters require at least one network interface configured with a static IP address on each node of the cluster for cluster communication. The cluster nodes also need at least one separate network interface to communicate to the clients on the LAN. Consult the Related Links section of this document for additional information on Microsoft Clusters.

Target Pre-Cluster Preparation Tasks

The volumes are created and associated with the iSCSI target that is mapped to the first node of the cluster. In the following example, the HP StorageWorks 1200 All-in-One was used for the iSCSI target storage for the cluster, which uses virtual disks for the volumes.

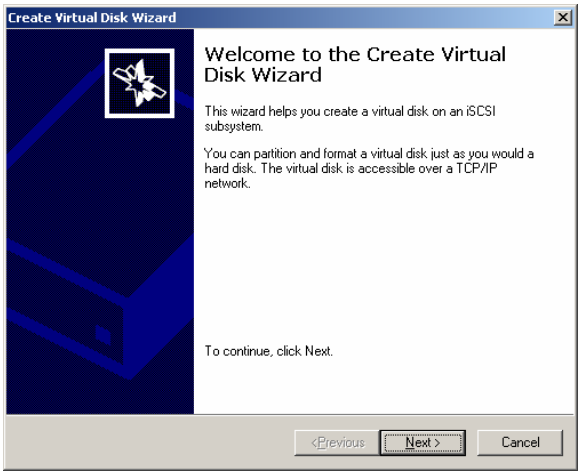

Create the virtual disk for the volume that will be used as the quorum disk.

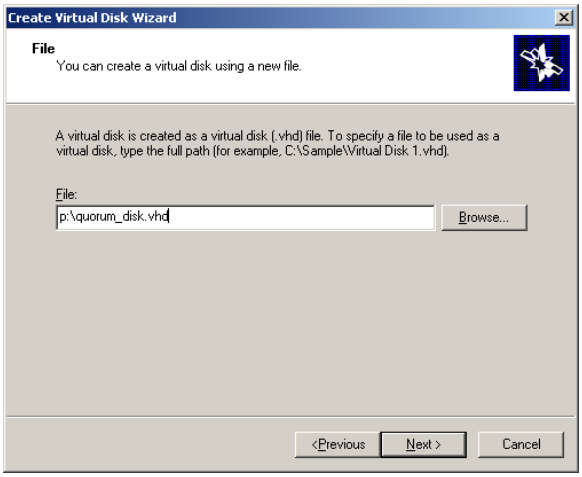
Create Virtual Disk Wizard
File
You can create a virtual disk using a new file.
A virtual disk is created as a virtual disk (.vhd) file. To specify a file to be used as a virtual disk, type the full path (for example, C:\Sample\Virtual Disk 1.vhd).
File:
p:\quorum_disk.vhd
Browse...
<Previous
Next >
Cancel

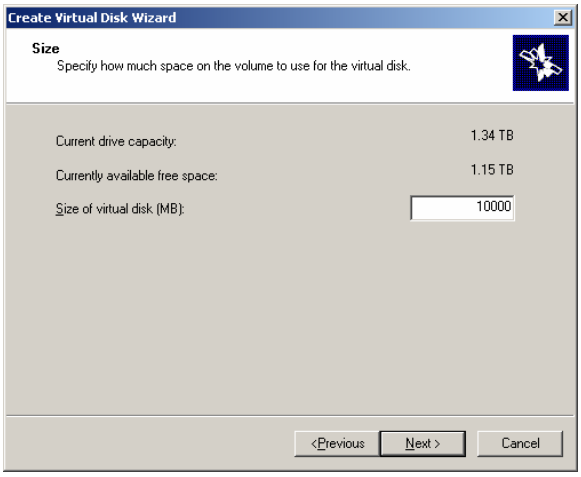
Create Virtual Disk Wizard
Size
Specify how much space on the volume to use for the virtual disk.
Current drive capacity:
1.34 TB
Currently available free space:
1.15 TB
Size of virtual disk (MB):
10000
<Previous
Next >
Cancel

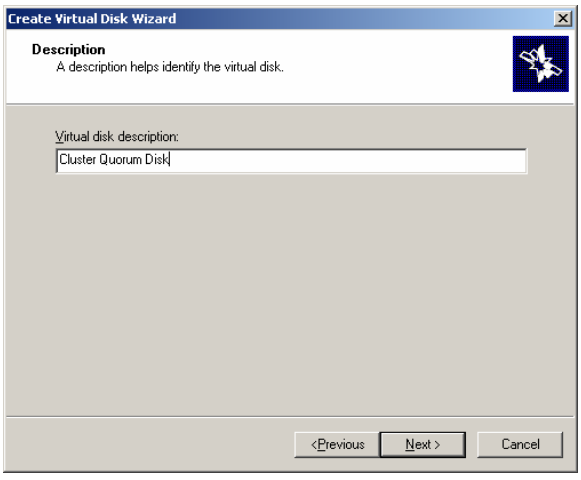
Create Virtual Disk Wizard
Description
A description helps identify the virtual disk.
Virtual disk description:
Cluster Quorum Disk
<Previous
Next >
Cancel

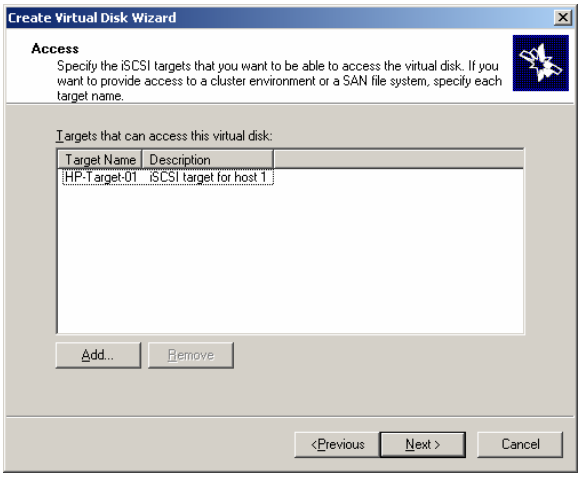
Create Virtual Disk Wizard
Access
Specify the iSCSI targets that you want to be able to access the virtual disk. If you want to provide access to a cluster environment or a SAN file system, specify each target name.
Targets that can access this virtual disk:
Target Name
Description
HP-Target-01
iSCSI target for host 1
Add...
Remove
<Previous
Next >
Cancel

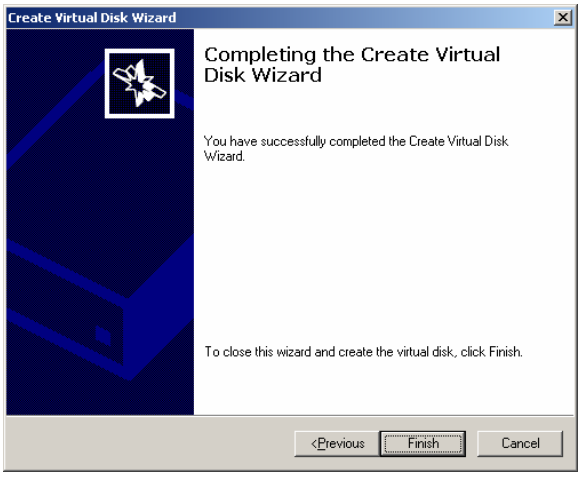
Create Virtual Disk Wizard
Completing the Create Virtual Disk Wizard
You have successfully completed the Create Virtual Disk Wizard.
To close this wizard and create the virtual disk, click Finish.
<Previous
Finish
Cancel

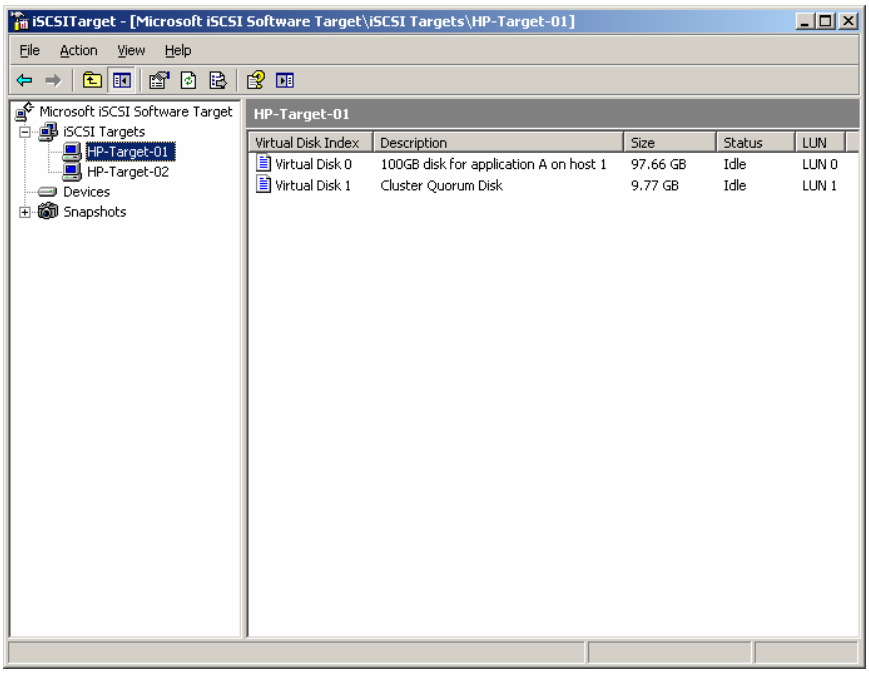

Other virtual disk volumes are created for applications that will use the cluster.

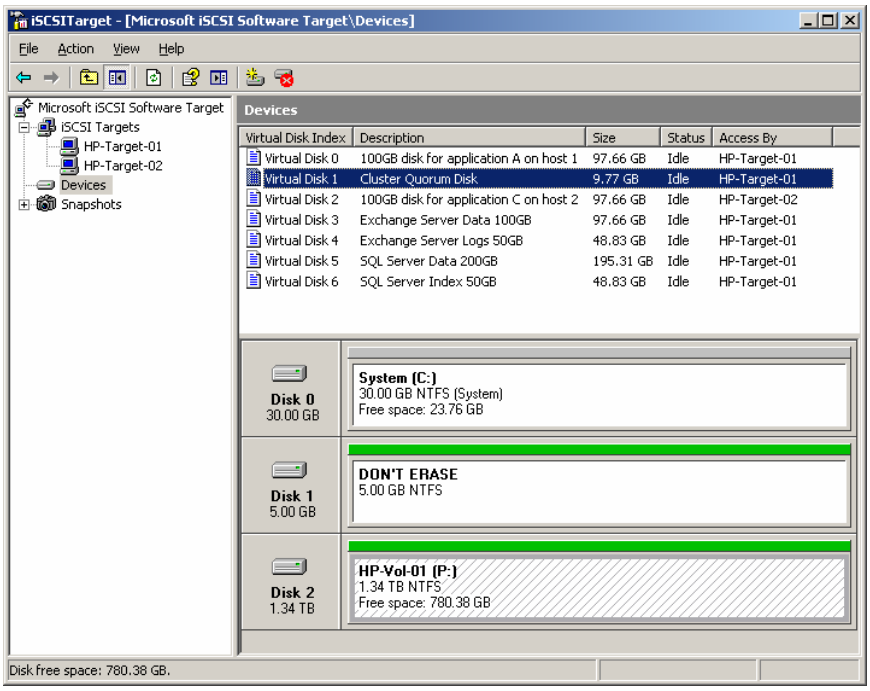

After the volumes are created, the second node is added to the iSCSI initiator list for the target.

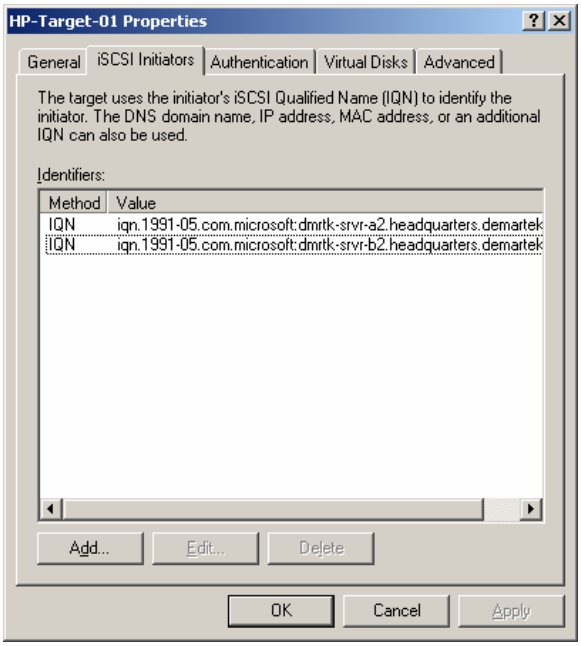

Because cluster nodes must be members of a domain, the full domain identifiers are added to the list of iSCSI initiators.

Cluster Creation Tasks

Once the storage is available for the cluster, the storage must be activated on the first node, and then the cluster itself can be created.

Disks 3 through 8 are volumes on the iSCSI target storage system. These disks are activated via Disk Manager on the first node, and then formatted in the standard fashion.

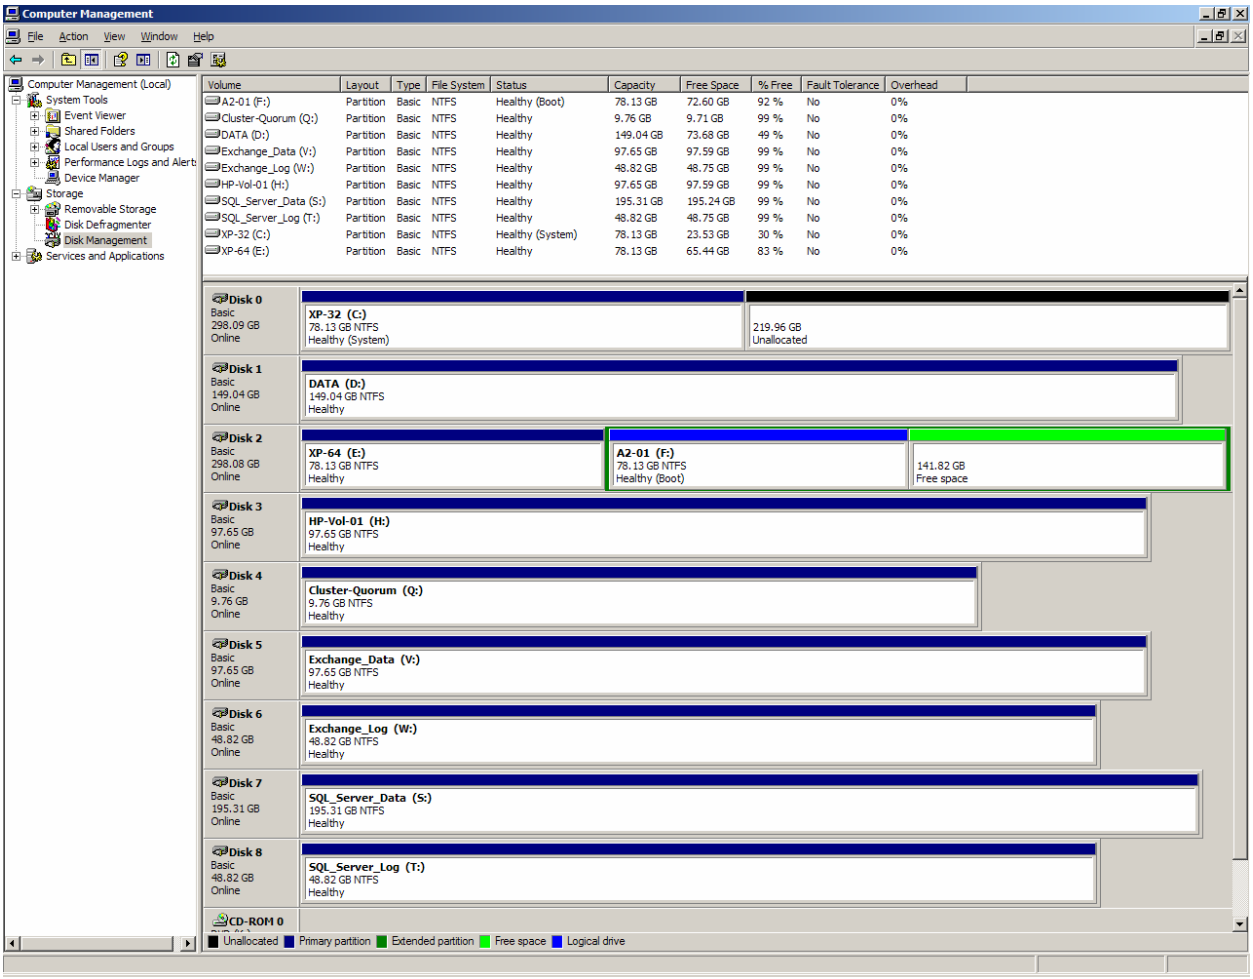

The Cluster Administrator wizard is used to create the cluster and assign the resources, and is launched on the first node.

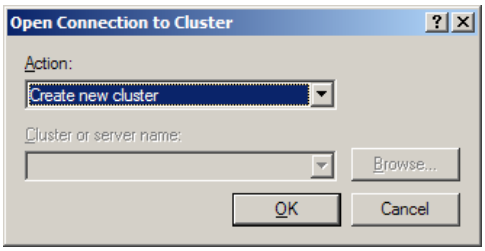

The cluster name is given to the cluster.

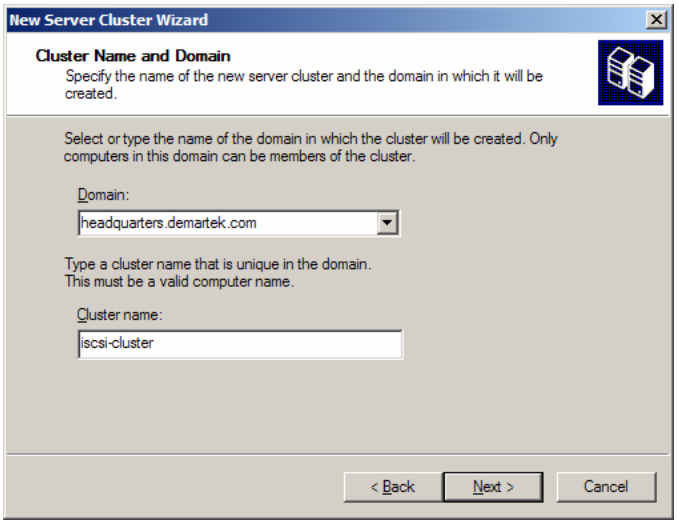

The first node of the cluster is identified. The first node name is “dmrtk-srvr-a2”.

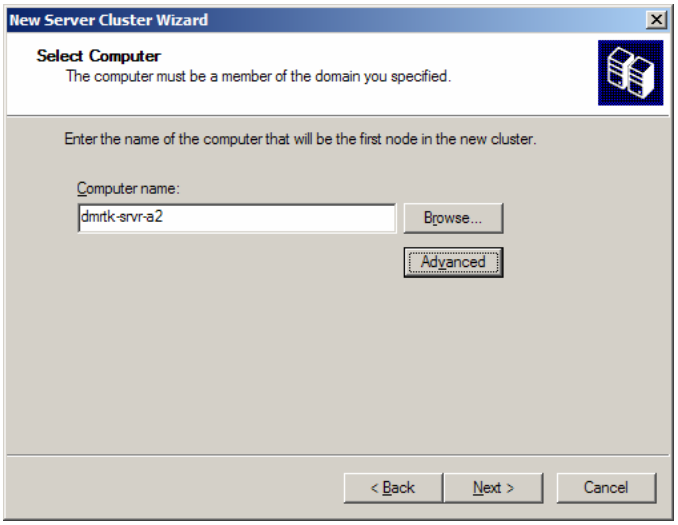

The cluster wizard performs its initial analysis of the node, the network, and the storage.

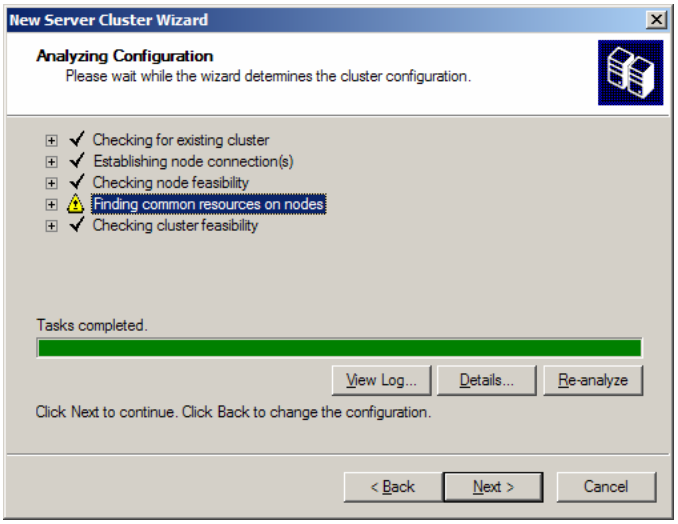

The cluster IP address is given.

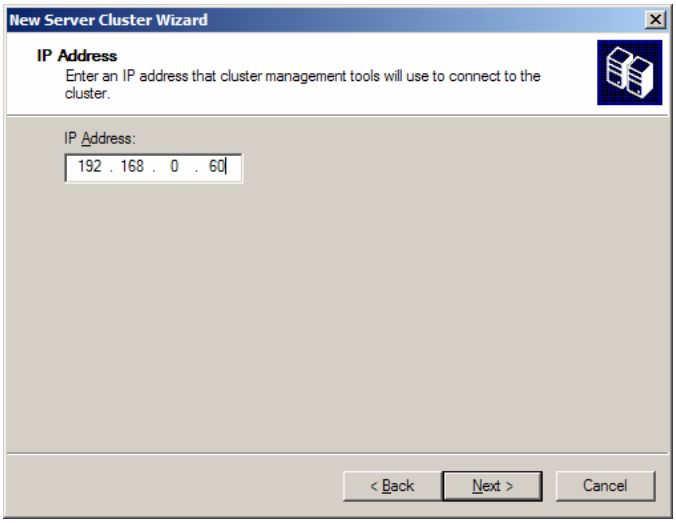

The cluster service domain account is provided.

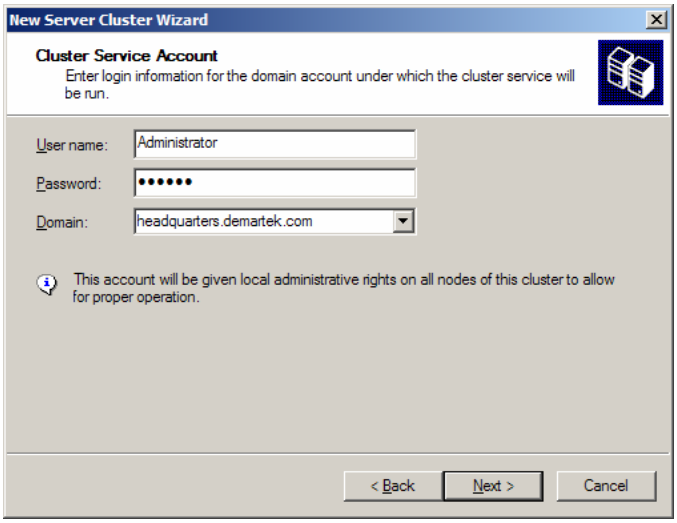

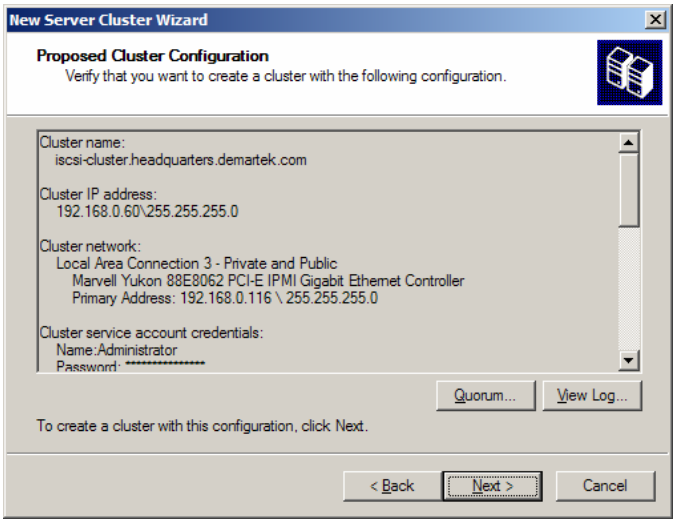

The cluster wizard performs its final analysis.

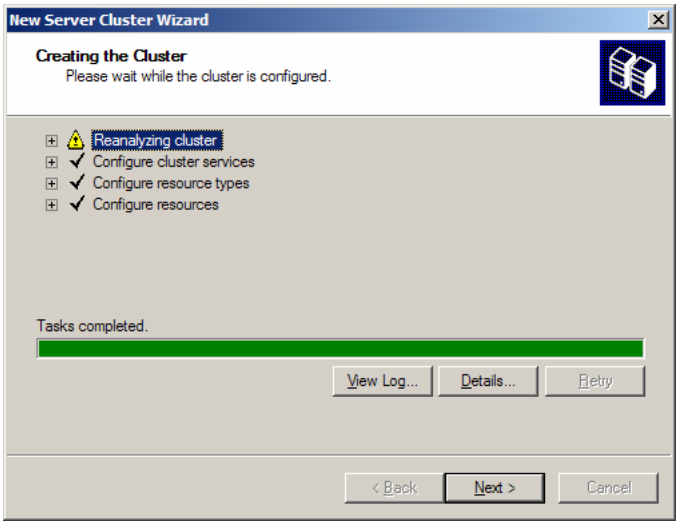

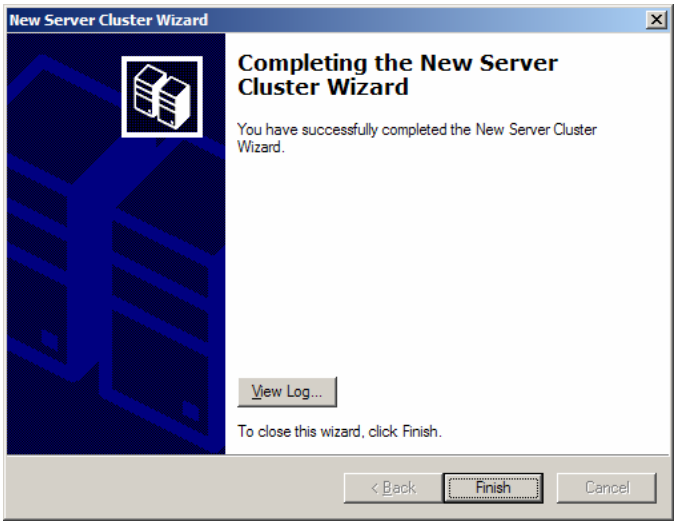

The Cluster Administrator has the first node of the cluster.

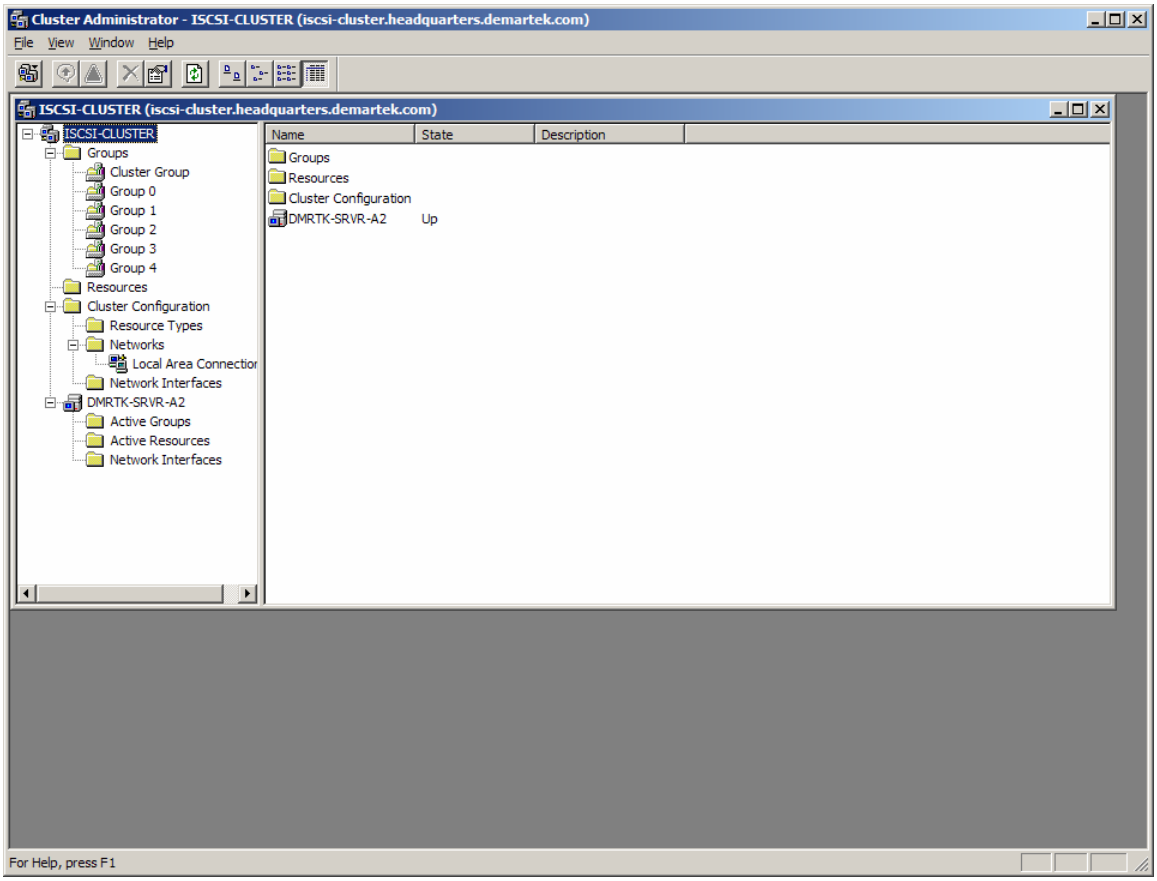

The next step is to add another node to the cluster.

The second node is “dmrtk-srvr-b2”.

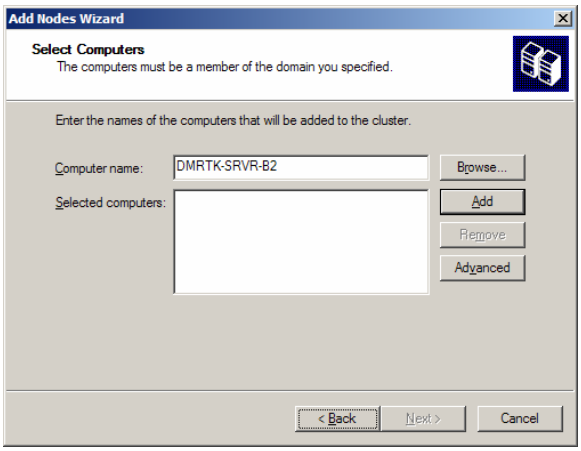

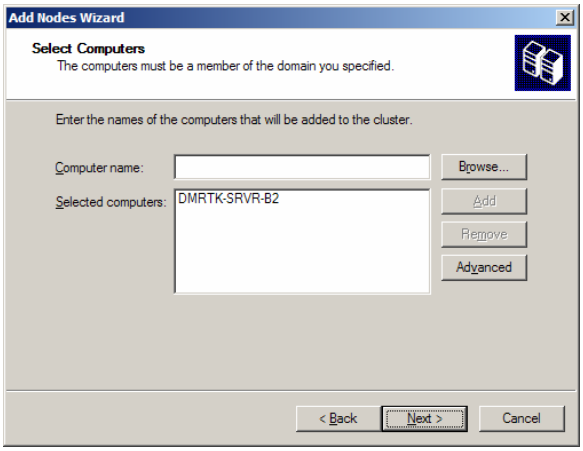

The wizard performs its initial analysis of the second node.

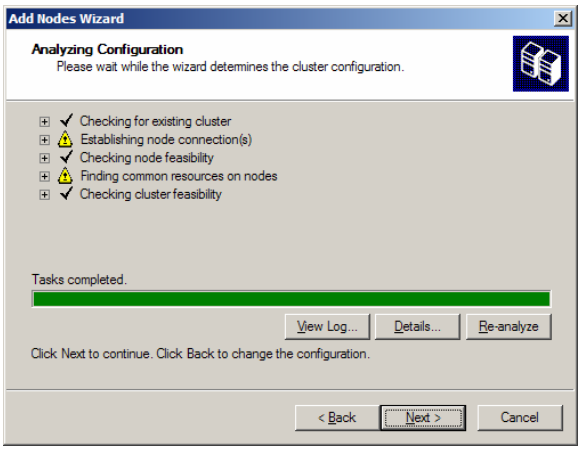
Add Nodes Wizard
Analyzing Configuration
Please wait while the wizard determines the cluster configuration.
Checking for existing cluster
Establishing node connection(s)
Checking node feasibility
Finding common resources on nodes
Checking cluster feasibility
Tasks completed.
View Log...
Details...
Re-analyze
Click Next to continue. Click Back to change the configuration.
< Back
Next >
Cancel

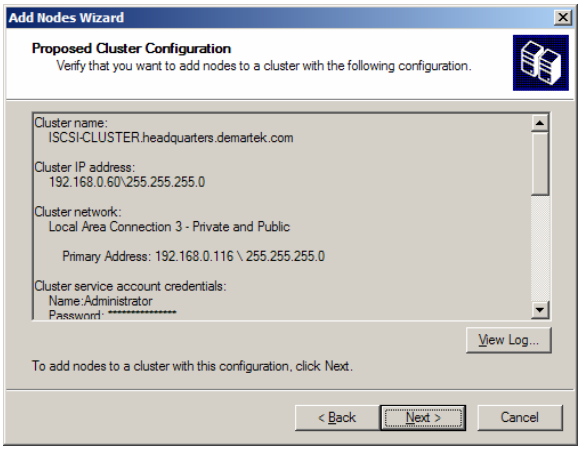
Add Nodes Wizard
Proposed Cluster Configuration
Verify that you want to add nodes to a cluster with the following configuration.
Cluster name:
ISCSI-CLUSTER.headquarters.demartek.com
Cluster IP address:
192.168.0.60\255.255.255.0
Cluster network:
Local Area Connection 3 - Private and Public
Primary Address: 192.168.0.116 \ 255.255.255.0
Cluster service account credentials:
Name:Administrator
View Log...
To add nodes to a cluster with this configuration, click Next.
< Back
Next >
Cancel

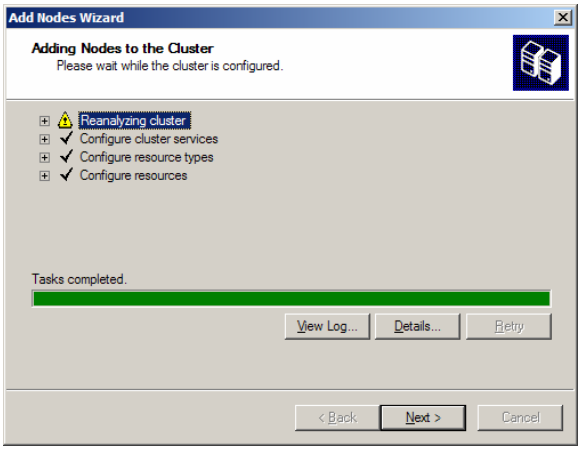
Add Nodes Wizard
Adding Nodes to the Cluster
Please wait while the cluster is configured.
Reanalyzing cluster
Configure cluster services
Configure resource types
Configure resources
Tasks completed.
View Log...
Details...
Retry
< Back
Next >
Cancel

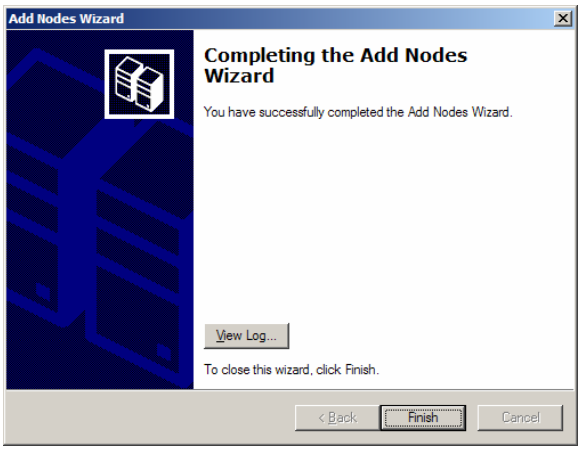

The Cluster Administrator now has the second node.

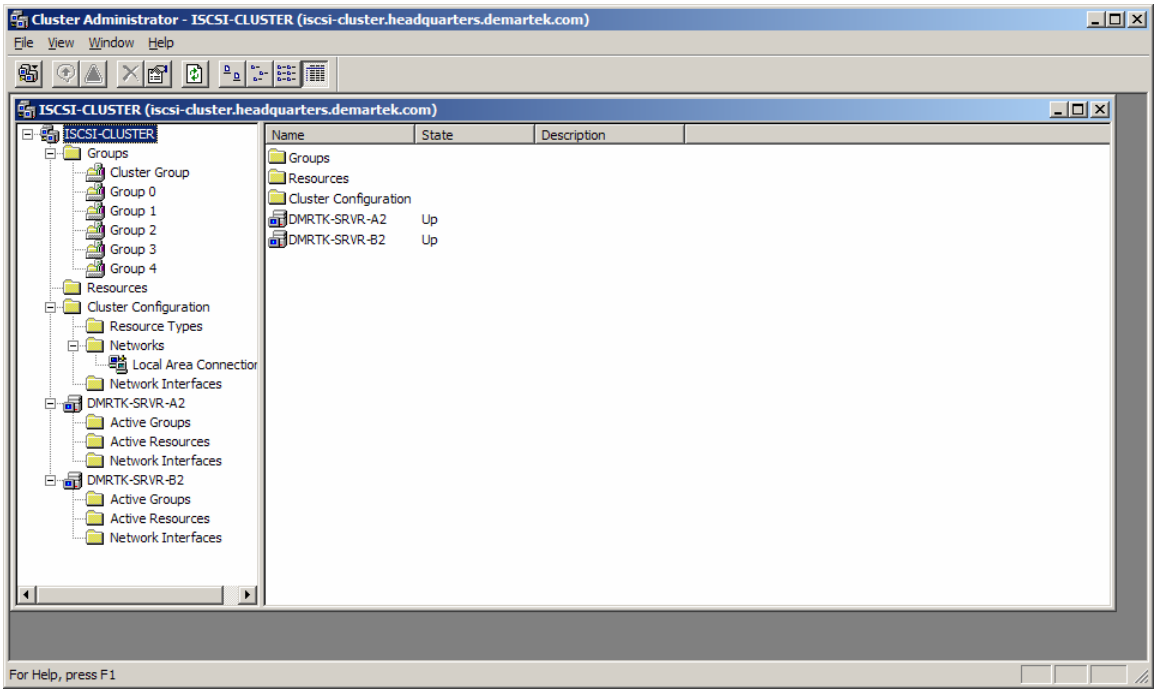

The cluster groups are displayed.

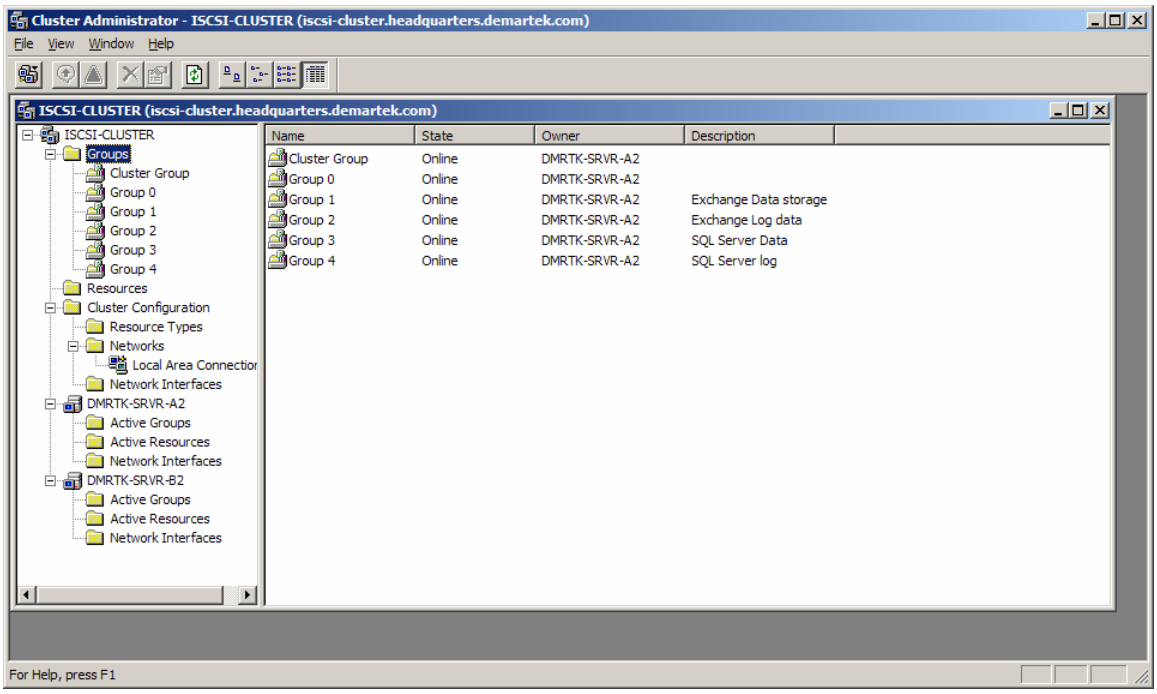

The groups can be assigned to the other node, as needed. Groups 3 and 4 are moved to the other node, as in a fail-over scenario. This storage is now visible to the second node, and no longer visible to the first node. The Disk Manager view below is from the first node after moving the resources to the second node. Notice that disks 7 and 8 are not accessible to the first node.

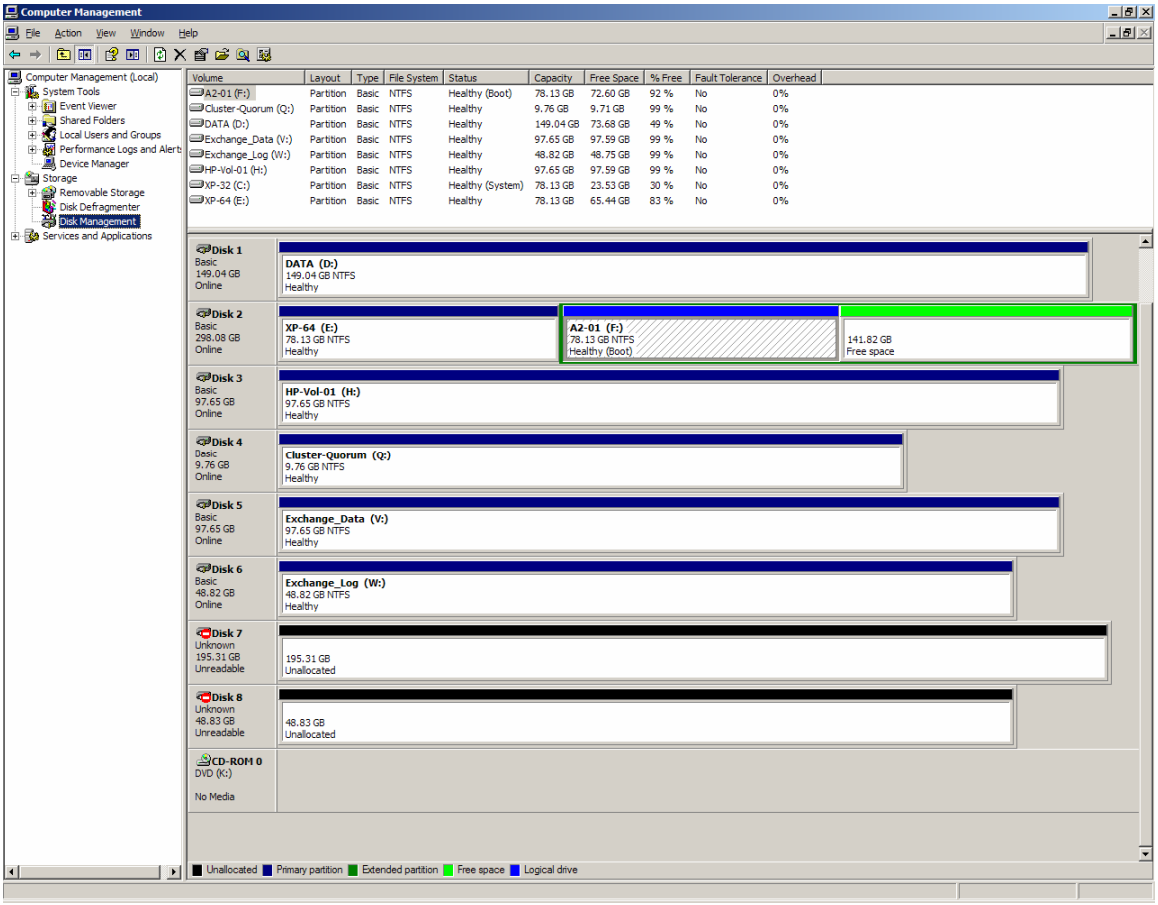

The cluster is now ready for applications and data.

Microsoft Office SharePoint Server 2007

Microsoft Office SharePoint Server 2007 supports the use of iSCSI connected disks as storage for the SharePoint Server and its associated data files.

For this installation, the disk drive letters “N”, “S” and “T” are all iSCSI target volumes allocated to server DMRTK-SRVR-B2, which is running Microsoft Windows Server 2003 R2 Enterprise x64 Edition. The iSCSI target volumes were created in the standard fashion described in other sections of this document.

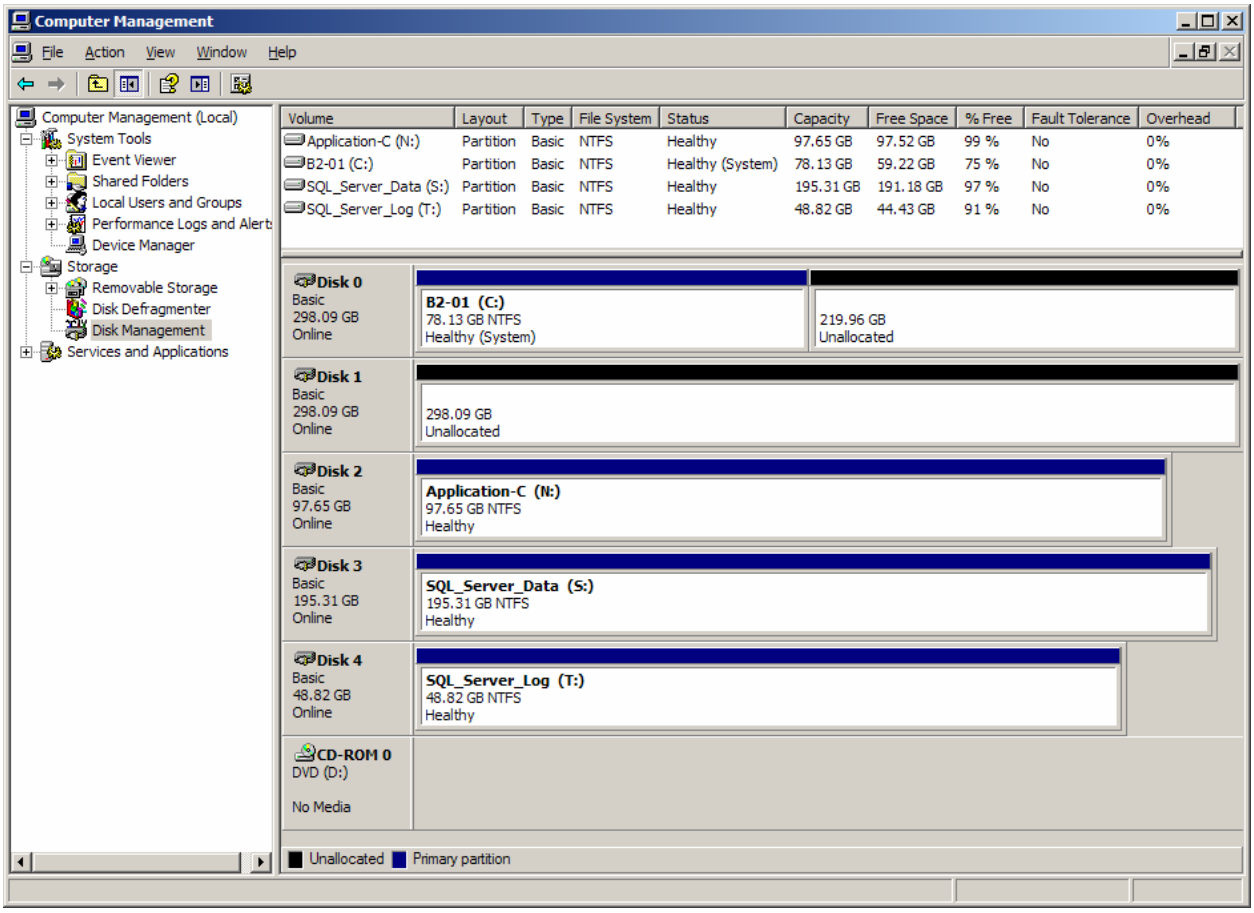

SharePoint Deployment

The installation wizard prompts steps through the installation with minimal interaction from the user.

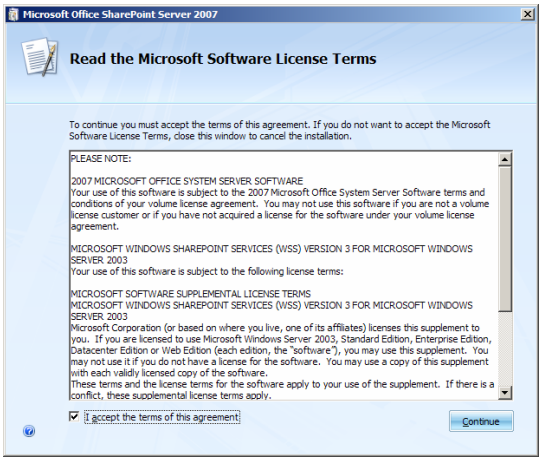

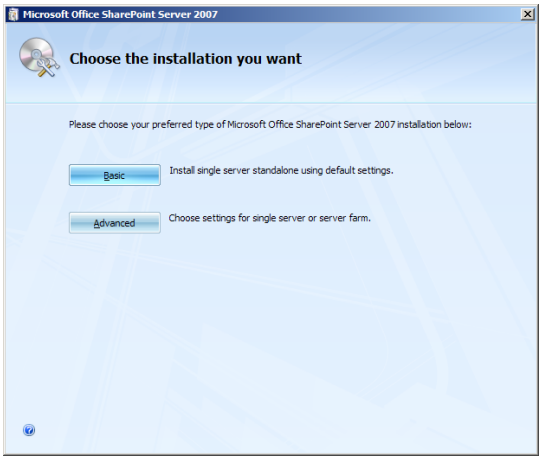

The Advanced installation option is chosen so that the iSCSI target device, drive letter “N”, can be used for the location of all the SharePoint programs and data. A stand-alone deployment will be performed.

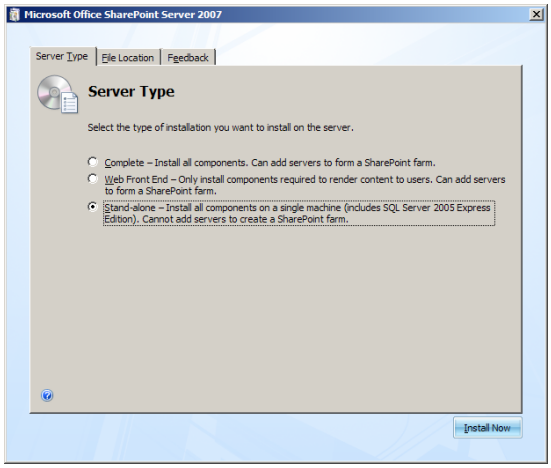

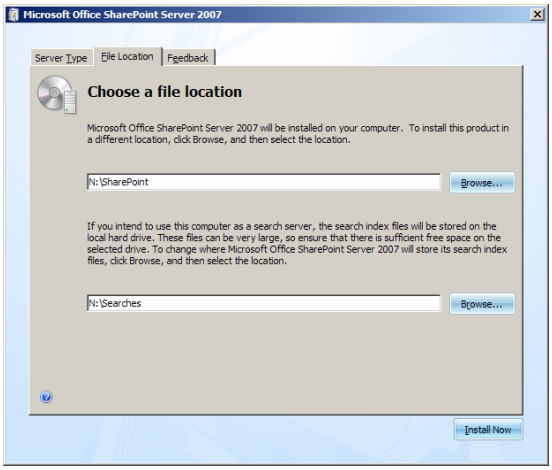

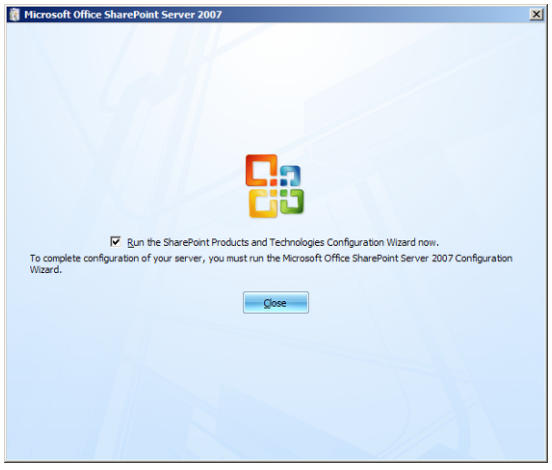

When the installation wizard completes, the SharePoint configuration wizard is launched.

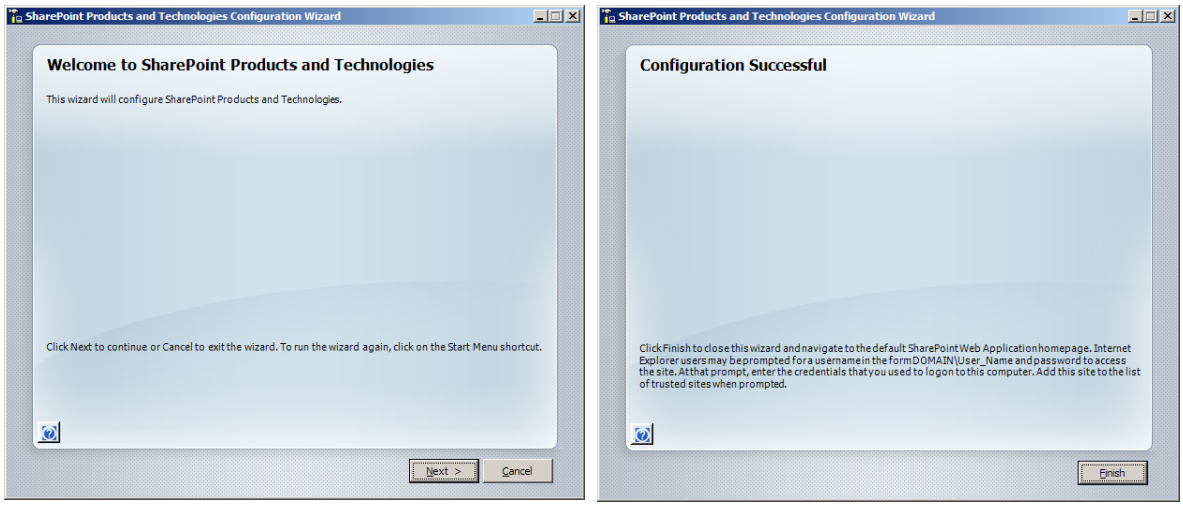

After the configuration wizard finishes, the SharePoint default home page is displayed and SharePoint is ready to use.

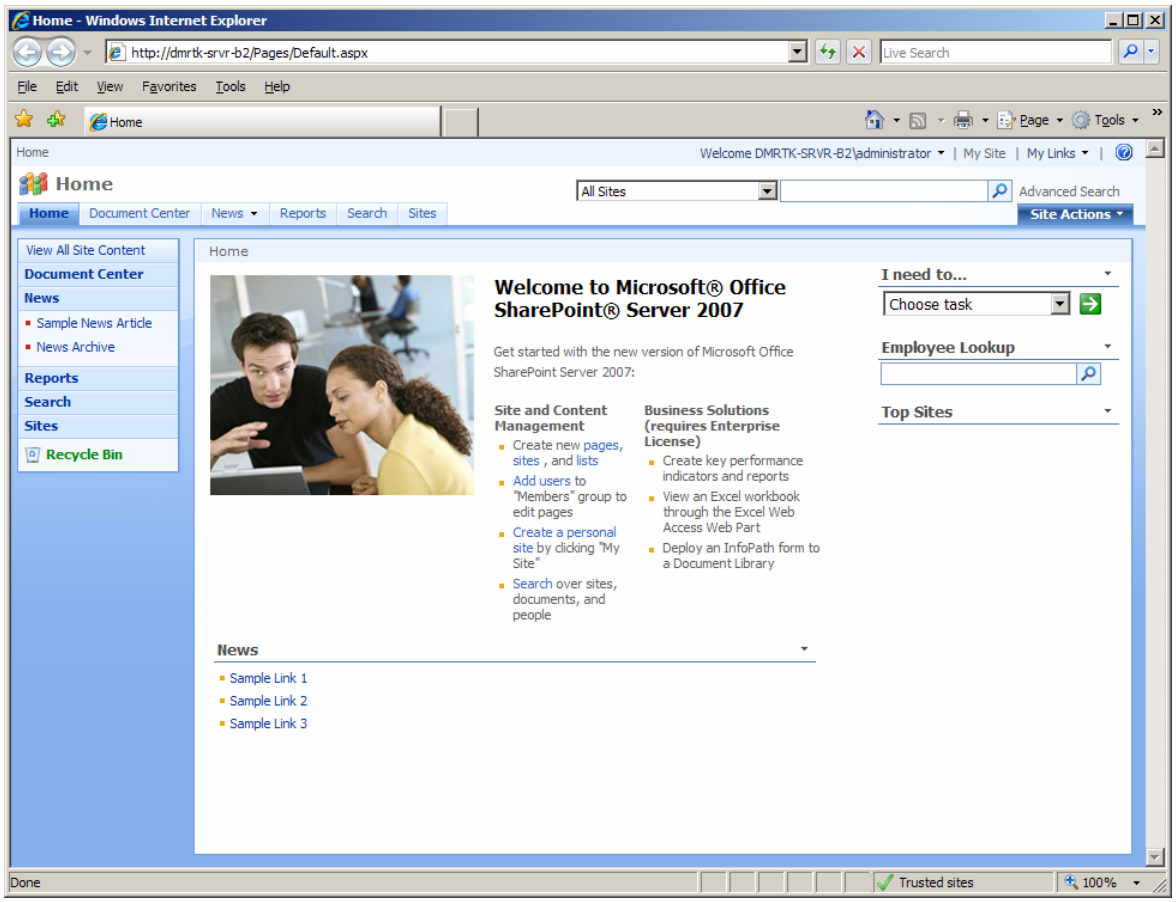

Microsoft Exchange Server 2007

Microsoft Exchange Server 2007 can be deployed using an iSCSI storage solution. The basic requirements for storage for Exchange Server 2007 are that there is enough capacity, acceptable disk latency and response time, and enough disk throughput to meet service level agreements. The disk storage must be formatted using NTFS. Due to its 64-bit nature, increases in database cache and other improvements, transactional I/O requirements for Exchange Server 2007 have been reduced from previous releases.

Unlike previous versions of Exchange Server, network-attached storage is not supported in Exchange Server 2007. The only network-based storage transports supported for Exchange Server 2007 are iSCSI and Fibre Channel. In Microsoft test labs, iSCSI has been proven as a capable storage transport for Exchange Server. When high throughput is required, Microsoft recommends multiple network adapters and the use of MPIO.

Microsoft provides detailed deployment guides for Exchange Server 2007. These deployment guides include a major chapter on the planning of disk storage for Exchange Server and cover a variety of topics including storage capacity, storage technology, RAID selection, LUN sizes, storage validation (i.e. Jetstress), and more. These deployment guides are available at: http://www.microsoft.com/exchange. In addition, Microsoft has created the *Exchange Solution Reviewed Program – Storage* (ESRP), which is a Microsoft Exchange Server program designed to facilitate third-party storage testing and solution publishing for Exchange Server. Details about ESRP are available at: http://technet.microsoft.com/en-us/exchange/bb412164.aspx.

Prerequisites for Deployment

The following steps were completed on the server before installing Exchange Server 2007:

- All updates applied to Windows 2003 Server R2 Enterprise x64 Edition
- Domain controller functional level raised to Windows Server 2003 level
- Server DMRTK-SRVR-B2 joined the DEMARTEK domain
- Microsoft .NET Framework 2.0 installed on server
- Microsoft .NET Framework 2.0 x64 HotFix KB926776 applied
- Microsoft Windows PowerShell installed
- iSCSI targets logged on as volumes “M” and “N”

Exchange Server 2007 Deployment

Exchange Server 2007 provides a wizard to perform the installation.

Steps 1, 2 and 3 had already been completed before the Exchange Server 2007 installation began. The wizard was launched, license agreement reviewed and installation choices made.

A “typical” installation was performed. The Exchange Server will be installed in the default location. The Exchange organization was given as “Demartek-Lab”. Outlook 2003 clients are available in this environment, so the public folder database will be created.

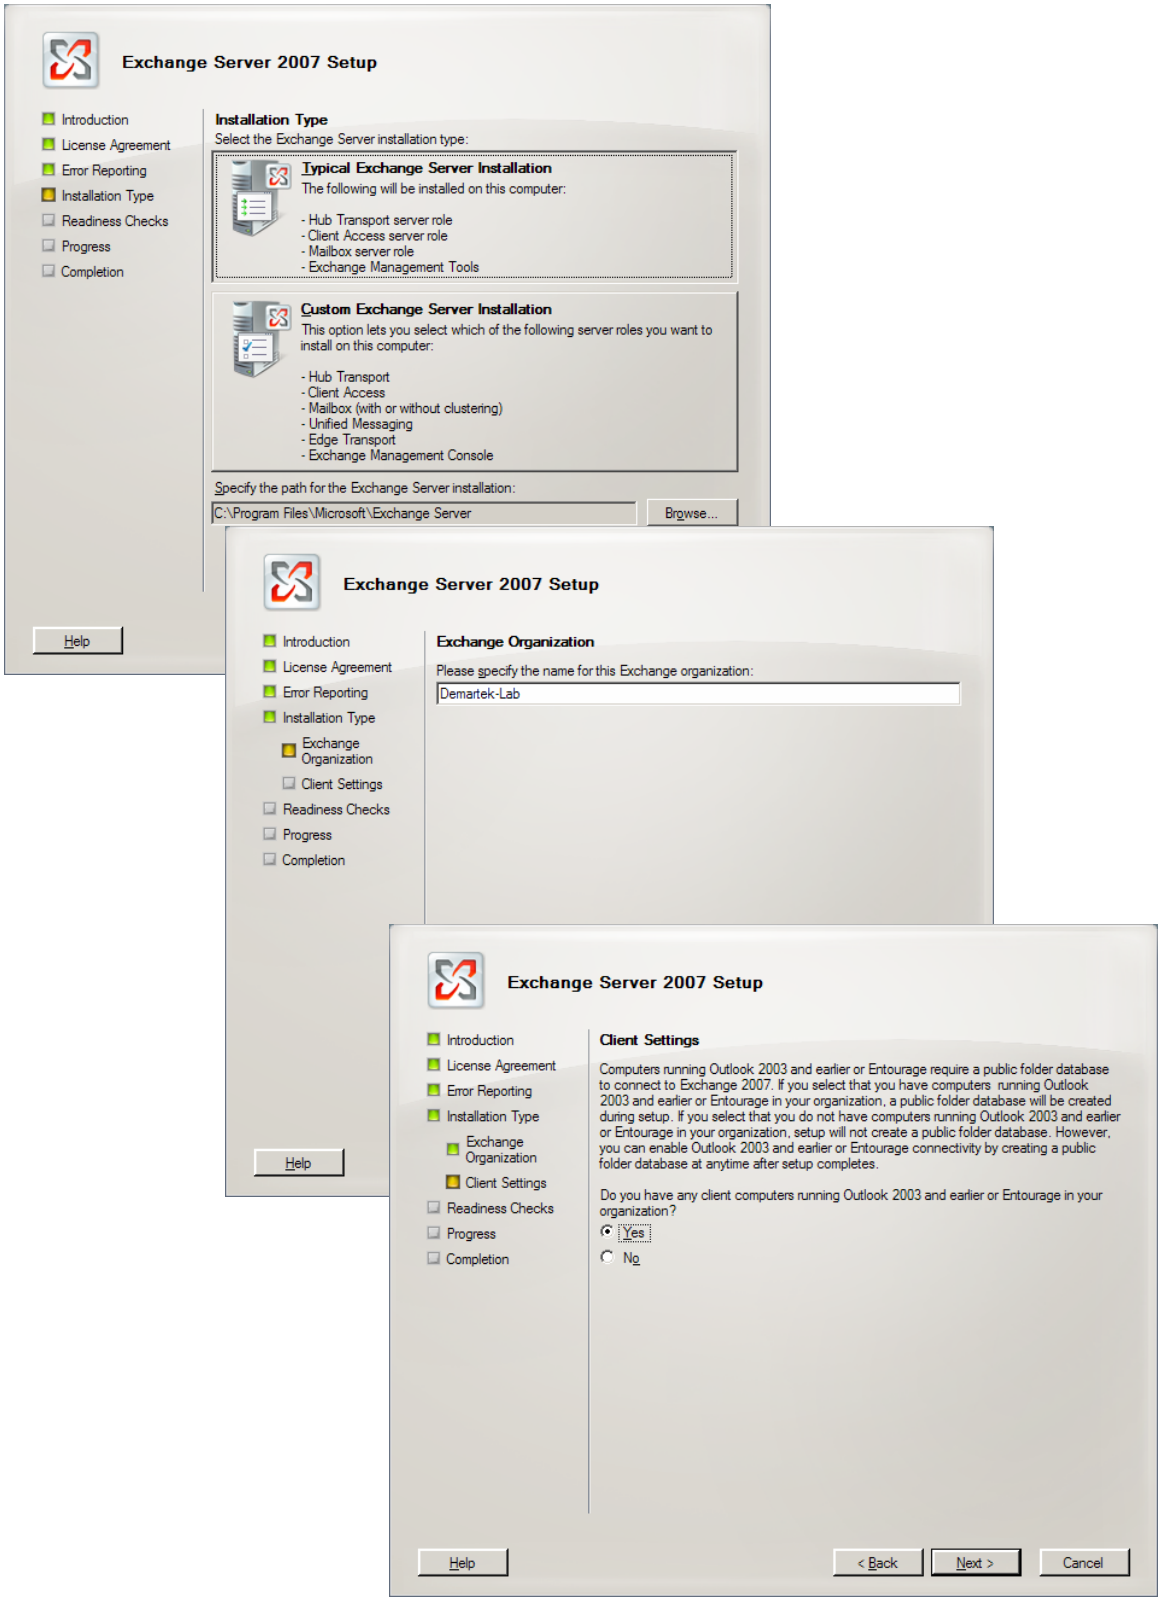

After providing the data needed, the wizard performed the Readiness Checks then performed the actual installation.

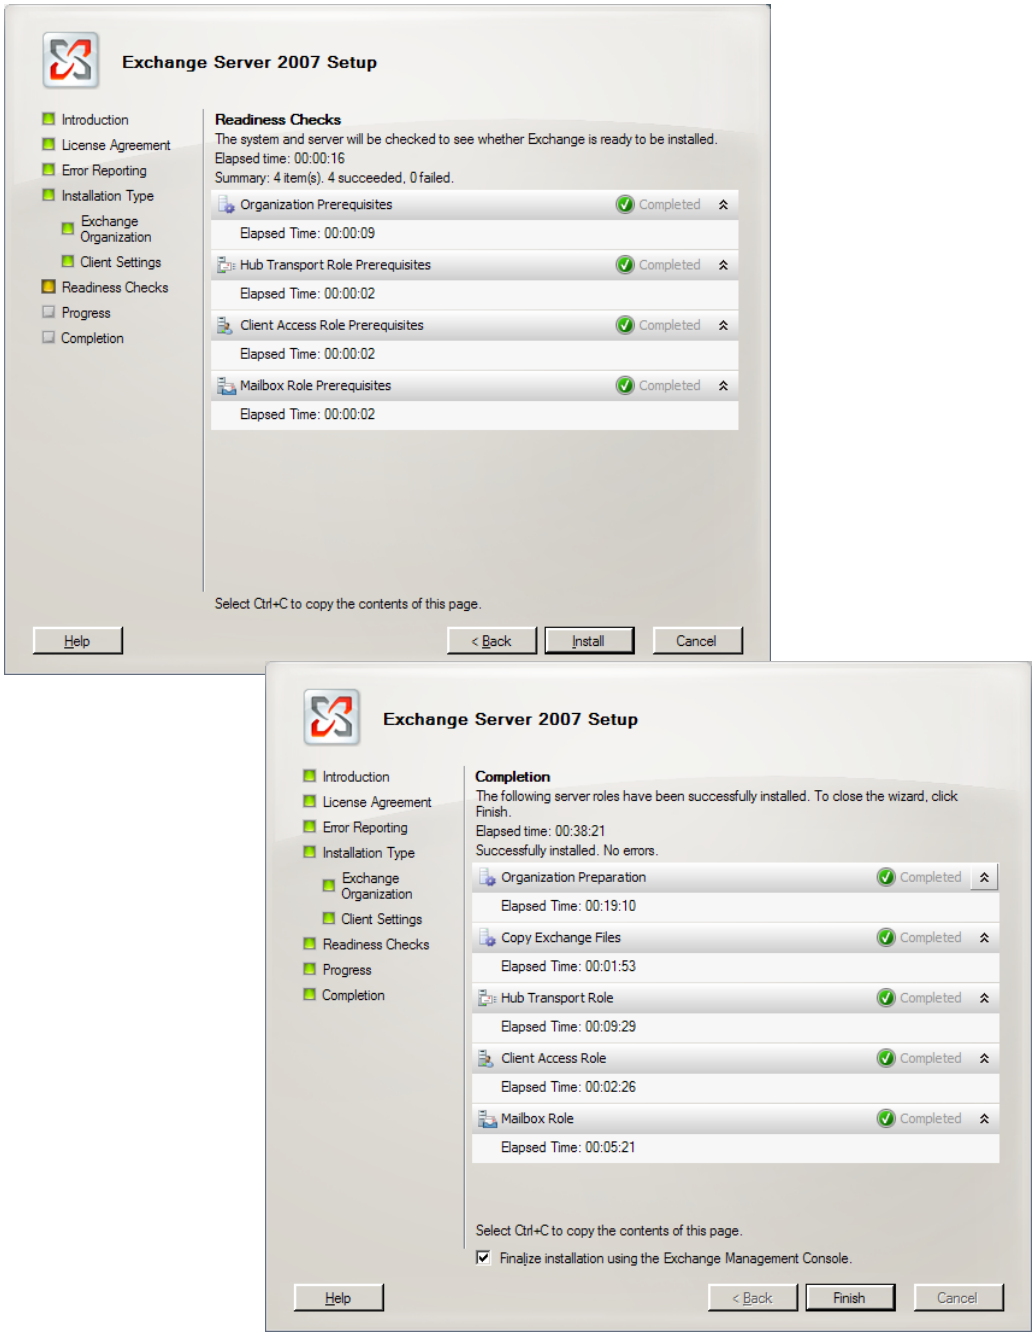

Update Rollup 1 for Exchange Server 2007 was applied, and then the Exchange Server Management Console was launched.

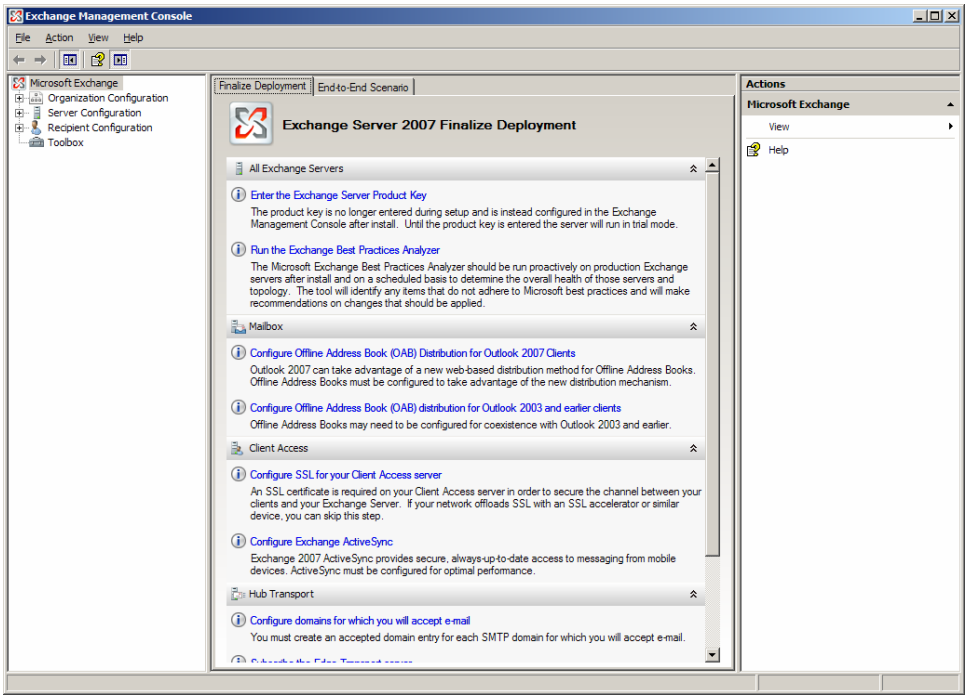

Clicking the “Mailbox” item in the left panel revealed the details of the Exchange storage groups.

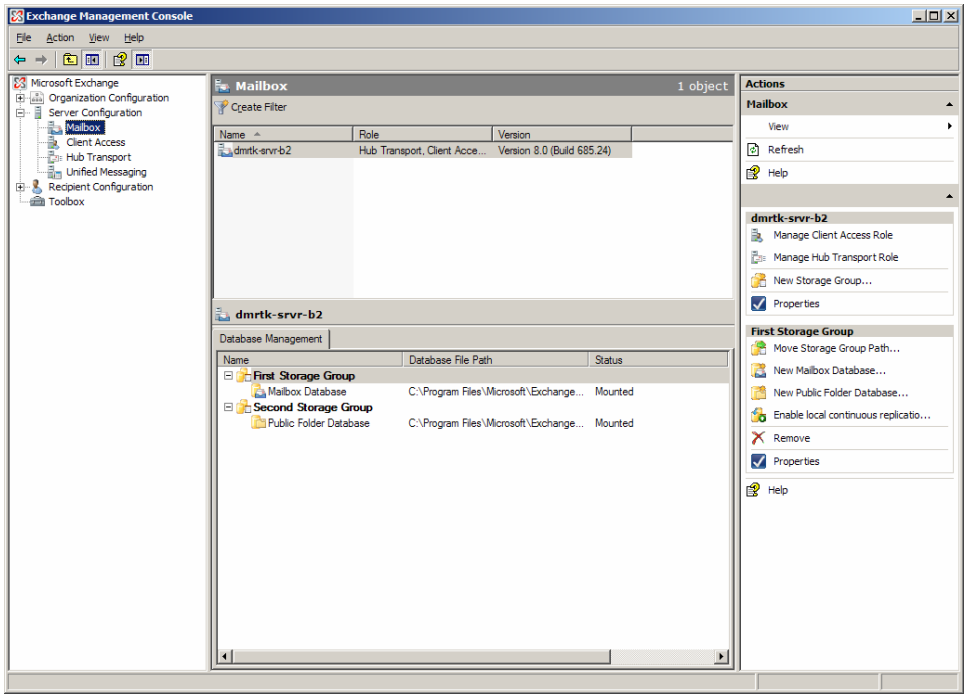

A new storage group will be created on the iSCSI target devices. The logs will be placed on the "M" volume and the data store will be on the "N" volume.

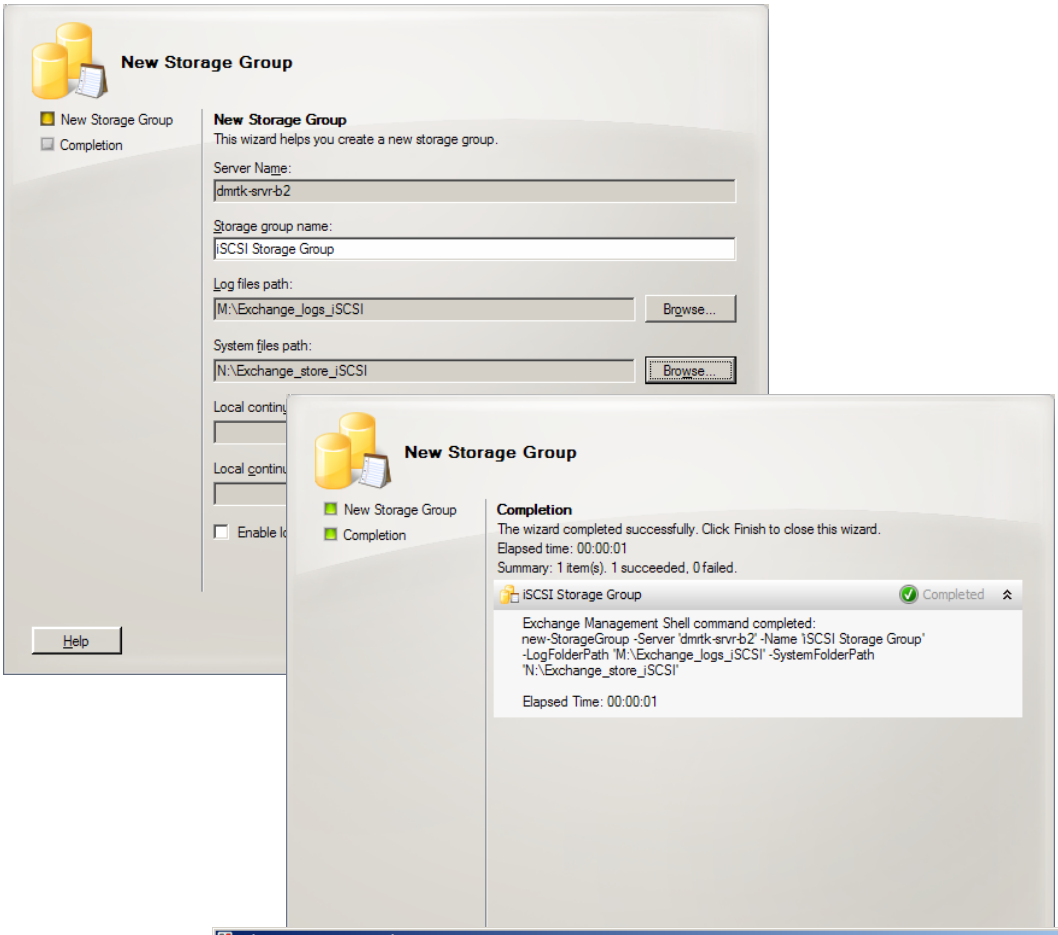

The iSCSI Storage Group now appears with the First Storage Group.

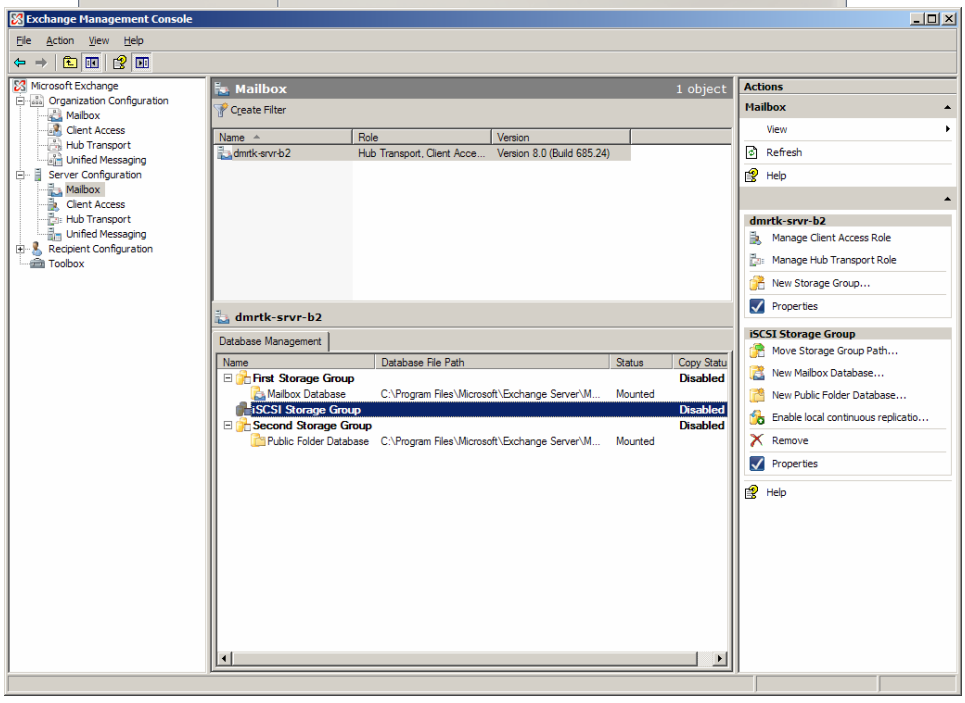

A new mailbox database will be created and mounted on the “N” iSCSI target volume.

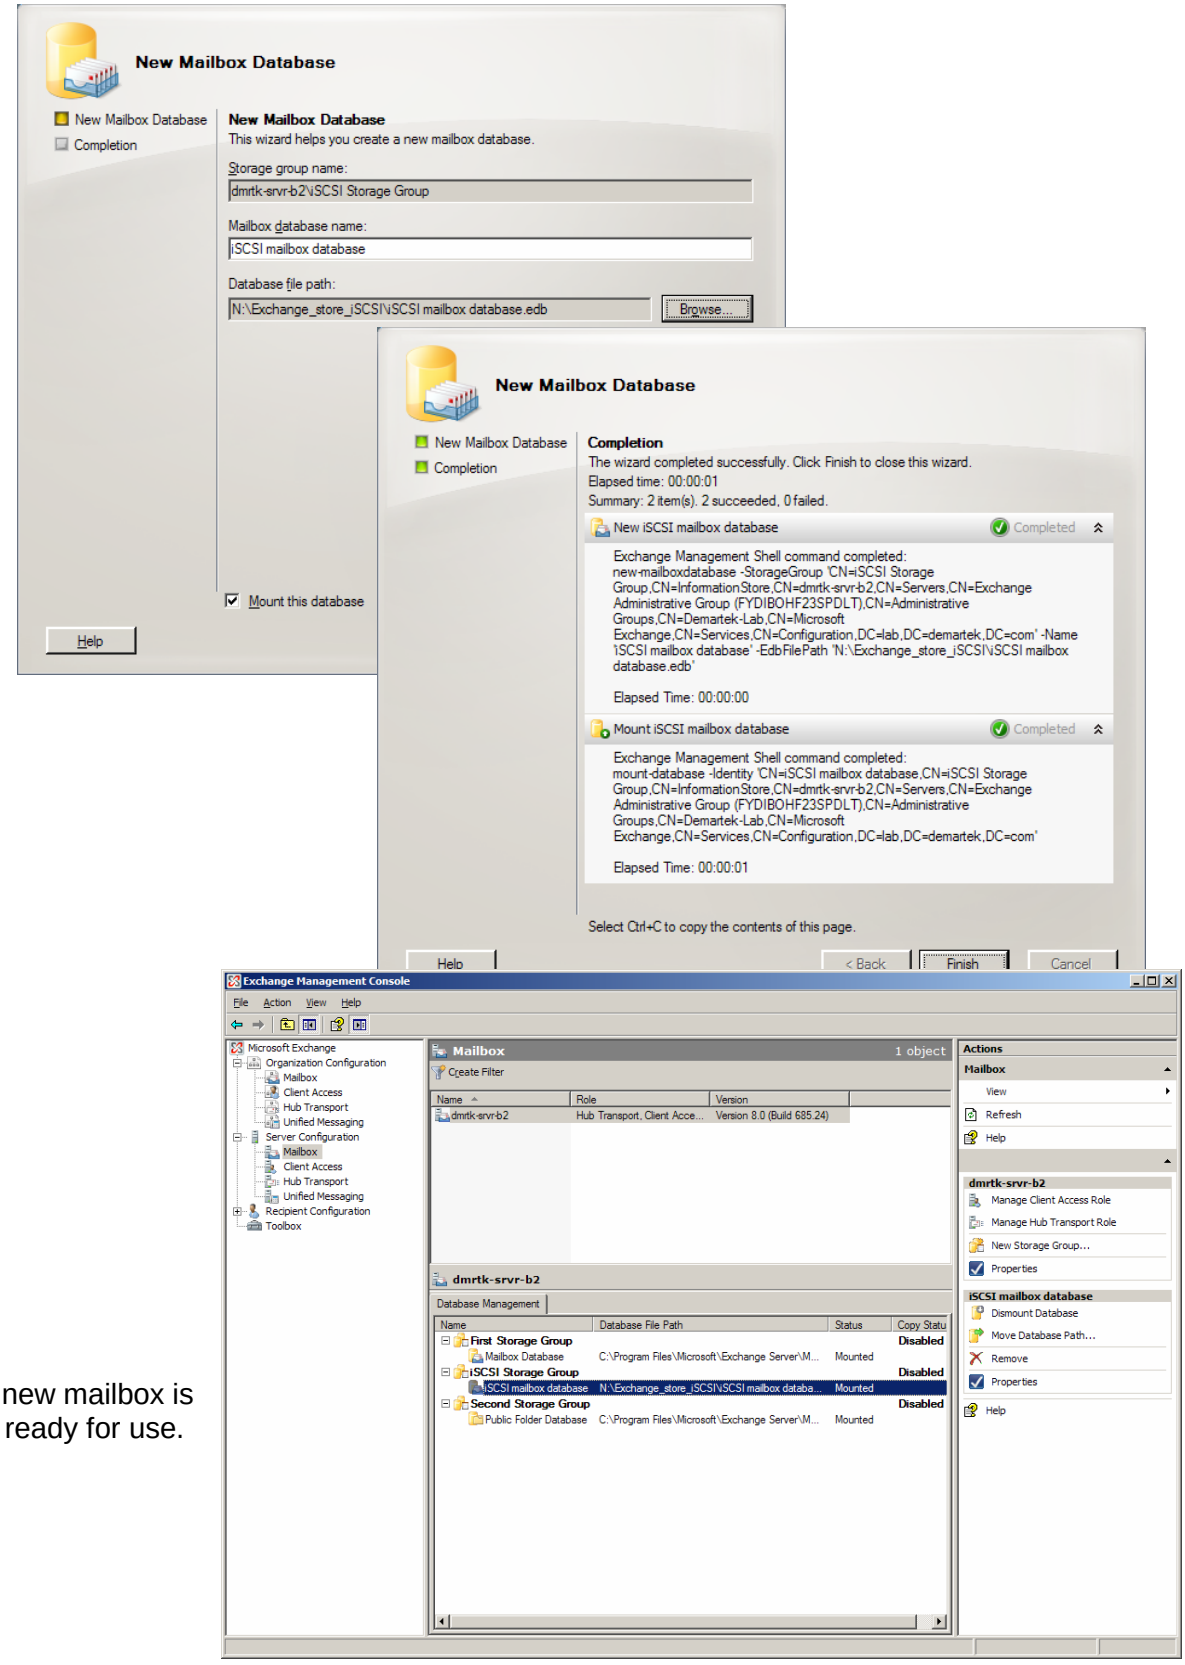

The new mailbox is now ready for use.

Microsoft SQL Server 2005

SQL Server 2005 can use iSCSI storage solutions for databases, log files and other SQL-related files in much the same way that it can use other storage technologies. Network-attached storage (NAS) storage is not recommended for SQL Server data, but iSCSI is a capable storage transport for SQL Server data. High throughput with iSCSI can be obtained by using multiple network adapters with MPIO. Microsoft provides extensive information regarding SQL Server at http://www.microsoft.com/sqlserver.

Prerequisites for Deployment

SQL Server 2005 was installed onto server DMRTK-SRVR-A2, which was running Windows 2003 Server R2 Enterprise x64 Edition. The following steps were completed on the server before installing SQL Server 2005:

- All updates applied to Windows 2003 Server R2 Enterprise x64 Edition
- Microsoft .NET Framework 2.0 installed on server
- iSCSI targets logged on as volumes “Q” and “R”

SQL Server 2005 Deployment

SQL Server 2005 provides a wizard to perform the installation. Most of the defaults were selected. First the wizard checked for SQL Server prerequisites.

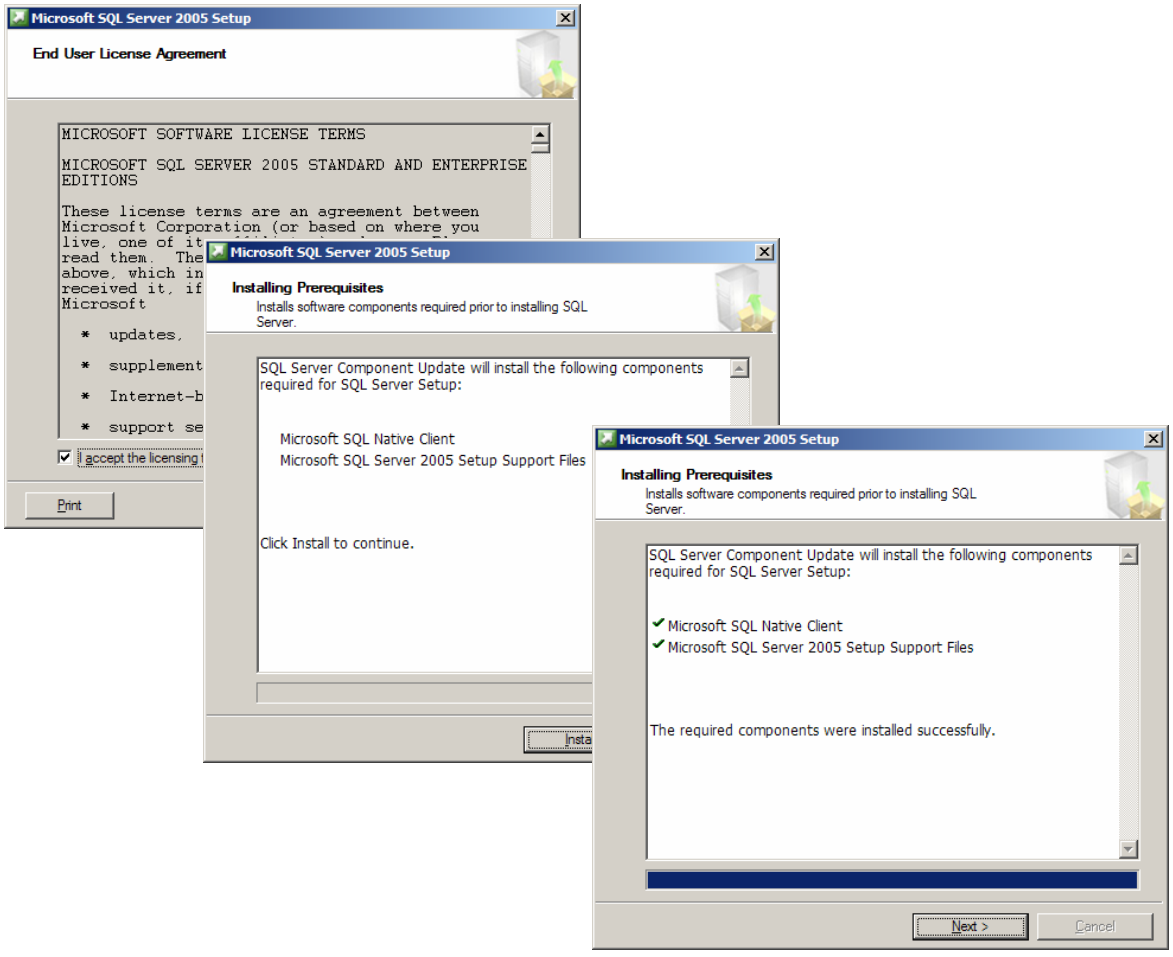

After the SQL Server prerequisites were completed, the main installation began.

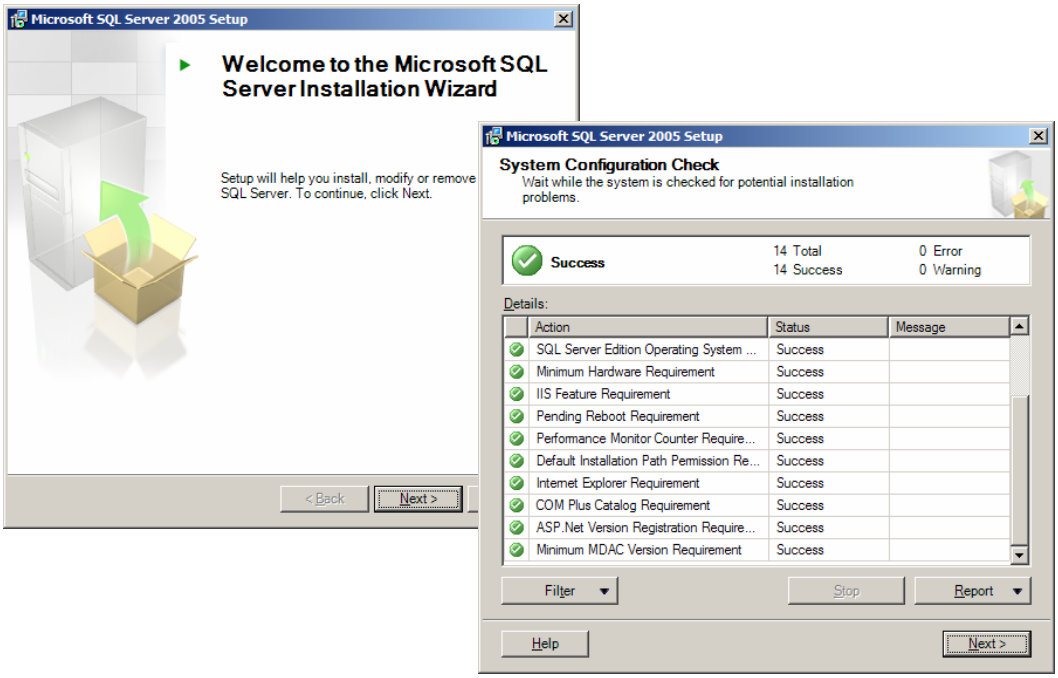

All the SQL Server components were installed. Most of the other defaults were chosen.

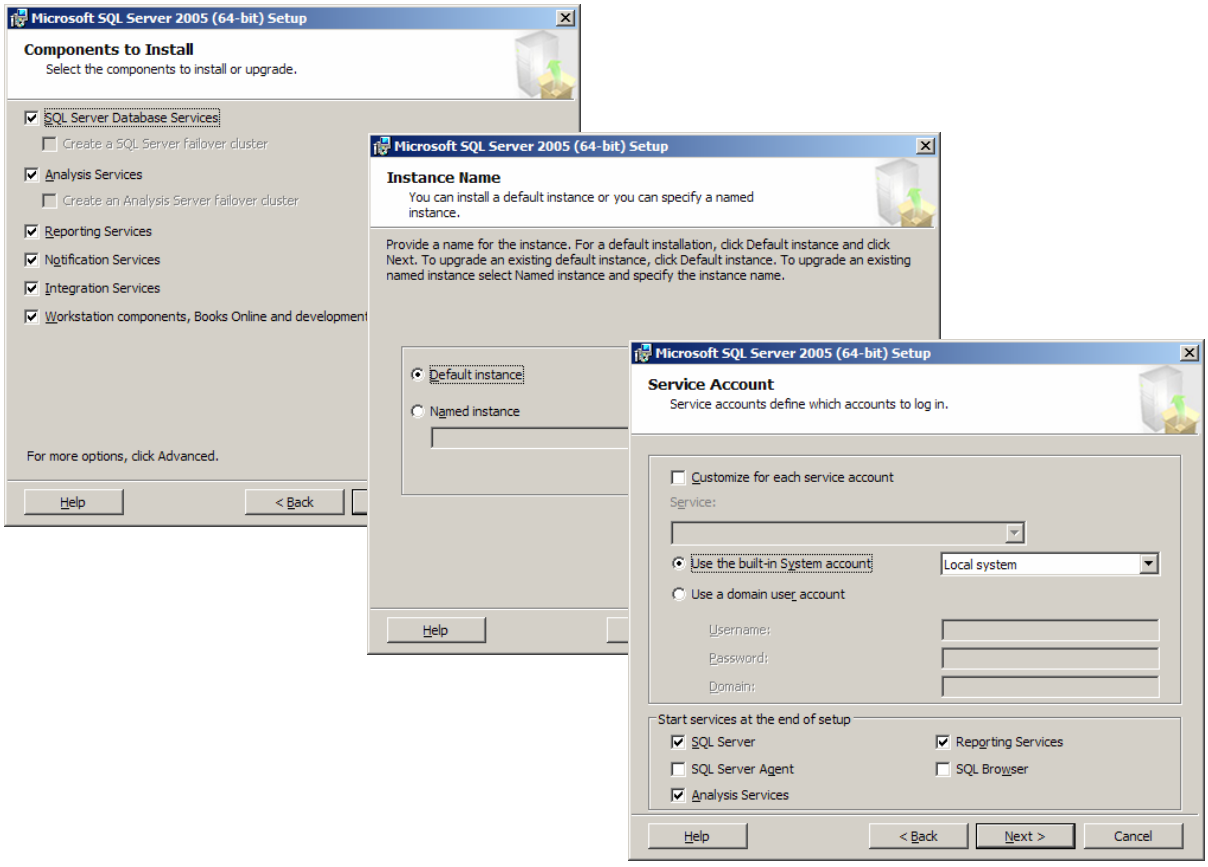

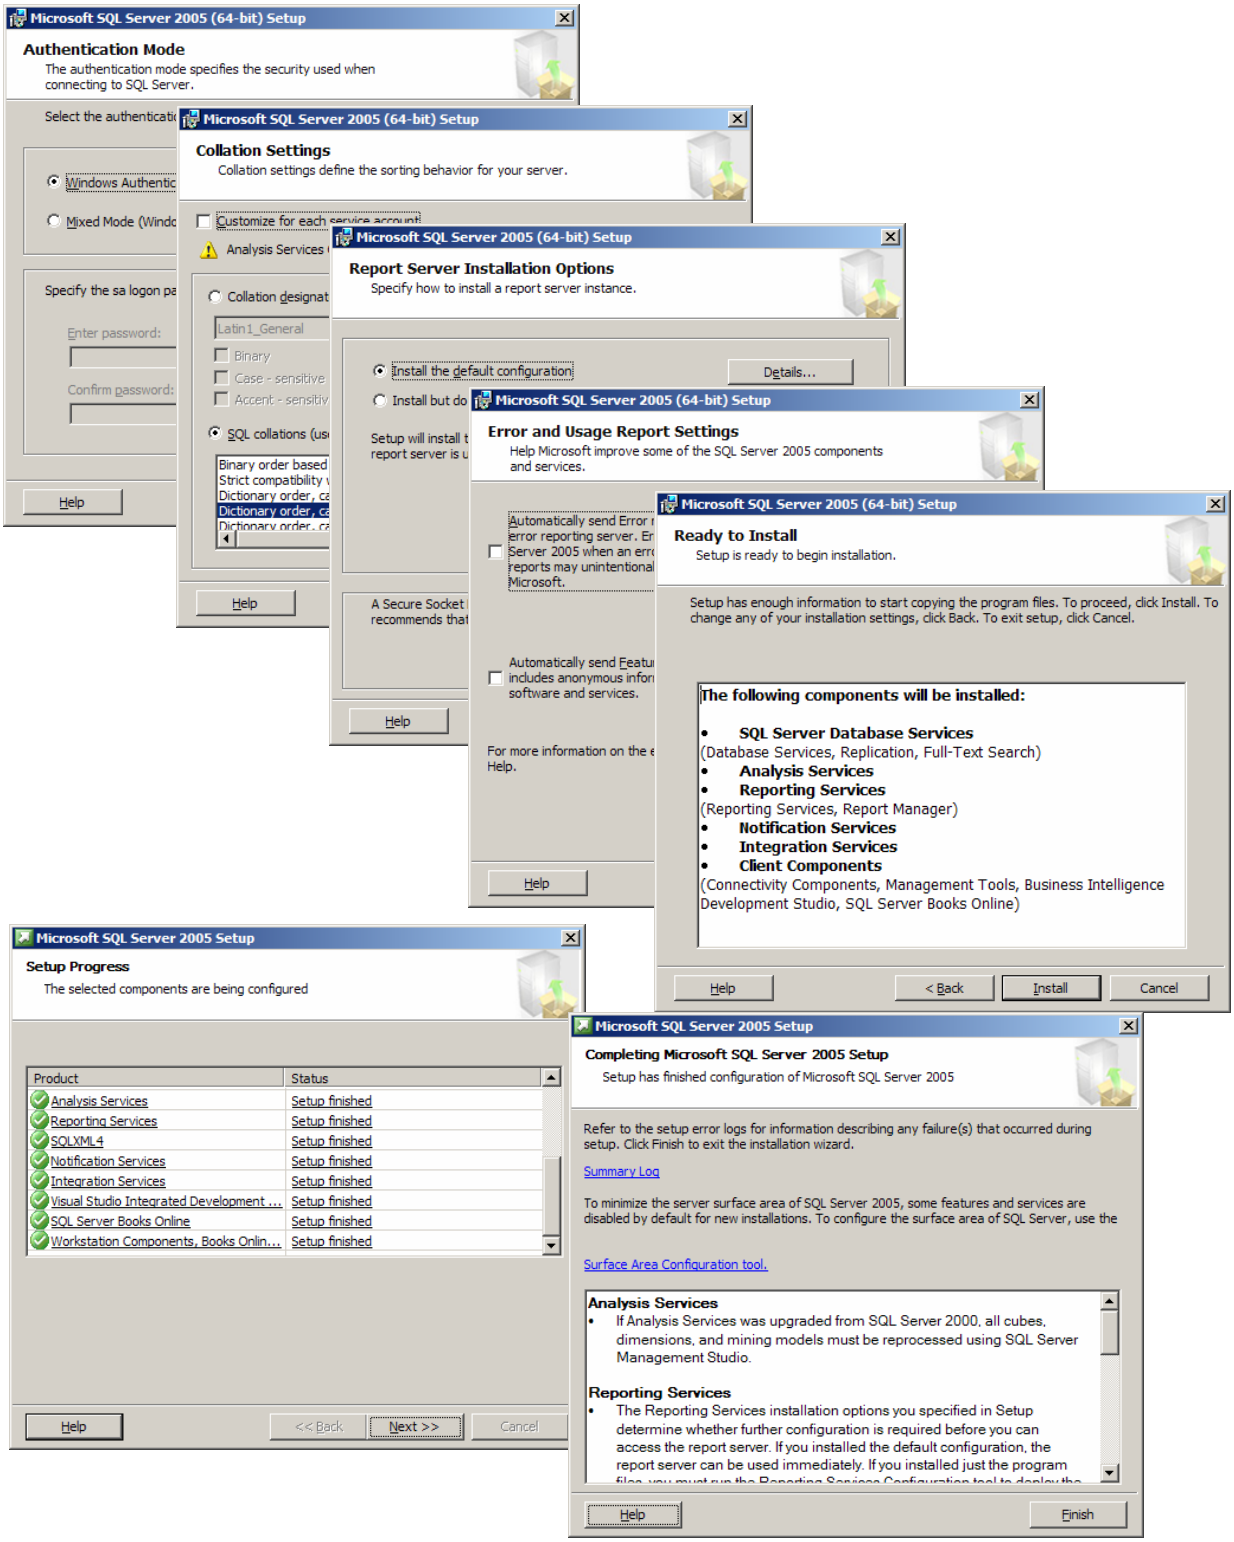

The installation completed successfully. After the installation, the system was rebooted. SQL Server 2005 Service Pack 2 was installed after the reboot.

The sample AdventureWorks databases were installed into the default SQL Server file location. The Copy Database Wizard was used from within SQL Server Management Studio to copy the AdventureWorks database and AdventureWorksDW (data warehouse) to the iSCSI target volumes “Q” and “R”. Volume “Q” was used for the data and volume “R” was used for the SQL logs. The source and destination server were the same, but the destination file location was set to the iSCSI target volumes.

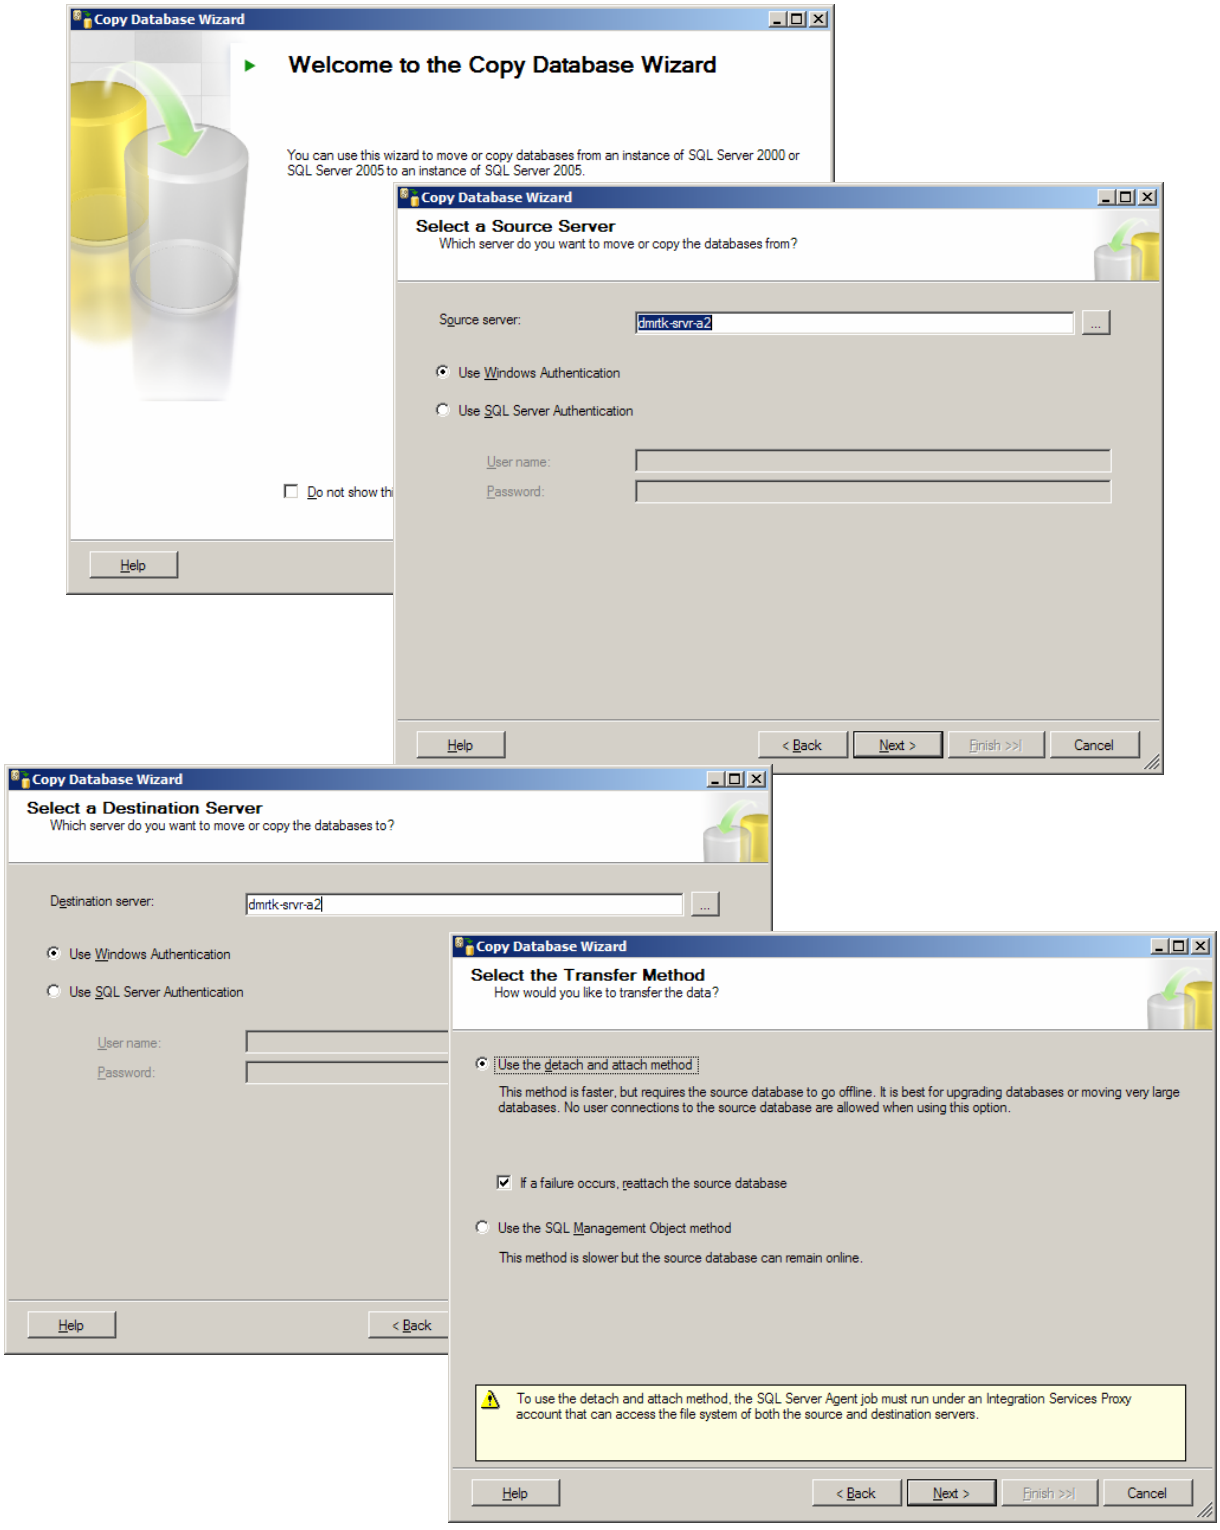

The wizard prompts for the file locations. The source files were in the default location from installation. The destination files were the iSCSI target volumes. Two databases, the AdventureWorks and AdventureWorksDW were copied. The data files were copied to the “Q” volume and the logs were copied to the “R” volume.

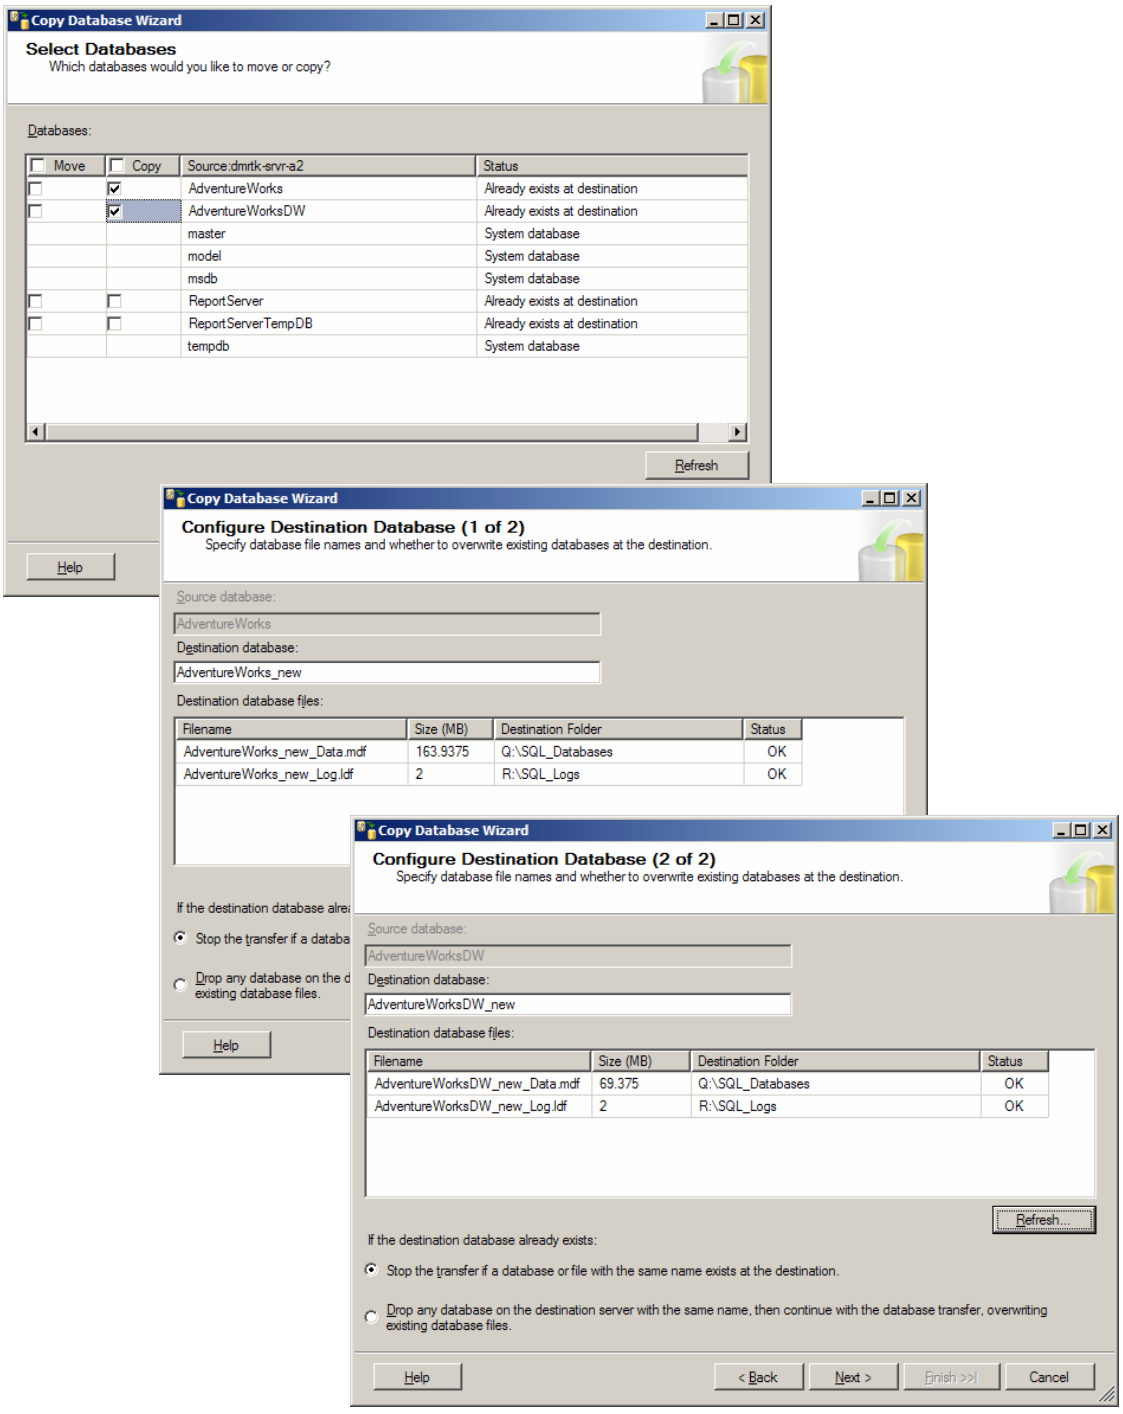

The Integration Services package was configured and run immediately with the SQL Server Agent service account.

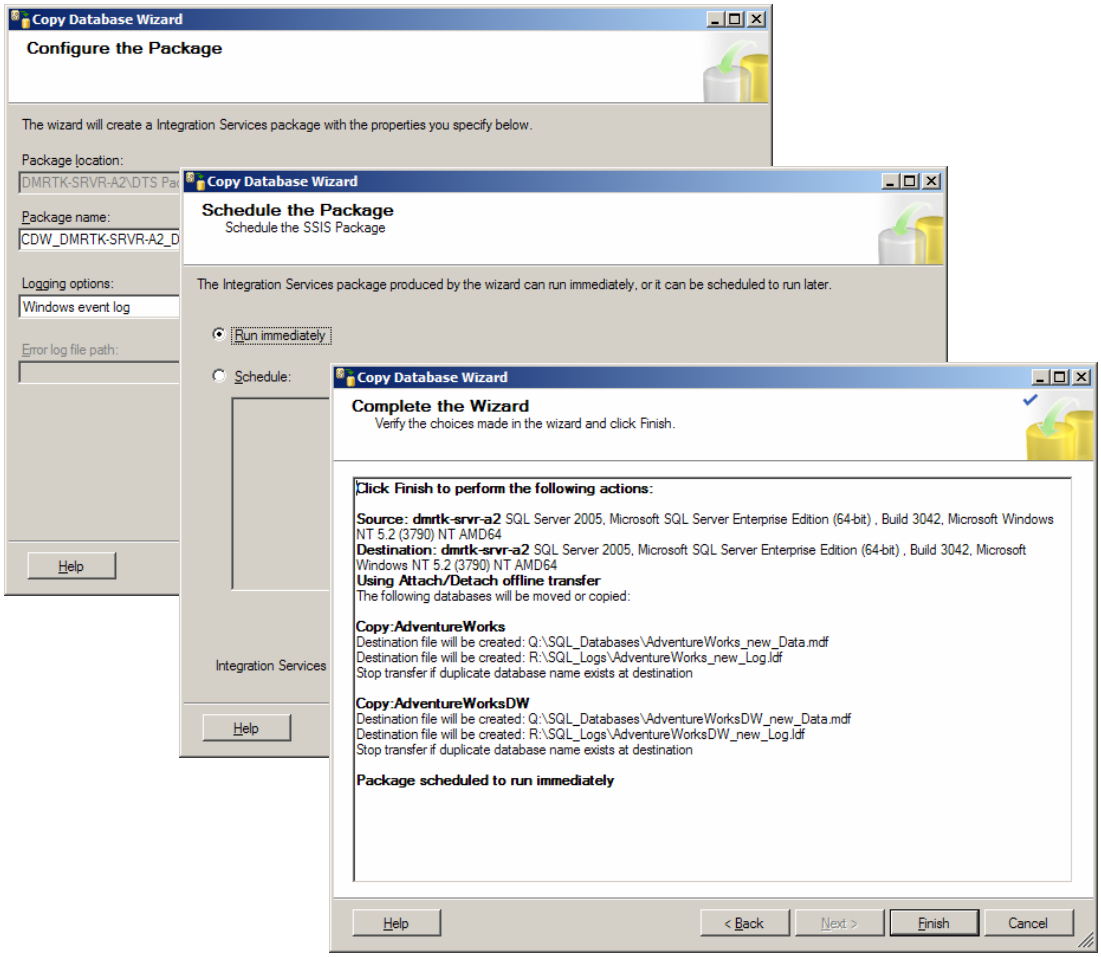

The copy was completed successfully to the iSCSI target volumes using SQL Server services.

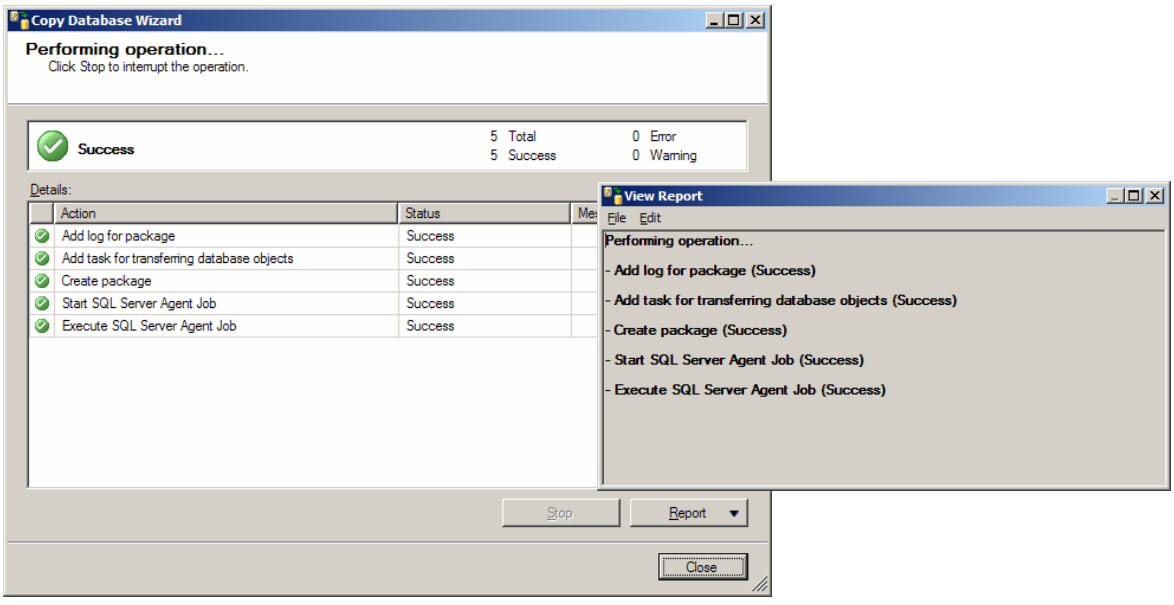

Dell™ PowerVault™ NX1950 Networked Storage Solution

The Dell PowerVault NX1950 is a networked storage system that features Microsoft Unified Data Storage Server (WUDSS). Optimized for performance and interoperability, the system supports both file (sharing) and block (disk) access to storage resources over Ethernet networks through the inclusion of the iSCSI protocol. The PowerVault NX1950 comes bundled with advanced software capabilities that include snapshots and remote replication and is configurable for clustering to maintain high system and data availability. The PowerVault NX1950 can be expanded by simply adding up to two additional expansion enclosures for a total of 45 drives and 13.5TB capacity, to support growing business needs.

Target Configuration Steps

1. Configure Network Settings for iSCSI Target Device

The Dell PowerVault NX1950 is configured to use DHCP for its default network settings. The basic unit is designed for multi-path operations and is equipped with four RJ45 Ethernet connectors. The initial configuration screen shows the basic settings.

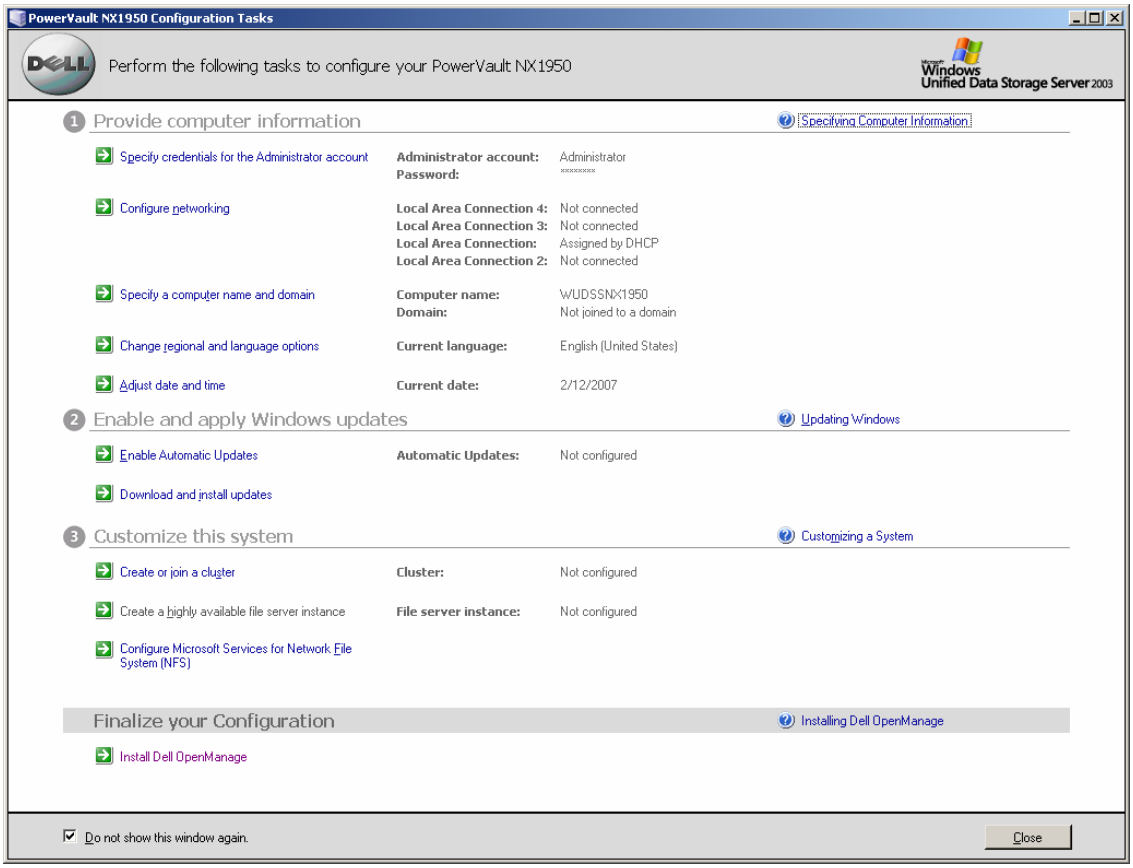

2. Launch Management Console

All the storage management functions for the Dell PowerVault NX1950 are performed from the management console, shown below.

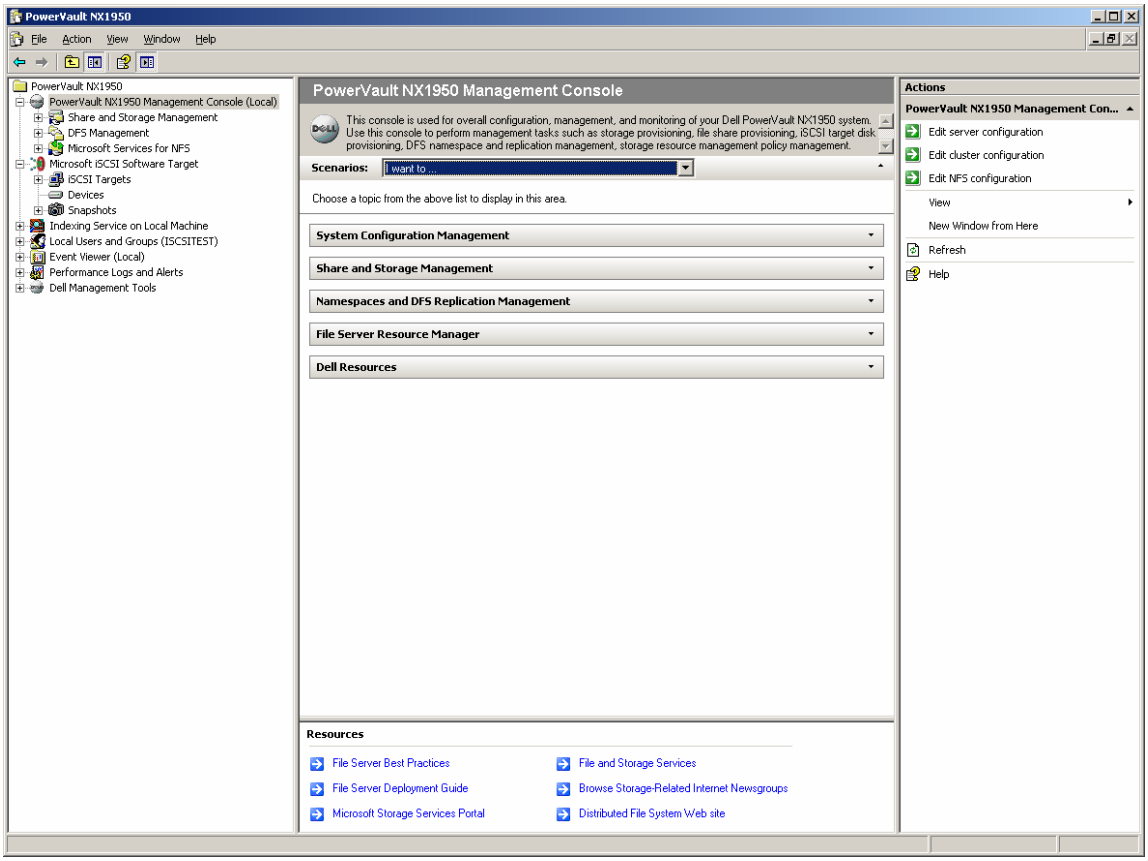

At the top of the center section of the management console, several scenarios are available that help step the administrator through each of the processes.

3. Create LUNs on Disk Array

To create the LUNs on the Disk Array, the administrator selects the “Provision Storage and Create Volume” scenario from the upper section of the management console. This directs the administrator in the appropriate steps to take.

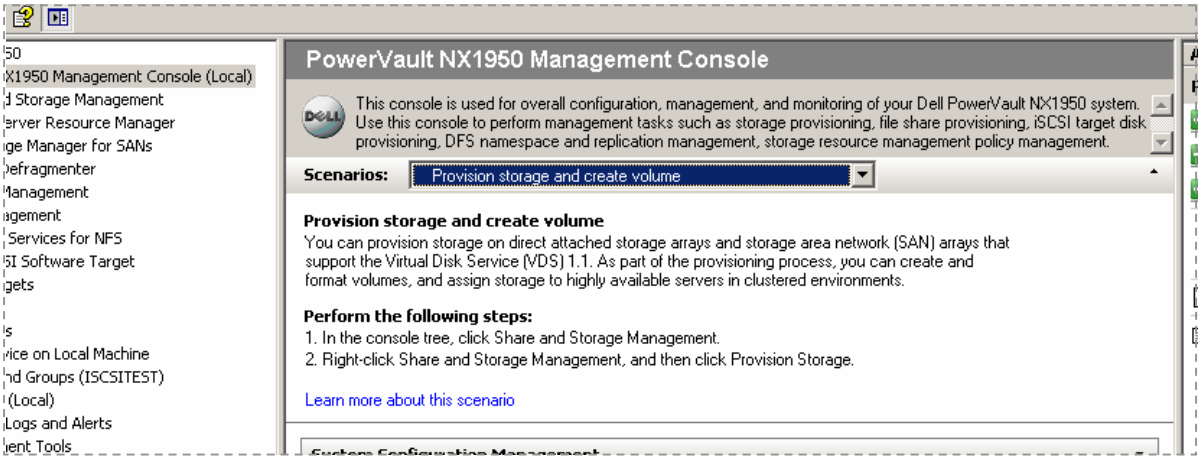

The right-side panel of the management console is context sensitive, and changes based on the item selected on the left side of the panel. Highlighting “Share and Storage Management” on the left panel console tree brings the “Provision Storage” wizard into view on the right side.

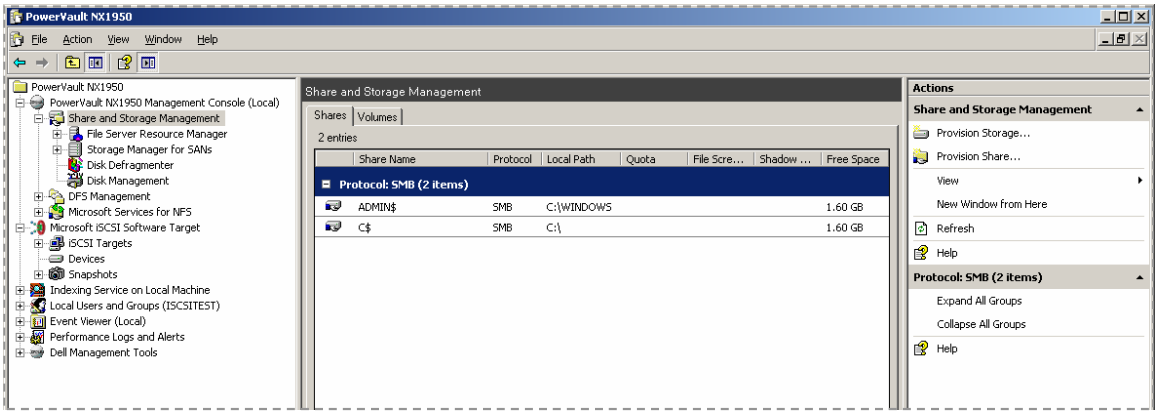

Selecting the “Provision Storage” action on the right side initializes the wizard for this function. The wizard allows the administrator to step through the provisioning process. Notice that the left side of the wizard lists each of the main steps in the process.

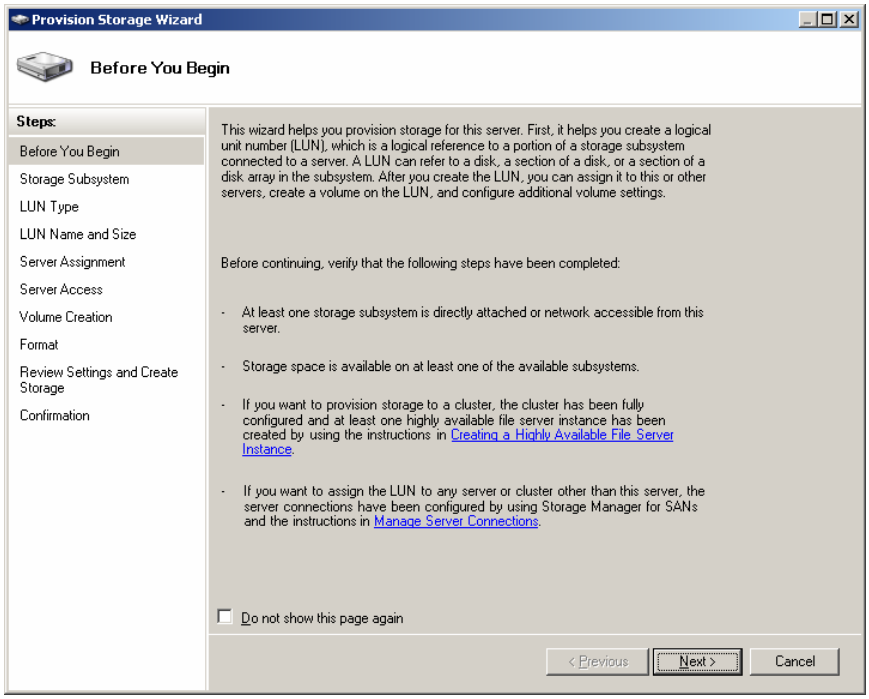

The storage subsystem must be selected and in this case there is only one subsystem to select. It is of type “Fibre Channel” because the internal interface to the disk subsystem is listed as Fibre Channel even if SAS is used. This will be changed in the next version.

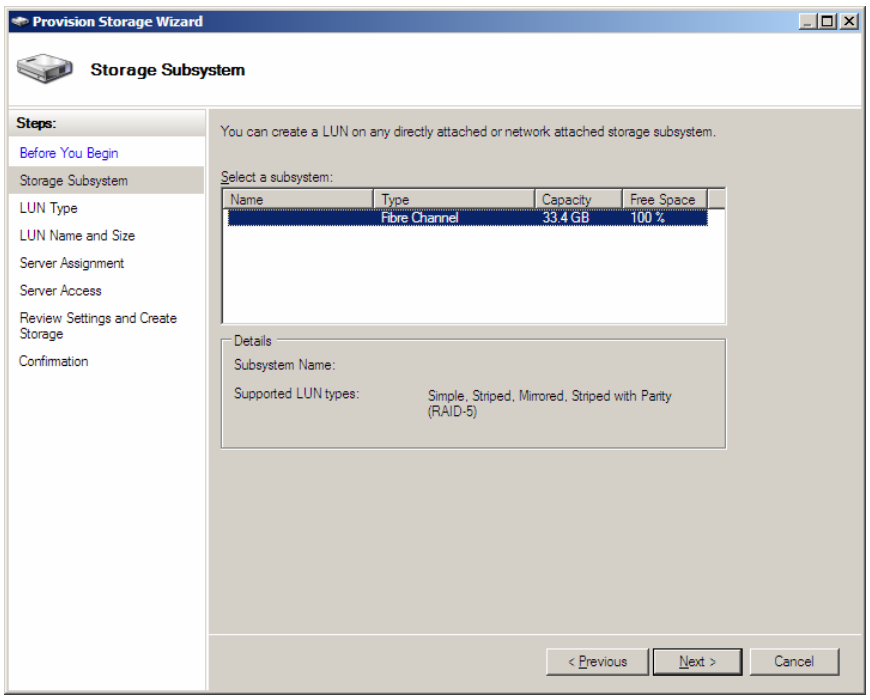

The LUN type must be chosen from among the various types of LUNs available. In this example, we have chosen a “striped” LUN type. Each type of LUN has a maximum size that depends on the type.

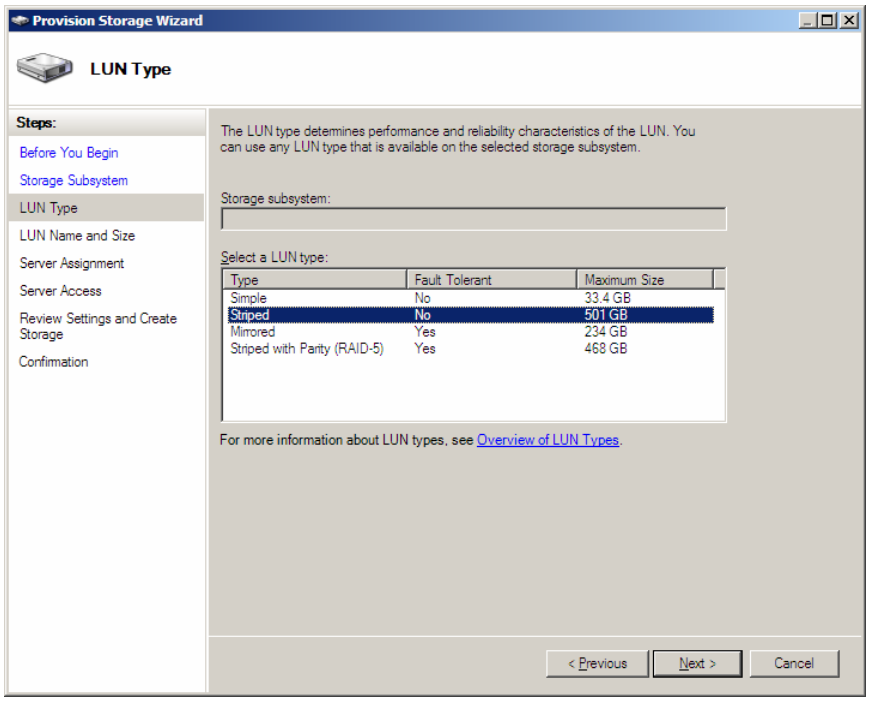

Note – It is important at this point to note that the storage solution LUN size should not be confused with the size of the iSCSI target. The iSCSI target will be configured in a later step and is associated with the storage needed for a particular application on the host server. It is recommended that the LUN size on the storage hardware be as large as reasonably possible to allow the storage subsystem to optimize the use of the

physical disks underlying the LUN that is created. In this case, as shown below, we are choosing to create one LUN at the maximum size available for this hardware. The iSCSI targets created later will fit into this one LUN, based on the needs of the host application.

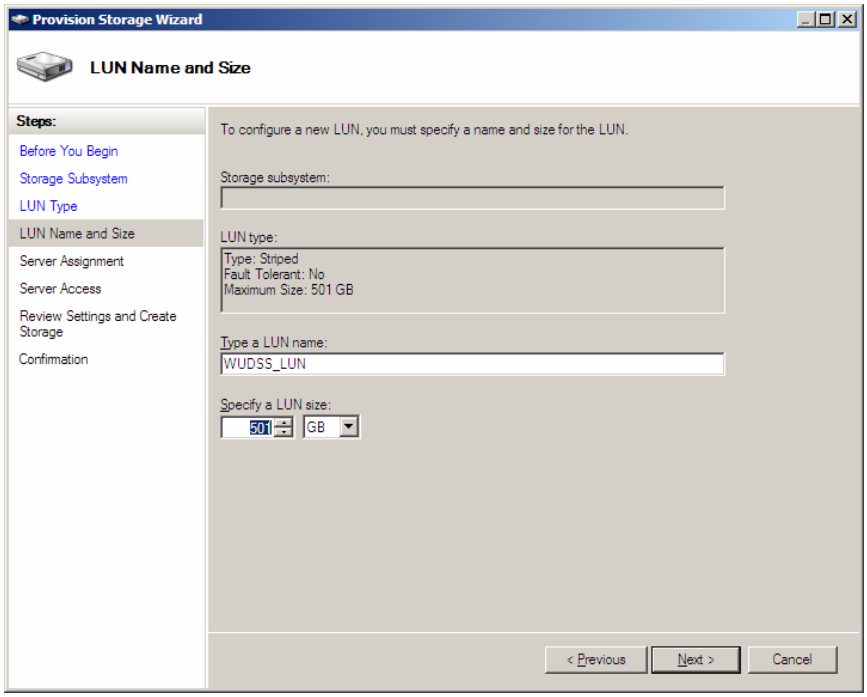

The LUN created will be assigned to this internal storage server only. In a later step, iSCSI targets will be created that will be assigned to external application servers.

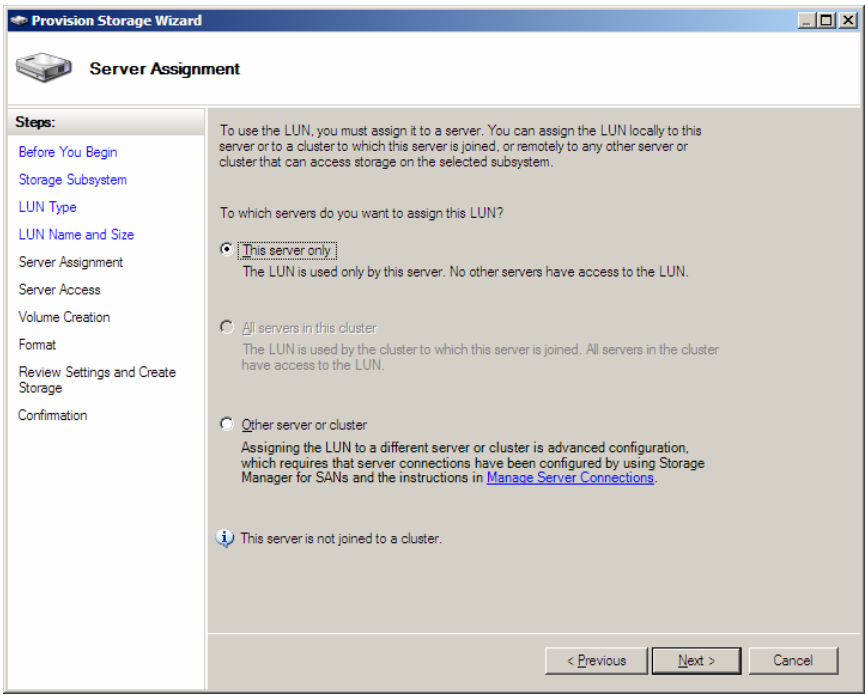

The name of this server needs to be provided for the assignment.

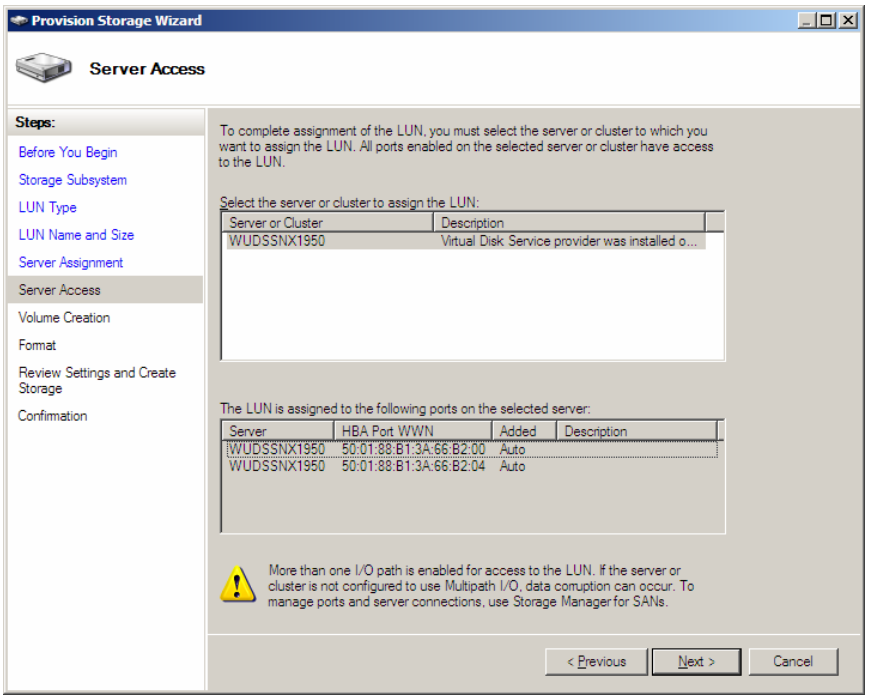

4. Make LUNs Ready for Use

Because this storage solution is running on a Microsoft Windows-based platform, the remaining steps would be familiar to a Windows administrator. This can be an advantage in environments where Windows is prevalent as it reduces specialized training needed for managing the storage devices. These include assigning a drive letter for the internal server, providing a volume name, etc. The wizard prompts for these items then provides a summary screen before performing all the necessary tasks to provision the storage.

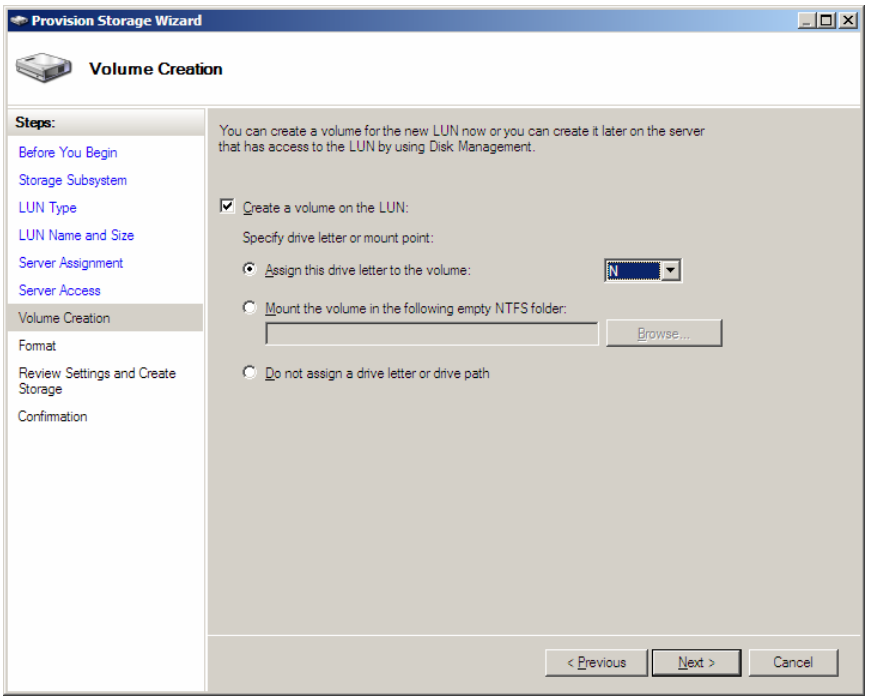

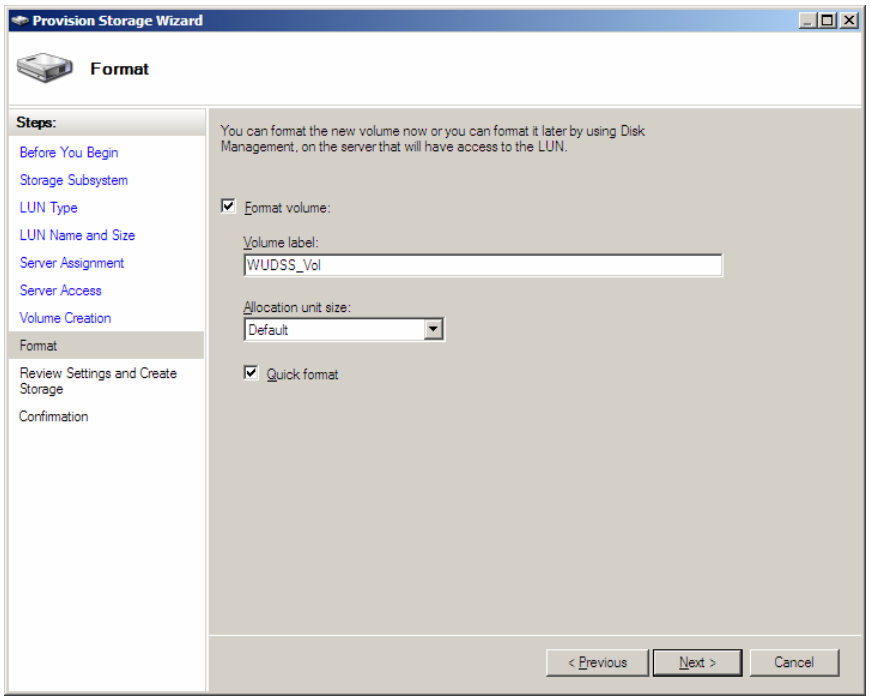

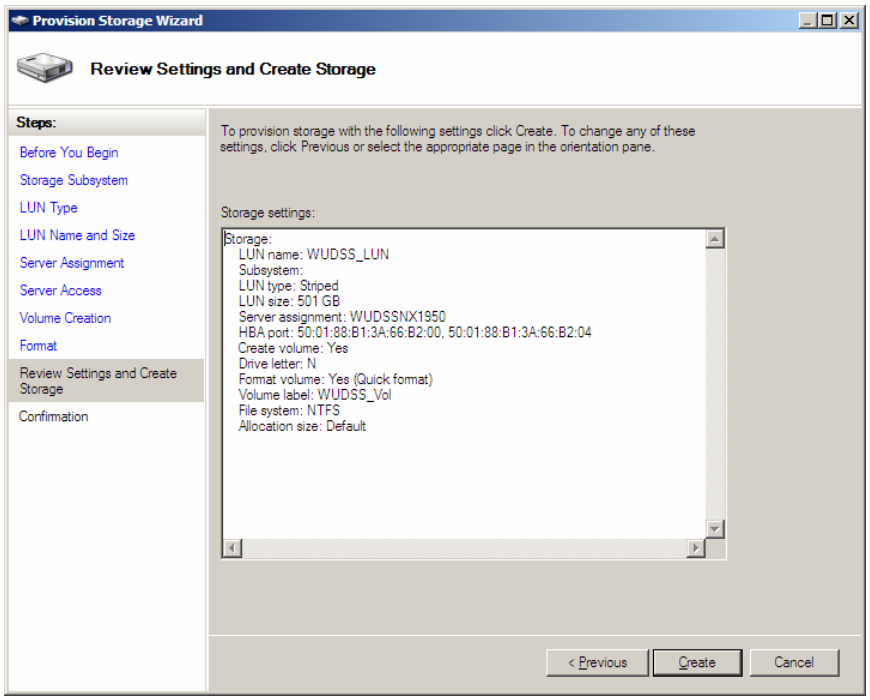

After a short while, the following screen indicates a successful provisioning operation.

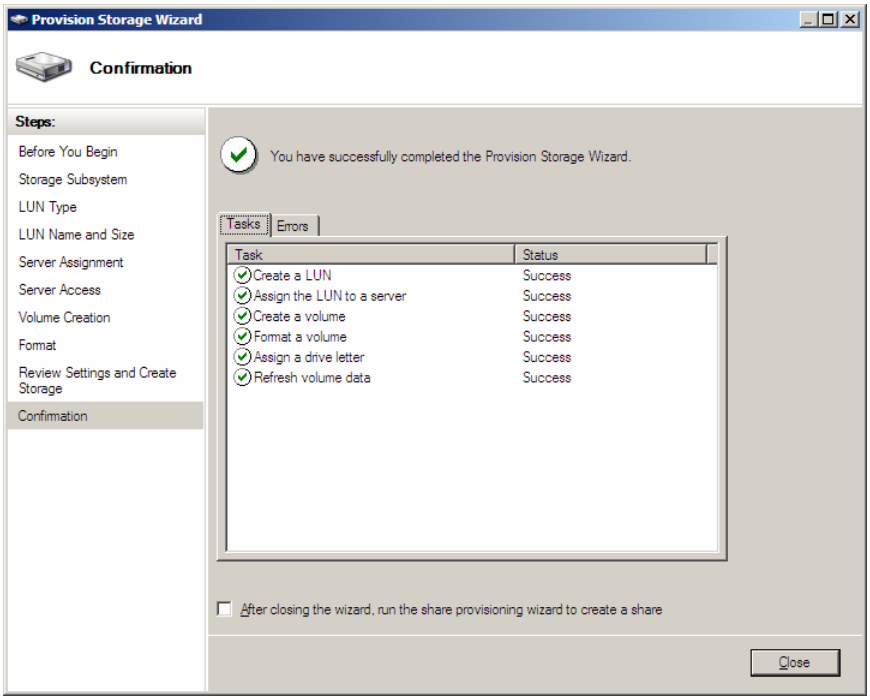

The LUN has now been created and is ready for use. The next step will create iSCSI targets and associate them with this newly-created LUN. This implementation of WUDSS uses the Microsoft Virtual Disk Service (VDS) internally on this server. The LUN can also be viewed in the Storage Manager for SANs section of the management console, as shown below.

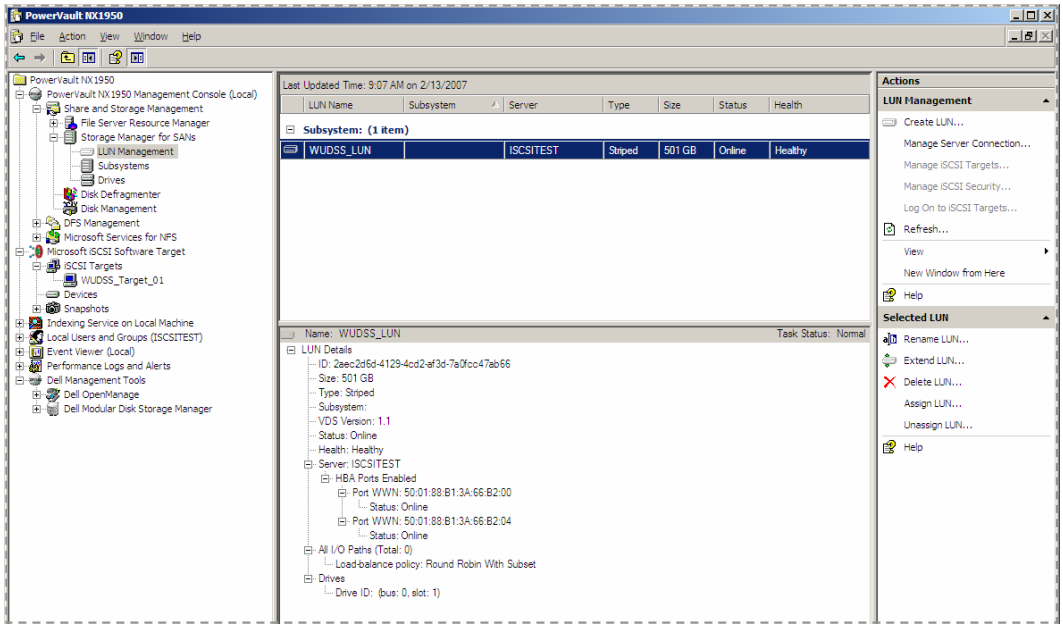

5. Create iSCSI Targets

Moving to the iSCSI target section of the management console, a wizard can be triggered (using the right-mouse click) that begins the iSCSI target creation process.

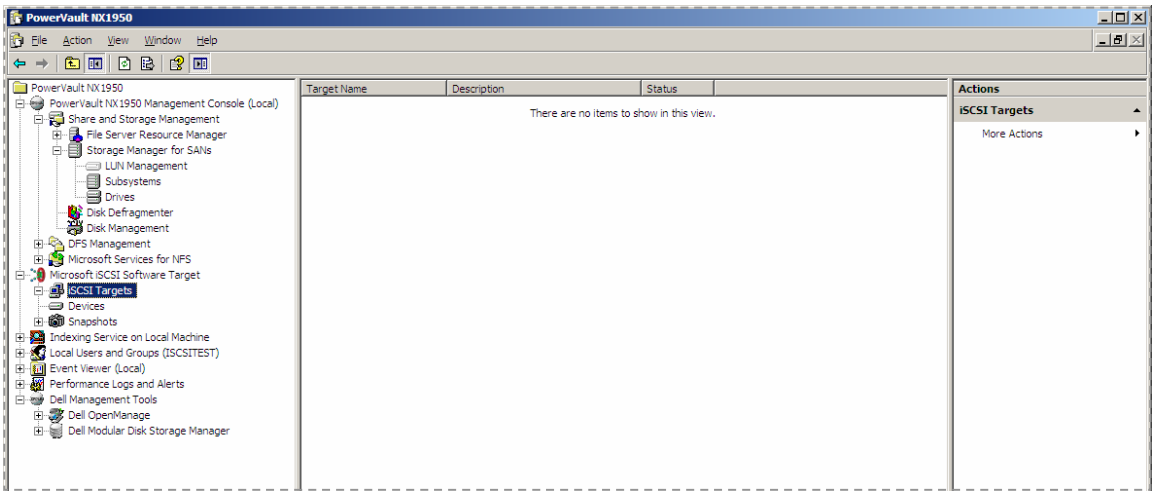

The iSCSI target wizard is launched. In this case, we will create two iSCSI targets. Each target will be made available to a different application on the host server. The target itself in the Microsoft-based iSCSI target solutions merely defines the path that the iSCSI storage traffic will use from the iSCSI initiator. The actual storage used by the target will be defined in a later step when the virtual disks are created.

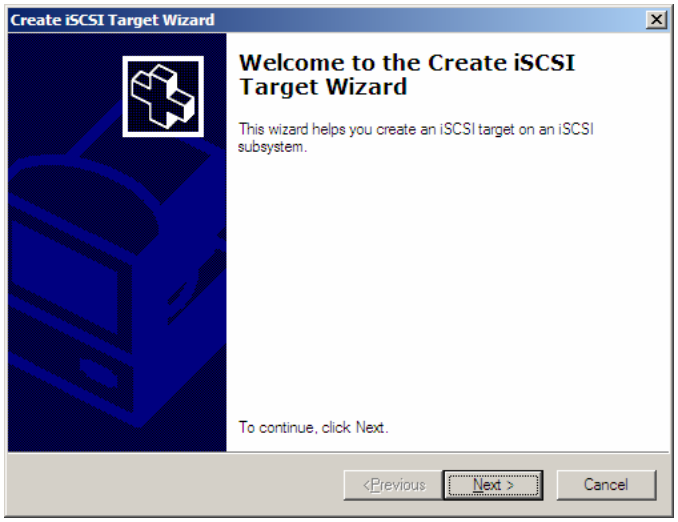

Each iSCSI target needs a name and optional description, which are supplied below.

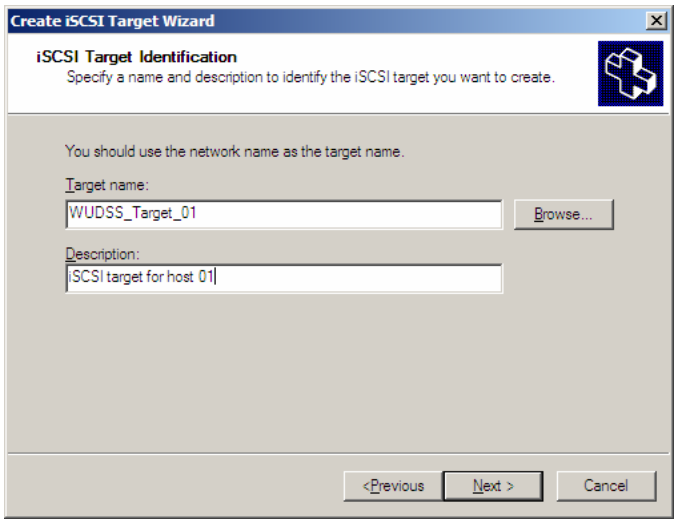

Each iSCSI target needs to be associated with an iSCSI initiator. The iSCSI initiator is the host that is requesting access to the storage represented by the iSCSI target name. The wizard prompts for the iSCSI Qualified Name (IQN) of the iSCSI initiator or provides alternate ways to identify the iSCSI initiator. In this case, we will choose to identify the iSCSI initiator by its IP address.

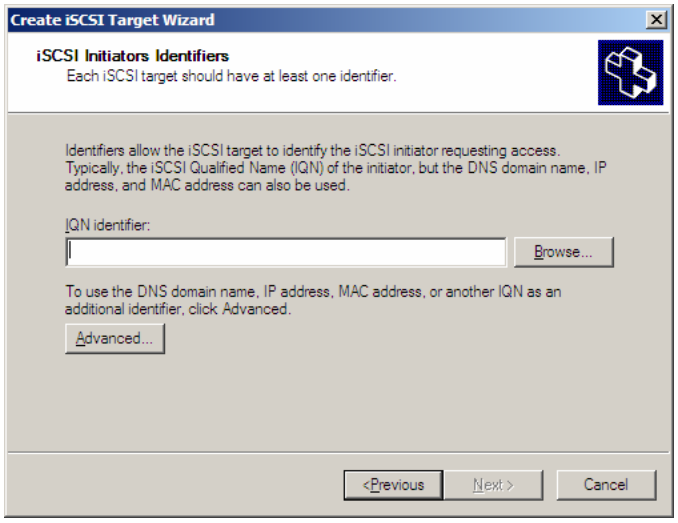

Clicking “Advanced” allows us to choose alternate methods of identification.

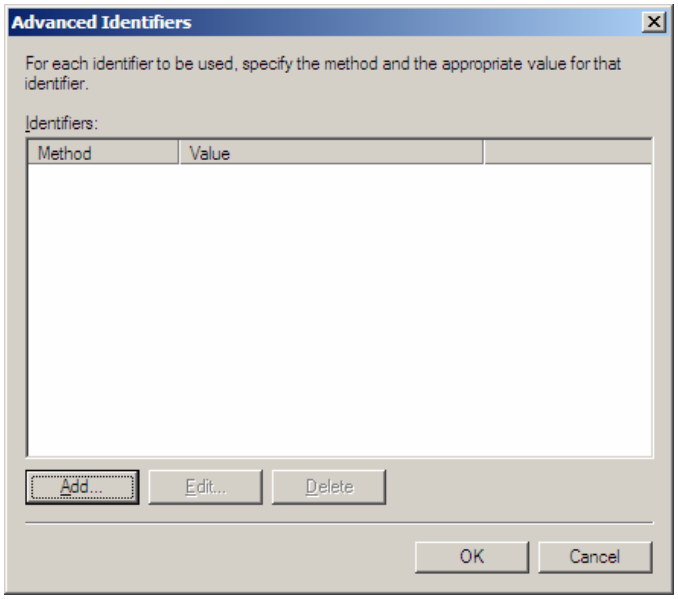

Clicking “Add” allows the type of identifier to be entered and the specific identifying information to be entered.

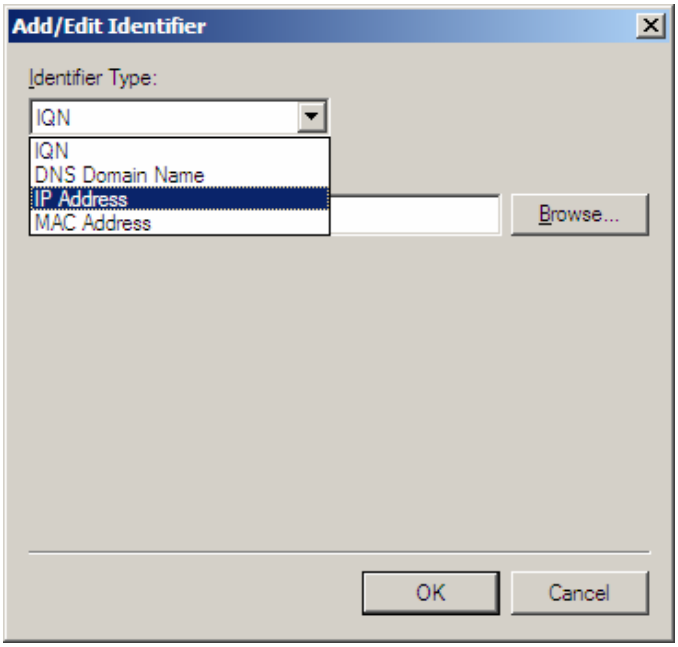

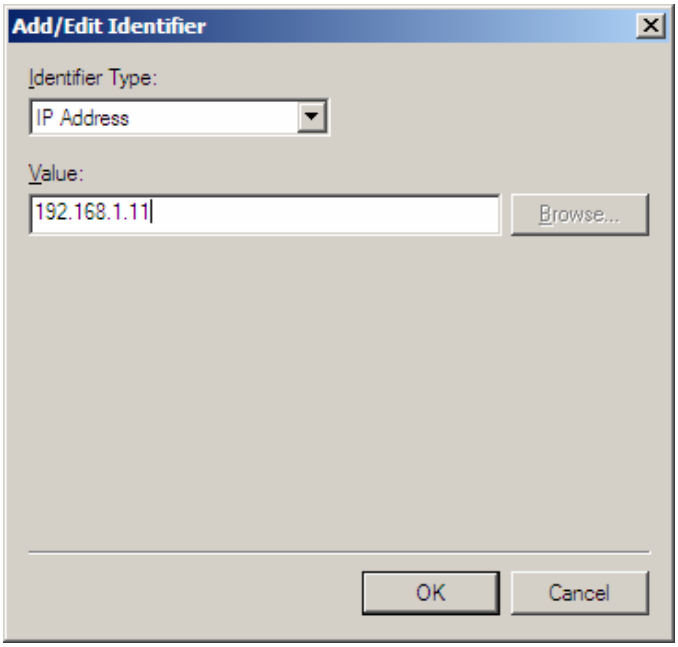

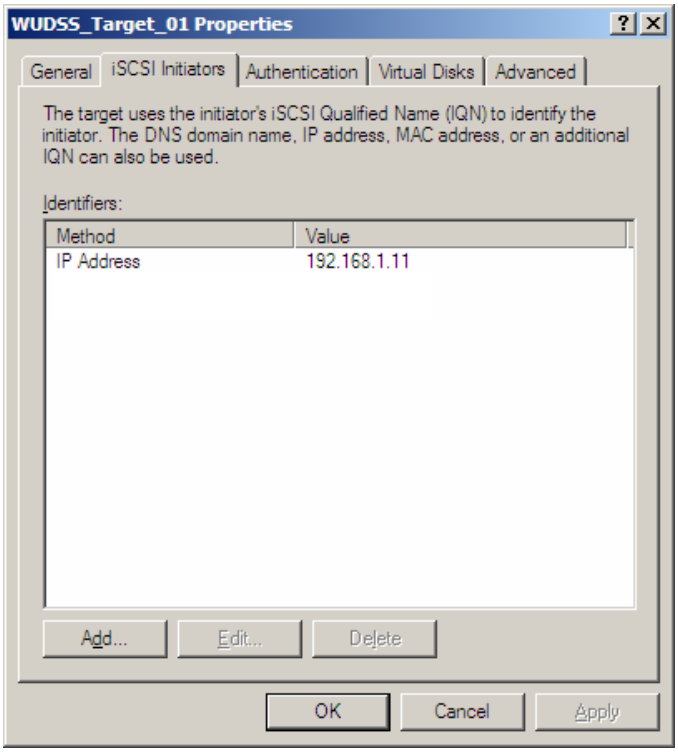

After identifying the iSCSI initiator, we are ready to proceed.

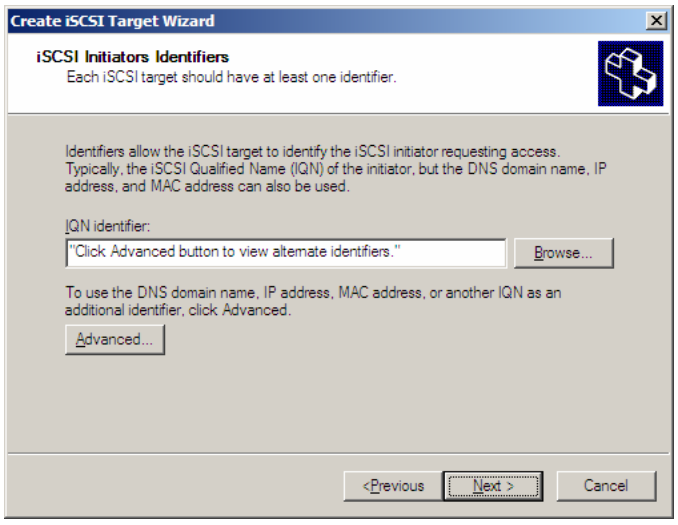

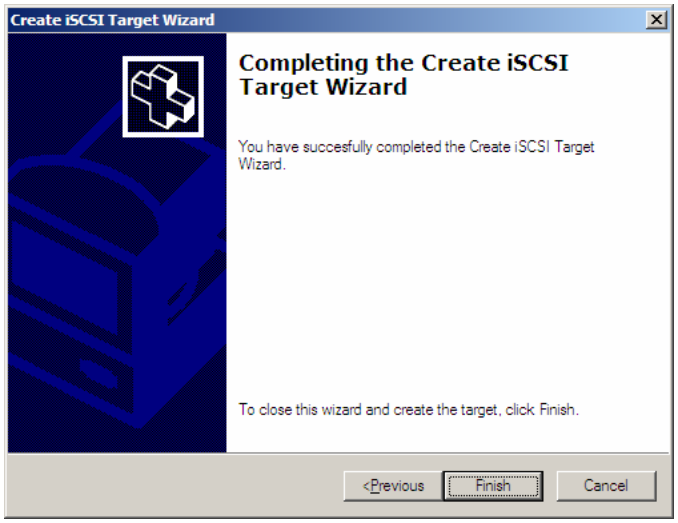

The management console now shows the newly created iSCSI target.

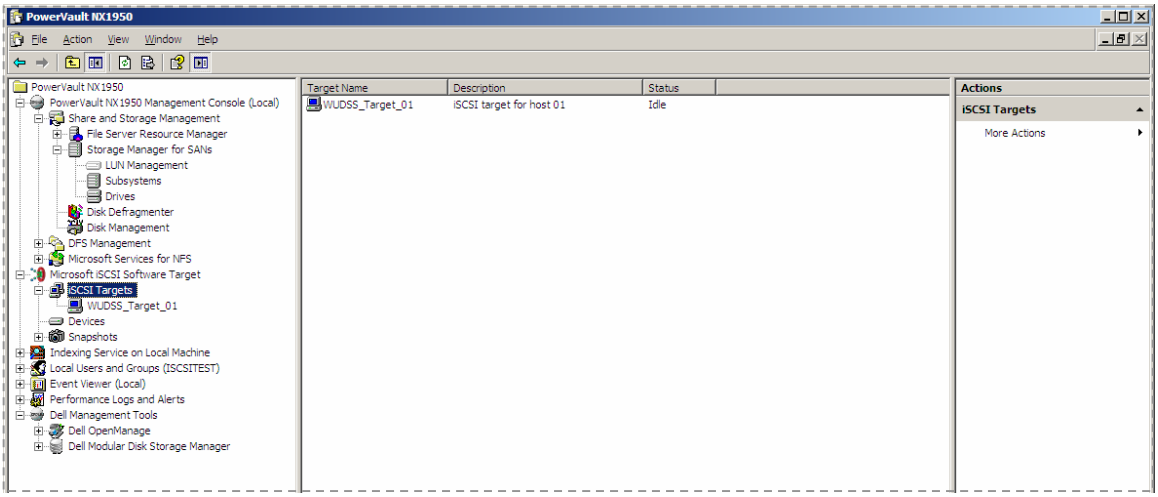

The management console also shows the underlying devices available for the iSCSI targets. The storage that will be used by the iSCSI initiators (application hosts) will be defined in a later step when the virtual disks are created.

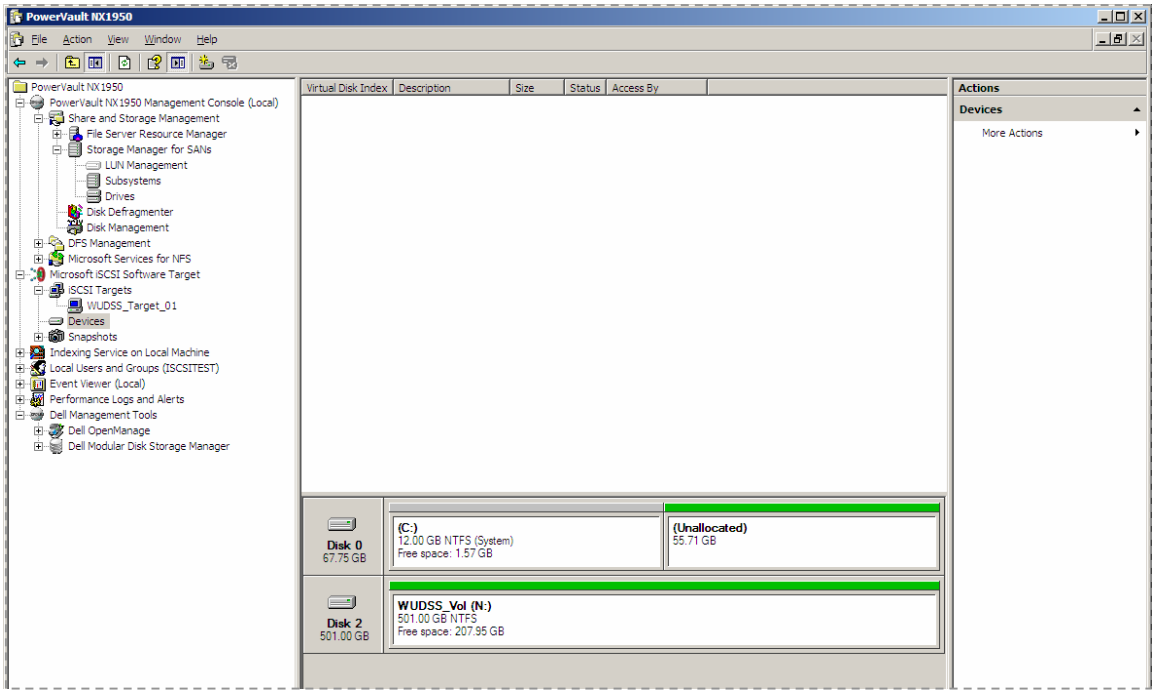

6. Create Multi-path I/O for iSCSI Targets (optional)

For a Microsoft-based target solution, multiple paths to the target device are created by providing multiple iSCSI initiator identifiers for the same target. Multiple paths to an iSCSI target can be creating by adding sessions to the iSCSI target (Microsoft MPIO) or adding additional connections to the iSCSI target (Multiple Connections per Session). In the example below, we use MS MPIO and add a second IP address that is to be associated with the same iSCSI initiator, so that there are two IP addresses that can access the target. These two addresses are associated with two Ethernet ports on the same host server.

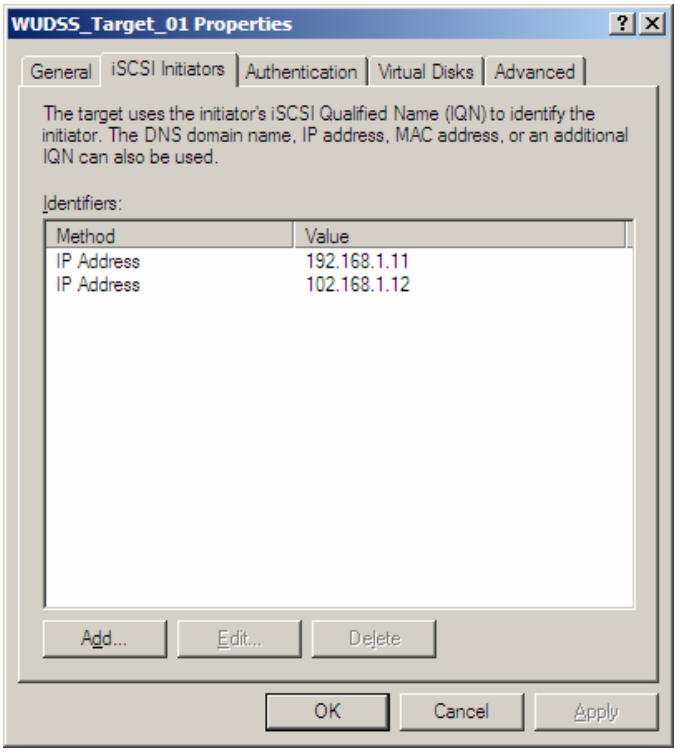

Other steps will need to be taken on the iSCSI initiator side to complete the multi-path configuration.

7. Configure Security for iSCSI Targets (optional)

CHAP can be configured with a password also known as the “CHAP secret”. The CHAP secret provides an additional level of security between the iSCSI initiator and target. The CHAP secret is provided on the “Authentication” tab of the target properties.

8. Make iSCSI Targets Ready for Use for iSCSI Initiators

Virtual disks need to be created on the iSCSI targets for Microsoft-based iSCSI target solutions. These virtual disks represent the storage volumes that the iSCSI initiators will use. The maximum capacity represented by all the virtual disks on a given iSCSI target on a Microsoft-based iSCSI target solution is two terabytes (2 TB) per target. In this example, we create a 100 GB and 200 GB virtual disk on the iSCSI target. These two virtual disks will be viewed as volumes by the iSCSI initiators over the TCP/IP network.

By right-clicking on the target name, the “Create Virtual Disk” wizard is launched.

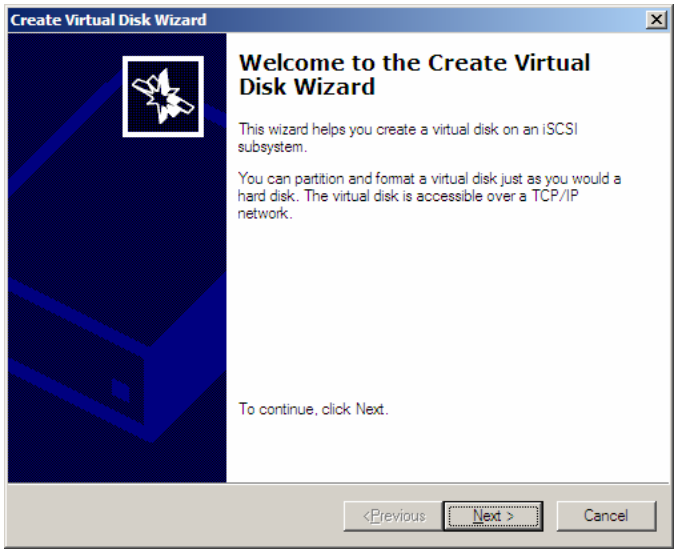

The virtual disk is created on the internal disk volume that is available to the iSCSI target. In this case, it is the “N” volume.

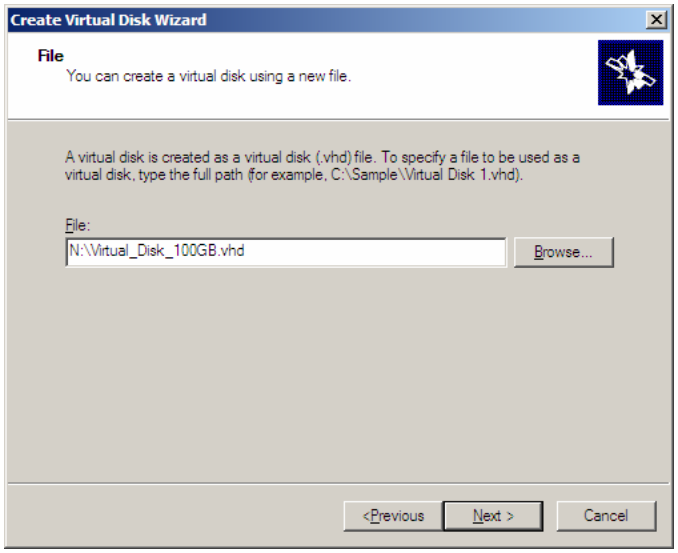

The size of the virtual disk depends on the needs of the application on the host server. For this virtual disk, we choose a size of 100 GB from the available 501 GB on this volume.

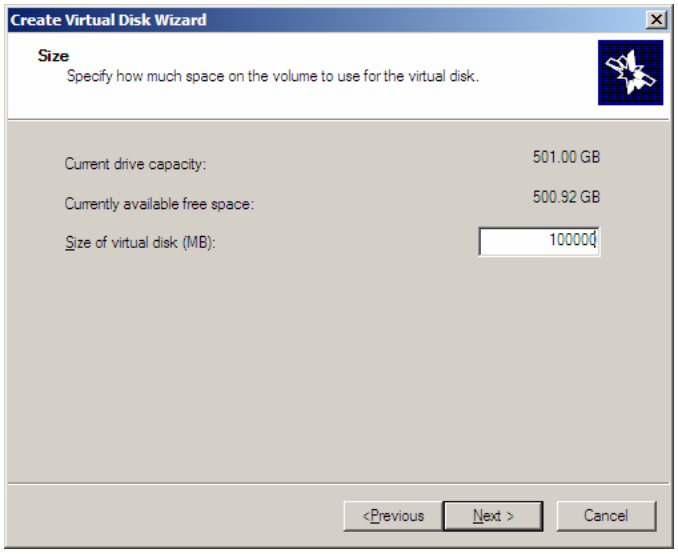

A description, although optional, is useful for better management.

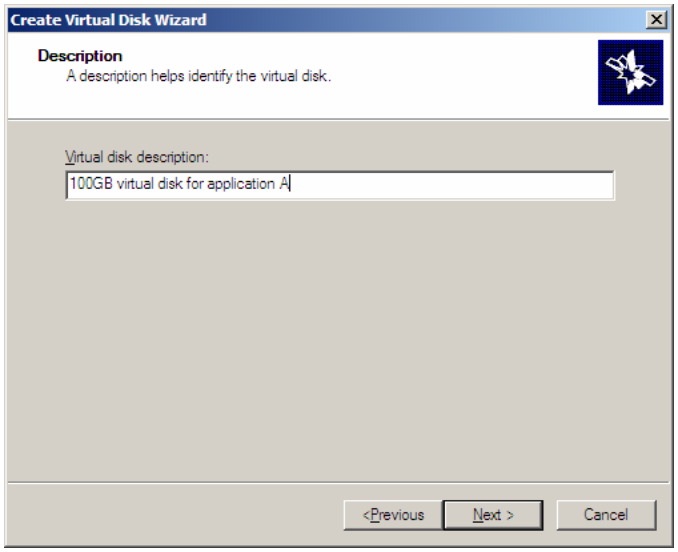

This virtual disk must be associated with an iSCSI target in order for the application host to use this storage as an iSCSI storage volume.

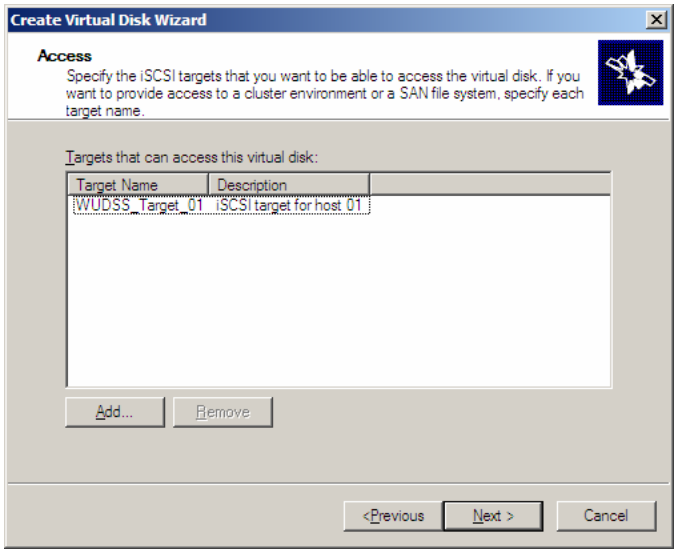

This completes the virtual disk configuration.

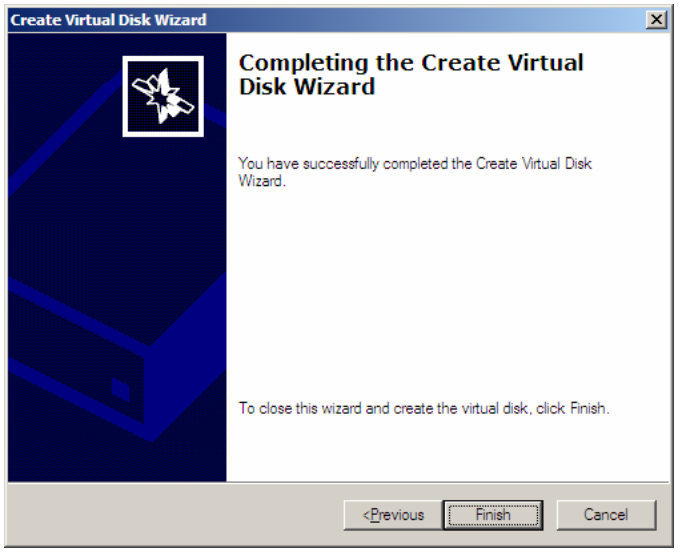

We repeat this process to create a second virtual disk of size 200 GB. After configuring the virtual disks, the management console shows the virtual disks associated with the iSCSI target.

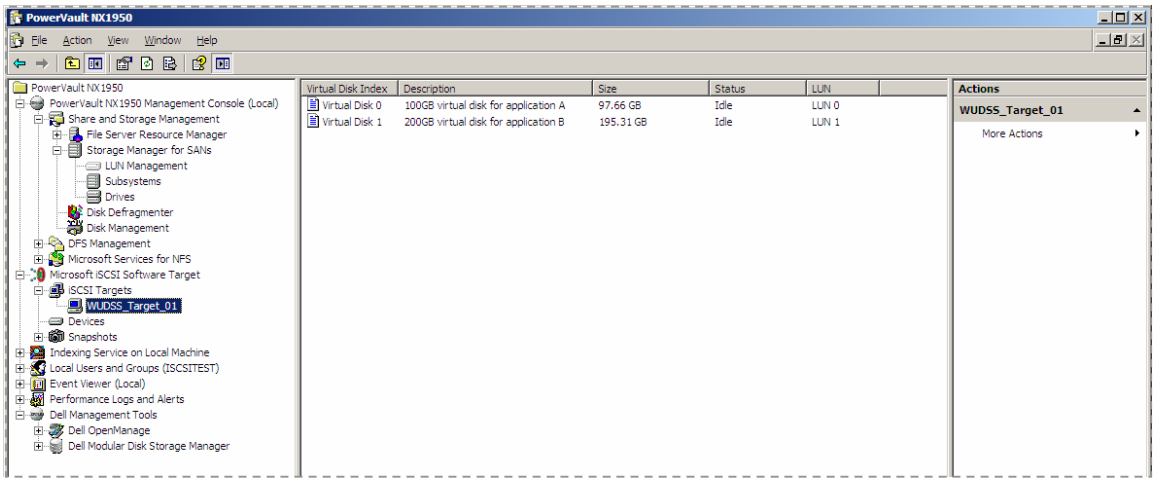

The iSCSI target device view shows the total volume size and the free space remaining on the device that is available for iSCSI targets.

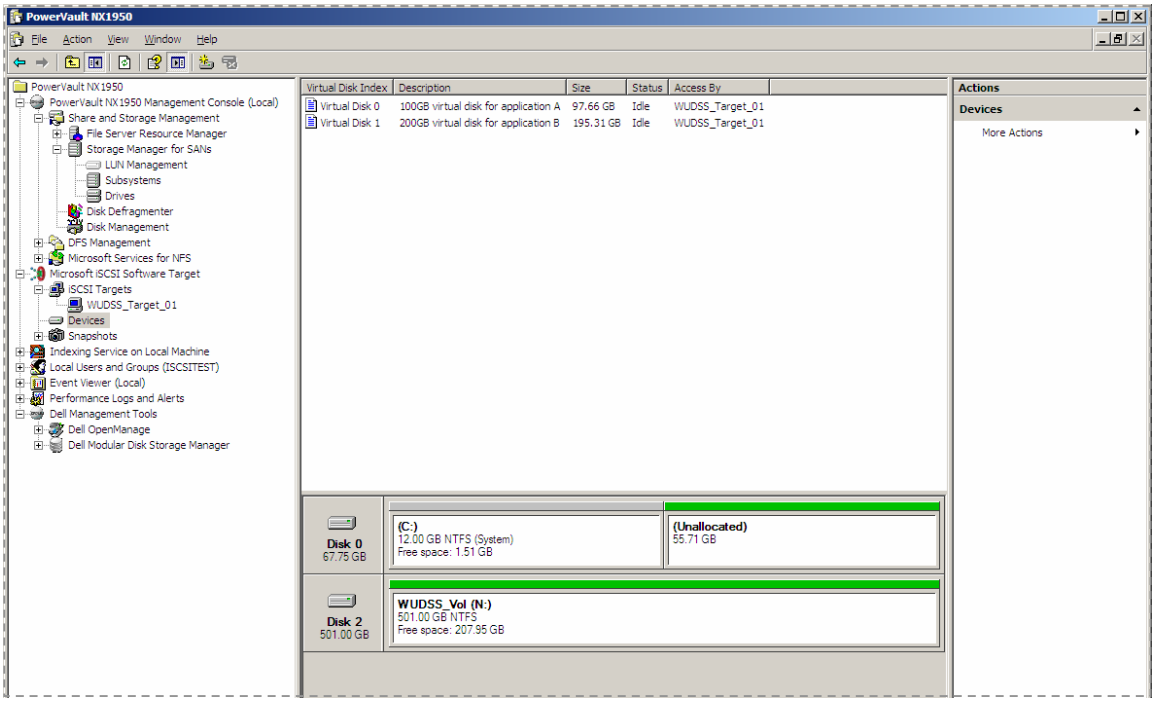

The target side configuration is now complete.

Initiator Configuration Steps

Configure Multi-path I/O from Application Host

To configure multi-path iSCSI I/O for the initiator that uses the Dell NX1950 iSCSI targets, follow the directions for Microsoft Multi-path I/O from the Deployment section of this document above.

Basic Performance Results

The following performance data is not intended to be viewed as a comprehensive performance benchmark, but to provide a general performance overview for the Dell PowerVault NX1950 Networked Storage Solution.

Selected performance results are shown below, using a standard server-class network adapter, without receive-side scaling on the host. This configuration used two paths from one host, two I/O workers, simultaneously accessing two target volumes and a queue depth of 20. Each volume shared round-robin access across both paths.

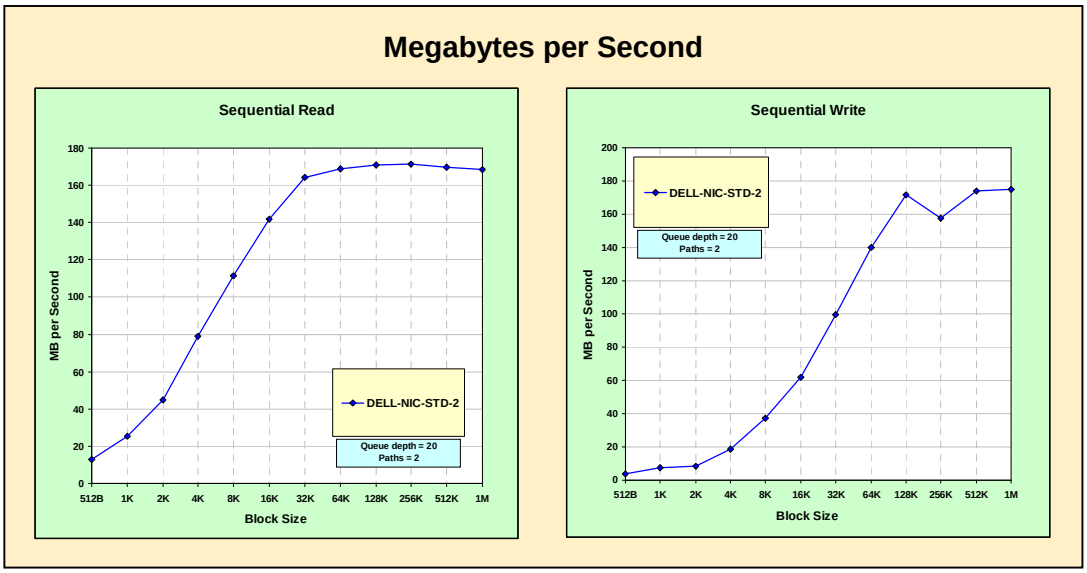

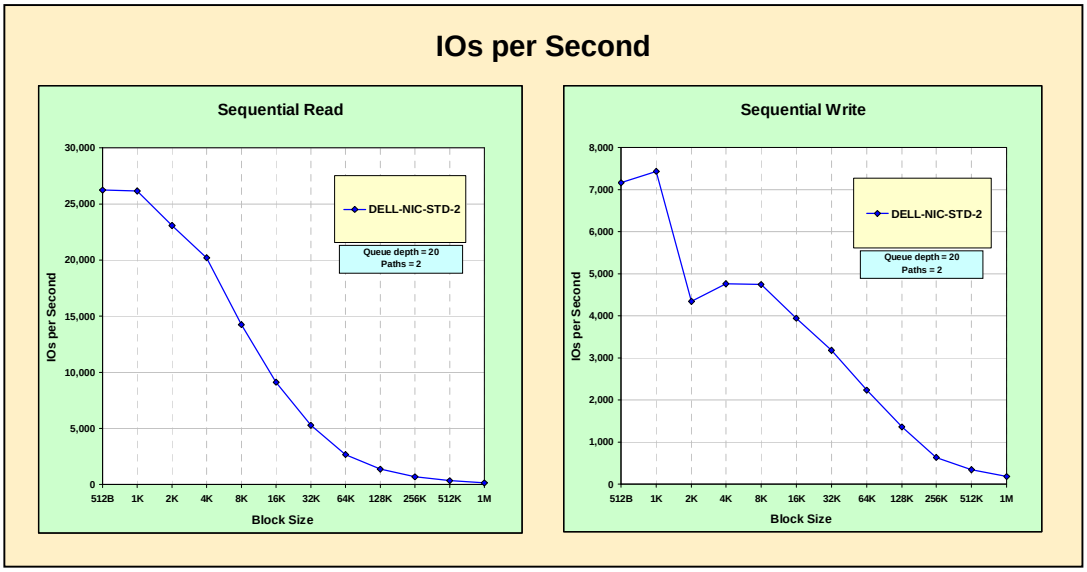

EqualLogic® PS3800XV

The EqualLogic PS3800XV is an iSCSI target solution that includes a RISC-based, dual-controller disk subsystem with 15K RPM, 146GB SAS disk drives totaling 2.3 TB of raw capacity per module. It includes battery-backed and mirrored RAID cache, and its storage can be configured for RAID 5, RAID 10 or RAID 50.

EqualLogic's iSCSI SAN solutions encompass advanced automated and intelligent management features, which continuously load balance and optimize the SAN to meet applications performance and capacity needs. In addition to the fault tolerant and fully redundant design of the PS Series, they also include an extensive set of data protection features such as snapshots, remote replication and fast fail back, providing the tools needed to setup a robust yet simple Data Protection and Disaster Recovery strategy. The PS Series offer multiple models with different drive capacities and speeds allowing for an efficient set up of different pools of storage to enable appropriate service levels for individual applications. EqualLogic has early on joined Microsoft Simple SAN program to provide, affordable, enterprise class, SAN solutions that are easy to setup manage and scale.

Target Configuration Steps

To install the EqualLogic PS3800XV solution from factory settings, a computer must be connected via the supplied serial cable to the PS3800XV array. The Quick Start Guide steps the administrator through the process of connecting all the cables properly and running the Group Manager Setup Utility to perform the basic system configuration.

EqualLogic also provides a Host Integration Tools CDROM that provides the VDS and MPIO drivers and the auto-snapshot manager to use with the EqualLogic array. These were installed from the CD on the iSCSI initiator hosts. If the EqualLogic Multi-path Device Specific Module (DSM) is installed, a system reboot will be required. The multi-path DSM is required for multi-path operations.

1. Configure Network Settings for iSCSI Target Device

The Group Manager Setup Utility, run from a computer connected via serial cable or Ethernet connection to the PS3800XV array, asks the administrator to specify the IP address and related network settings for the primary Ethernet connection and the name of the logical group and its IP address. The utility program HyperTerminal was used to communicate with the array.

```
Enter the network configuration for the array:

 Member name []: Demartek
 Network interface [eth0]:
 IP address for network interface []: 192.168.0.231
 Netmask [255.255.255.0]:
 Default gateway [192.17.2.1]: 192.168.0.1

Enter the name and IP address of the group that the array will join.

 Group name []: Demartek
 Group IP address []: 192.168.0.230
```

2. Launch Management Console

The management console is first launched via a 32-bit web browser. After logout, an option is displayed that provides a stand-alone JAVA application to use for the management console.

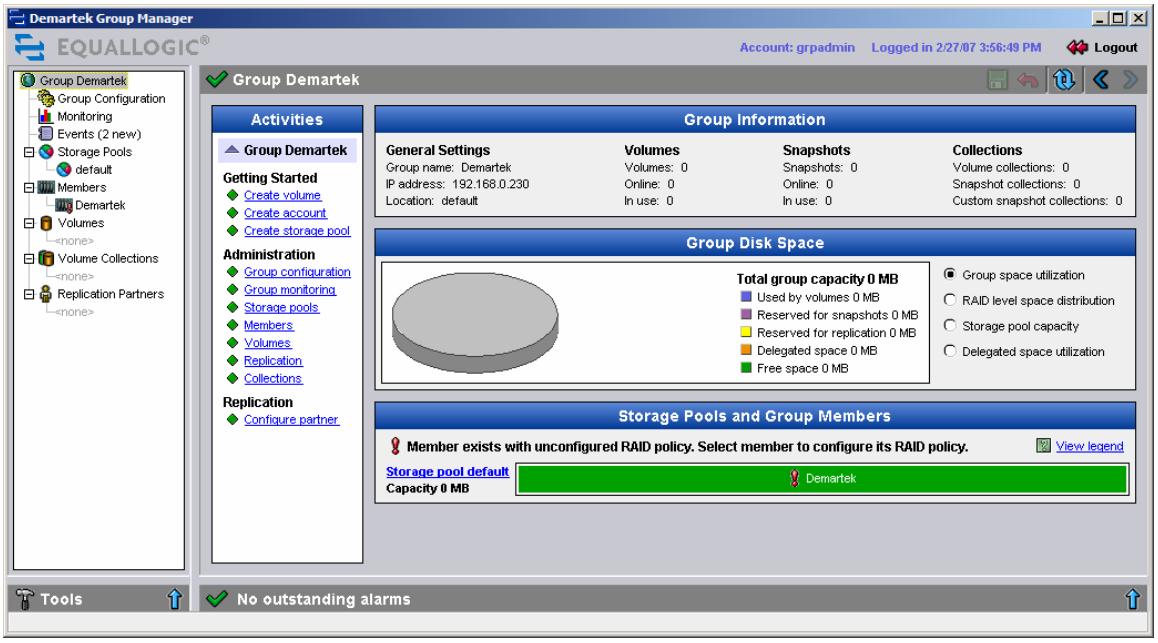

3. Create Volumes on Disk Array

Before Volumes can be created on the array, the RAID policy must be set. The Quick Start Guide provides the step-by-step instructions to set the RAID policy.

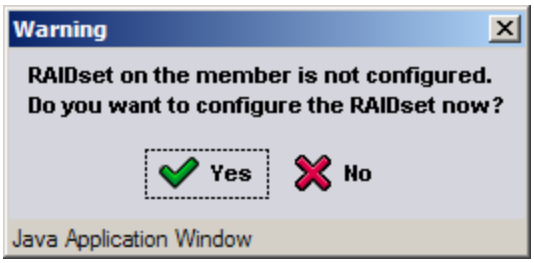
Warning
RAIDset on the member is not configured.
Do you want to configure the RAIDset now?
Yes
No
Java Application Window

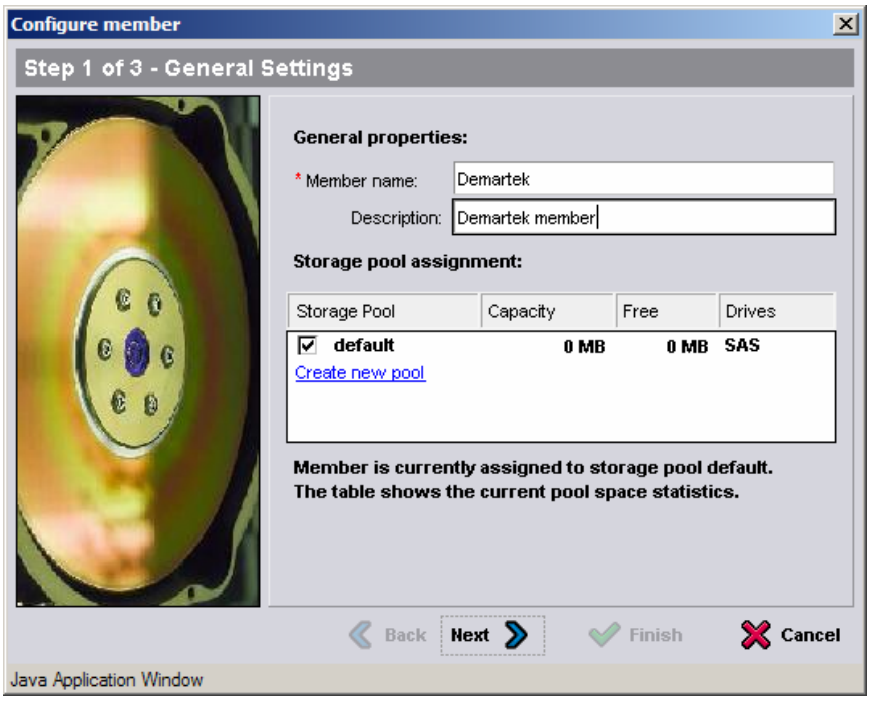
Configure member
Step 1 of 3 - General Settings
General properties:
* Member name:
Demartek
Description:
Demartek member
Storage pool assignment:
Storage Pool
Capacity
Free
Drives
default
0 MB
0 MB
SAS
Create new pool
Member is currently assigned to storage pool default.
The table shows the current pool space statistics.
Back
Next
Finish
Cancel
Java Application Window

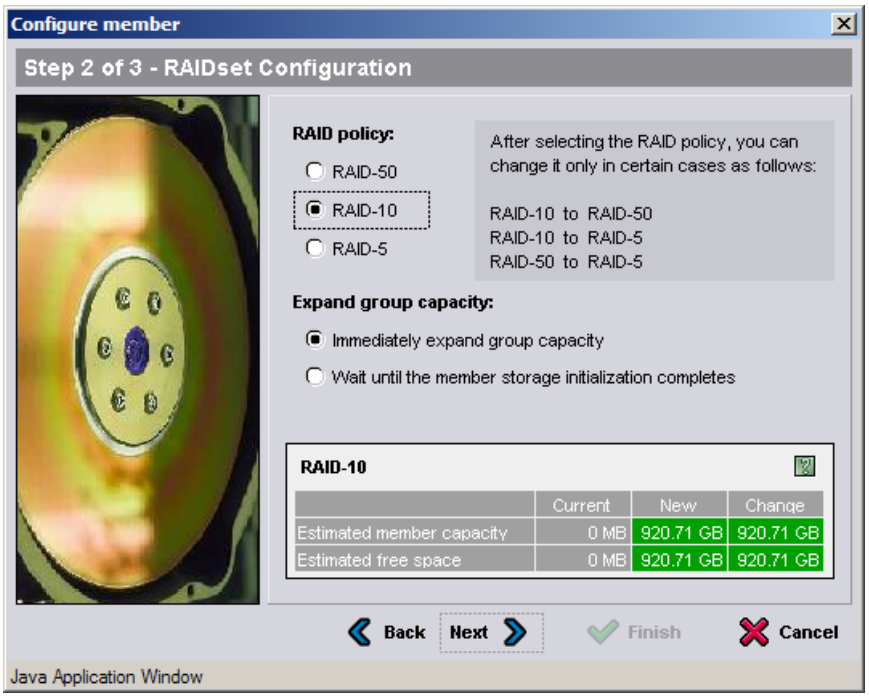
Configure member
Step 2 of 3 - RAIDset Configuration
RAID policy:
RAID-50
RAID-10
RAID-5
After selecting the RAID policy, you can change it only in certain cases as follows:
RAID-10 to RAID-50
RAID-10 to RAID-5
RAID-50 to RAID-5
Expand group capacity:
Immediately expand group capacity
Wait until the member storage initialization completes
RAID-10
Current
New
Change
Estimated member capacity
0 MB
920.71 GB
920.71 GB
Estimated free space
0 MB
920.71 GB
920.71 GB
Back
Next
Finish
Cancel
Java Application Window

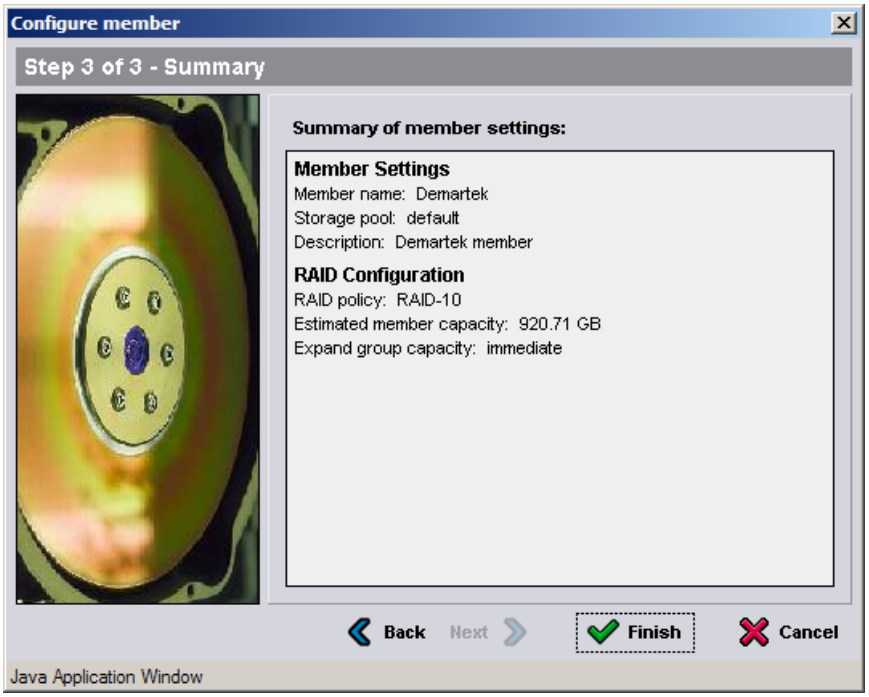
Configure member
Step 3 of 3 - Summary
Summary of member settings:
Member Settings
Member name: Demartek
Storage pool: default
Description: Demartek member
RAID Configuration
RAID policy: RAID-10
Estimated member capacity: 920.71 GB
Expand group capacity: immediate
Back
Next
Finish
Cancel
Java Application Window

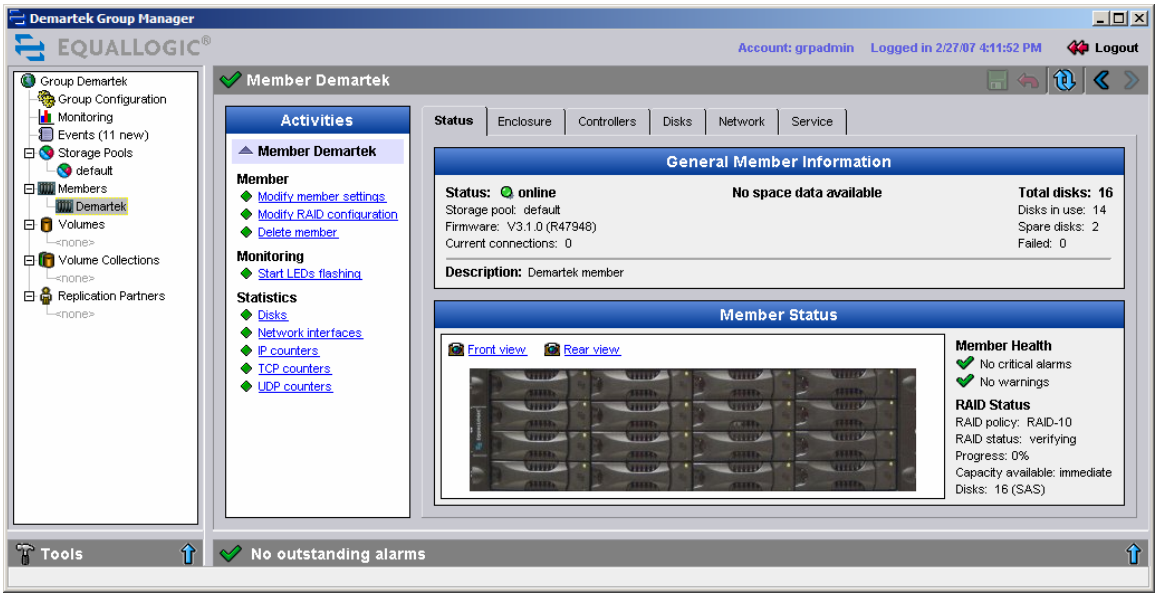

The array began to apply the RAID policy of RAID-10 across all disks.

Volumes are created using the “Activities” panel of the management console.

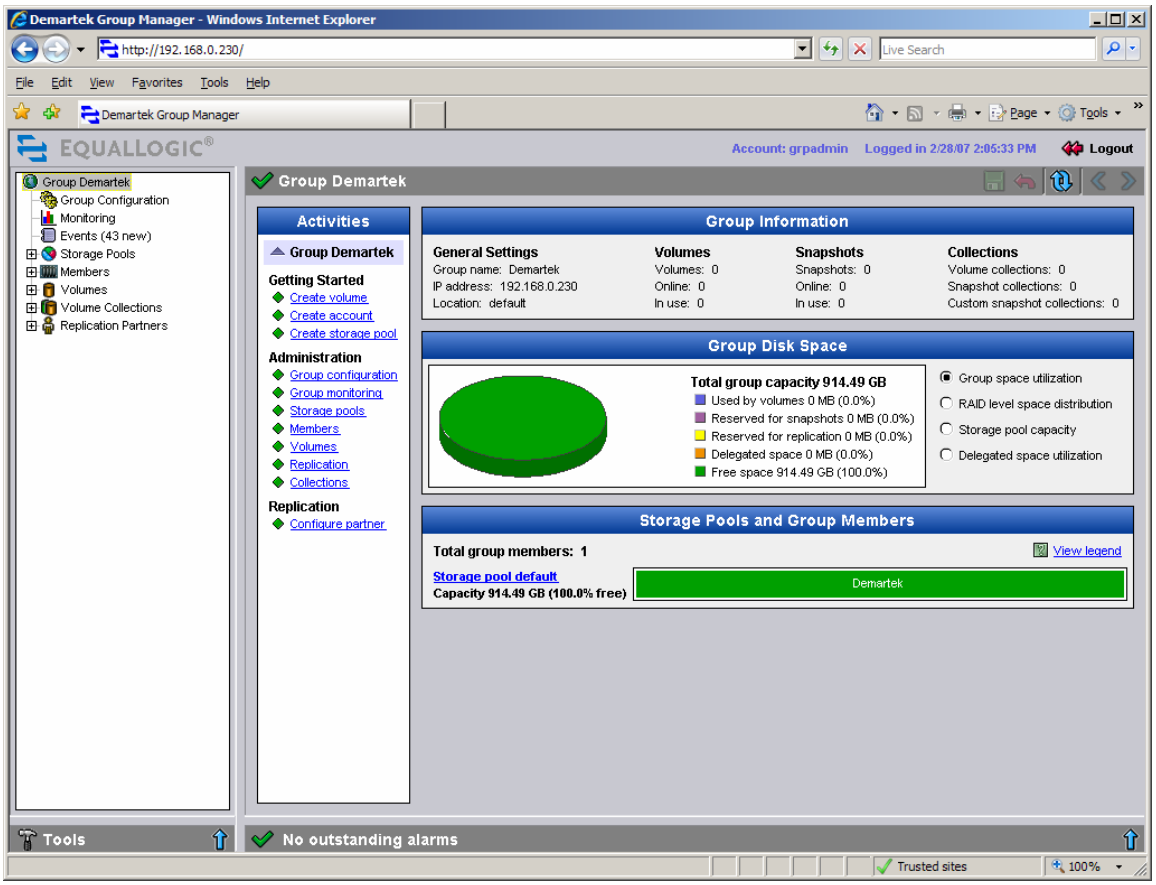

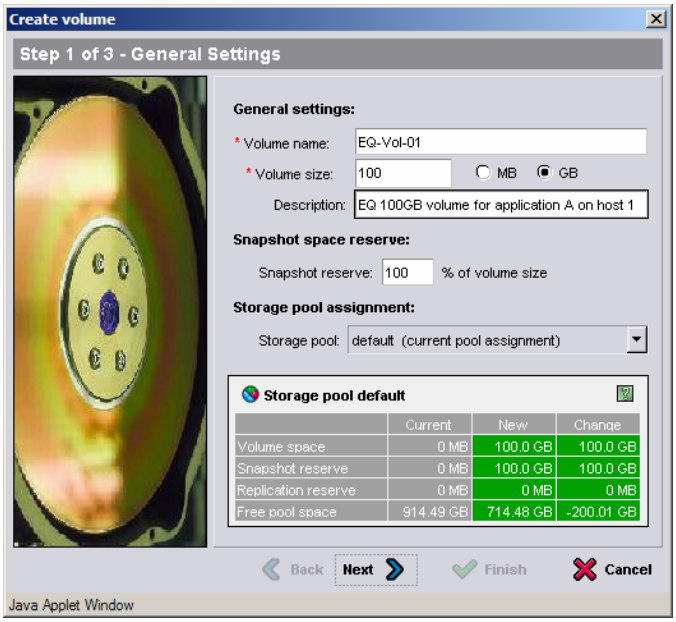
Create volume
Step 1 of 3 - General Settings
General settings:
* Volume name: EQ-Vol-01
* Volume size: 100 MB GB
Description: EQ 100GB volume for application A on host 1
Snapshot space reserve:
Snapshot reserve: 100 % of volume size
Storage pool assignment:
Storage pool: default (current pool assignment)
Storage pool default
Current New Change
Volume space 0 MB 100.0 GB 100.0 GB
Snapshot reserve 0 MB 100.0 GB 100.0 GB
Replication reserve 0 MB 0 MB 0 MB
Free pool space 914.49 GB 714.48 GB -200.01 GB
Back Next Finish Cancel
Java Applet Window

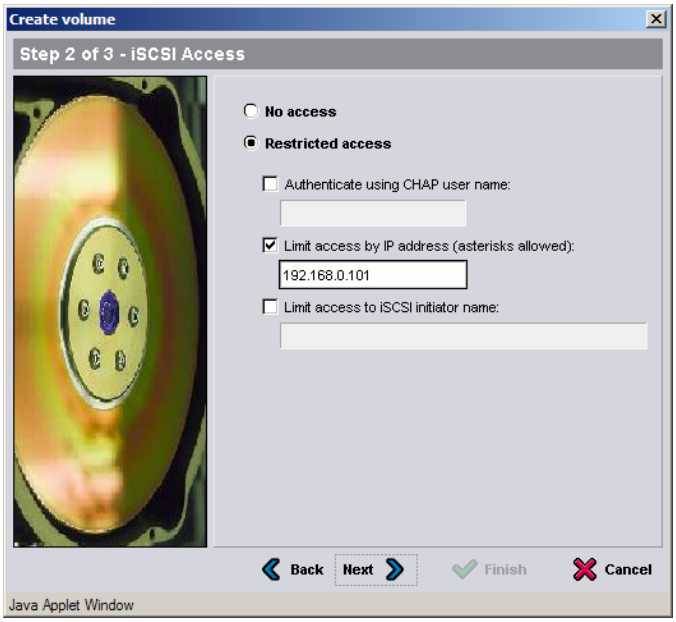
Create volume
Step 2 of 3 - iSCSI Access
No access
Restricted access
Authenticate using CHAP user name:
Limit access by IP address (asterisks allowed):
192.168.0.101
Limit access to iSCSI initiator name:
Back Next Finish Cancel
Java Applet Window

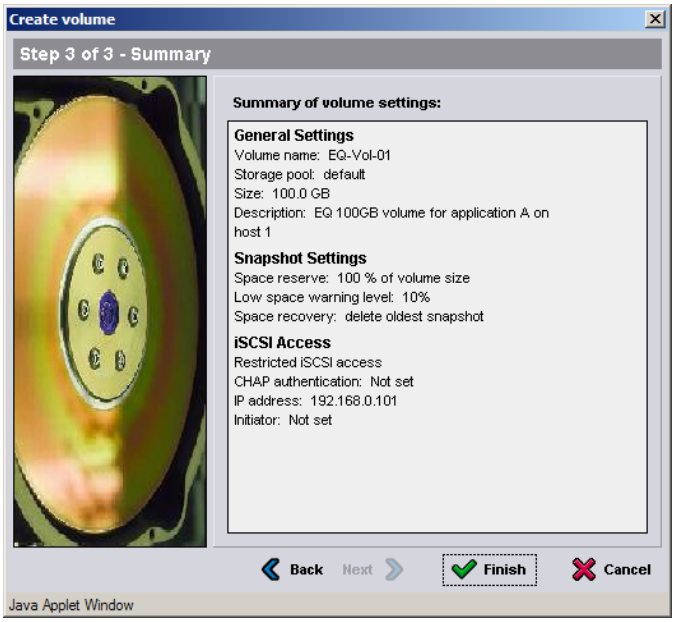

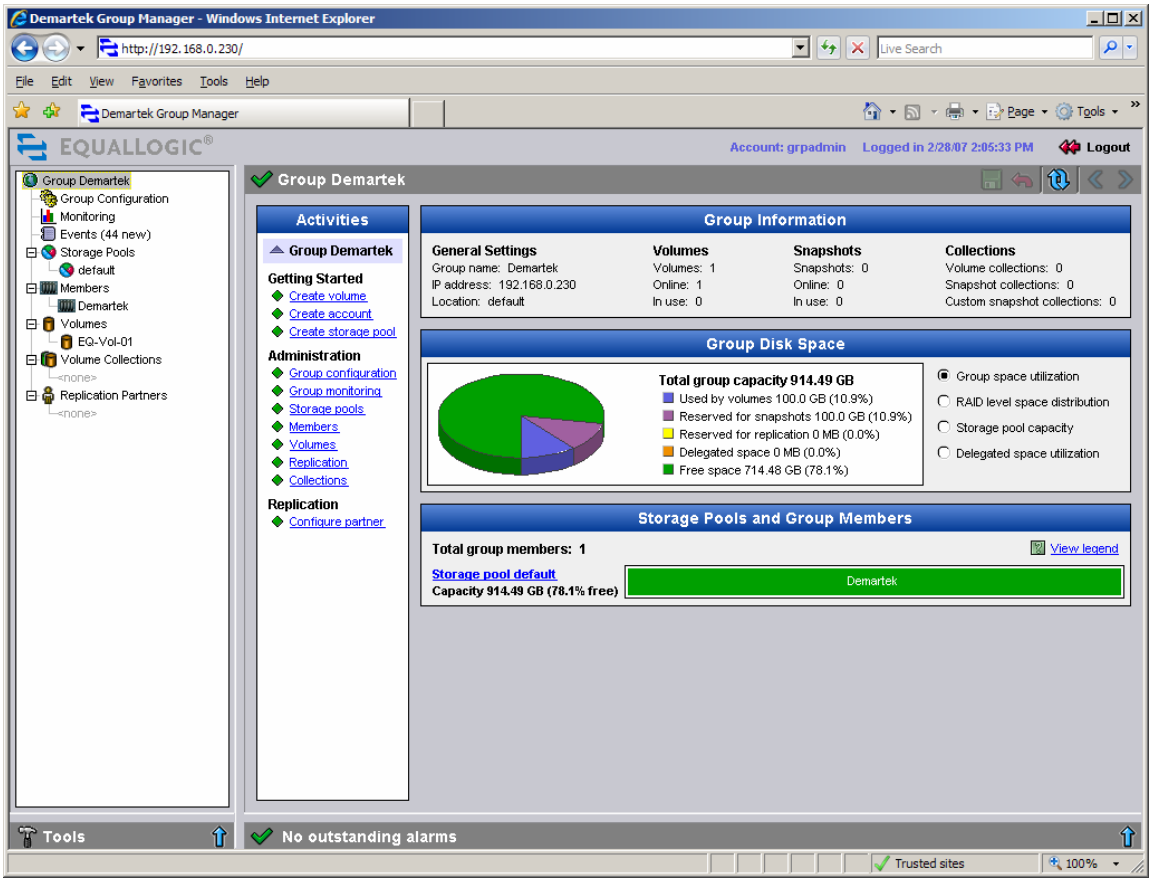

Four volumes were created.

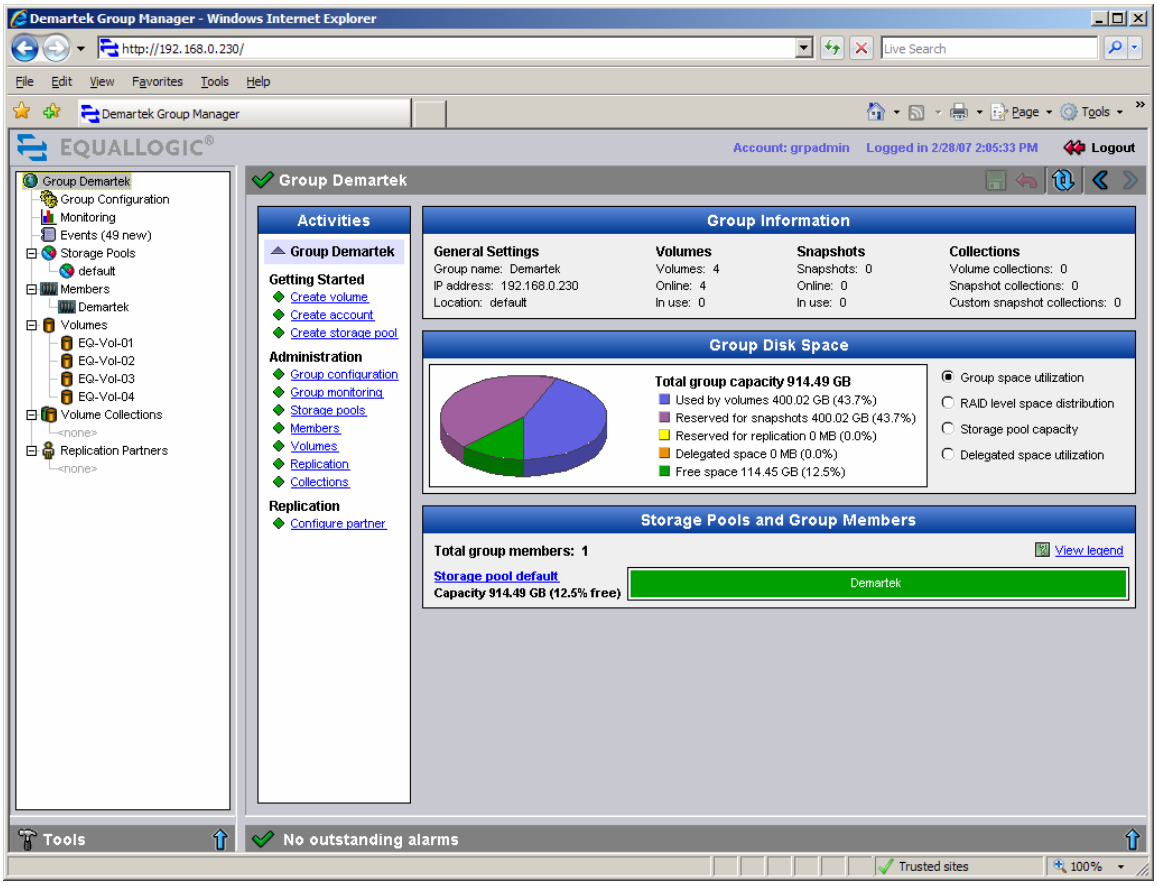
Demartek Group Manager - Windows Internet Explorer
http://192.168.0.230/
Live Search
File Edit View Favorites Tools Help
Demartek Group Manager
Page Tools
EQUALLOGIC®
Account: grpadmin Logged in 2/28/07 2:05:33 PM
Logout
Group Demartek
Group Configuration
Monitoring
Events (49 new)
Storage Pools
default
Members
Demartek
Volumes
EQ-Vol-01
EQ-Vol-02
EQ-Vol-03
EQ-Vol-04
Volume Collections
<none>
Replication Partners
<none>
Group Demartek
Activities
Group Demartek
Getting Started
Create volume
Create account
Create storage pool
Administration
Group configuration
Group monitoring
Storage pools
Members
Volumes
Replication
Collections
Replication
Configure partner
Group Information
General Settings
Group name: Demartek
IP address: 192.168.0.230
Location: default
Volumes
Volumes: 4
Online: 4
In use: 0
Snapshots
Snapshots: 0
Online: 0
In use: 0
Collections
Volume collections: 0
Snapshot collections: 0
Custom snapshot collections: 0
Group Disk Space
Total group capacity 914.49 GB
Used by volumes 400.02 GB (43.7%)
Reserved for snapshots 400.02 GB (43.7%)
Reserved for replication 0 MB (0.0%)
Delegated space 0 MB (0.0%)
Free space 114.45 GB (12.5%)
Group space utilization
RAID level space distribution
Storage pool capacity
Delegated space utilization
Storage Pools and Group Members
Total group members: 1
View legend
Storage pool default
Capacity 914.49 GB (12.5% free)
Demartek
Tools
No outstanding alarms
Trusted sites
100%

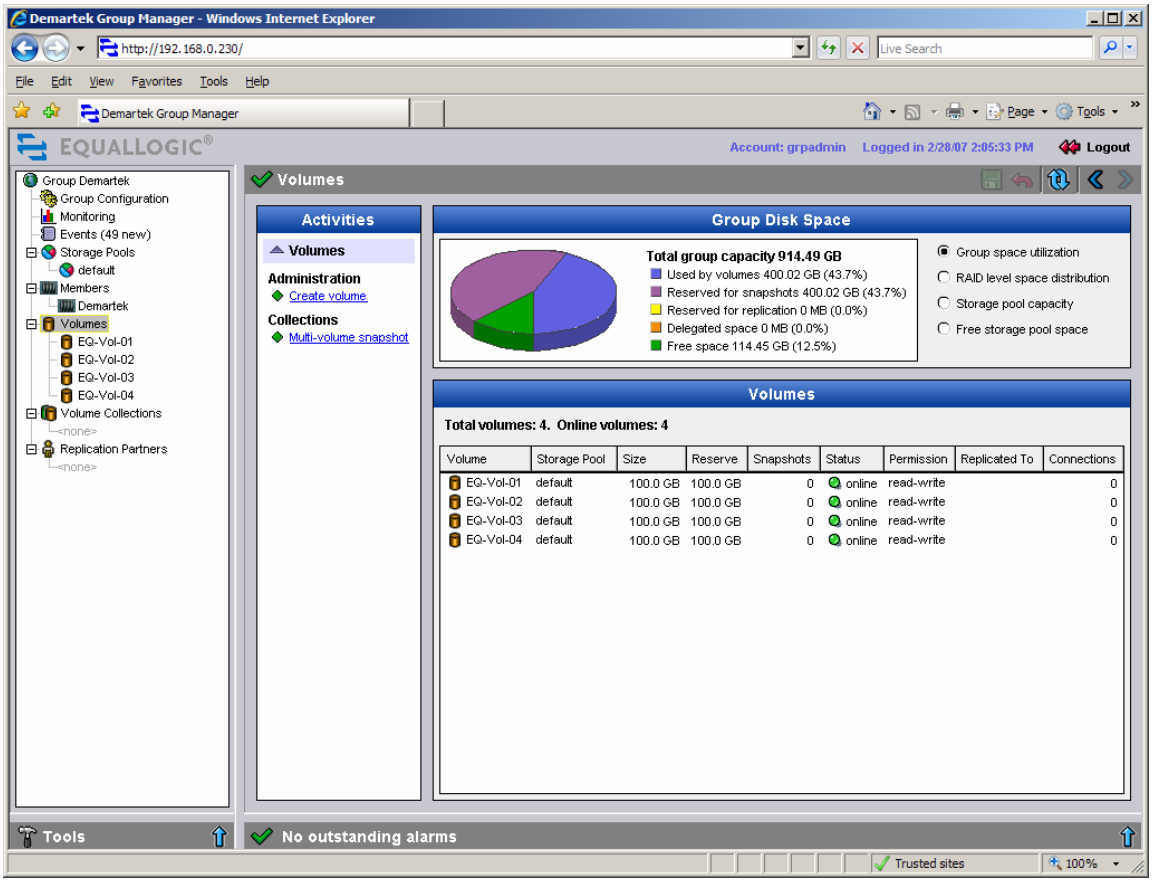
Demartek Group Manager - Windows Internet Explorer
http://192.168.0.230/
Live Search
File Edit View Favorites Tools Help
Demartek Group Manager
Page Tools
EQUALLOGIC®
Account: grpadmin Logged in 2/28/07 2:05:33 PM
Logout
Group Demartek
Group Configuration
Monitoring
Events (49 new)
Storage Pools
default
Members
Demartek
Volumes
EQ-Vol-01
EQ-Vol-02
EQ-Vol-03
EQ-Vol-04
Volume Collections
<none>
Replication Partners
<none>
Volumes
Activities
Volumes
Administration
Create volume
Collections
Multi-volume snapshot
Group Disk Space
Total group capacity 914.49 GB
Used by volumes 400.02 GB (43.7%)
Reserved for snapshots 400.02 GB (43.7%)
Reserved for replication 0 MB (0.0%)
Delegated space 0 MB (0.0%)
Free space 114.45 GB (12.5%)
Group space utilization
RAID level space distribution
Storage pool capacity
Free storage pool space
Volumes
Total volumes: 4. Online volumes: 4
Volume Storage Pool Size Reserve Snapshots Status Permission Replicated To Connections
EQ-Vol-01 default 100.0 GB 100.0 GB 0 online read-write 0
EQ-Vol-02 default 100.0 GB 100.0 GB 0 online read-write 0
EQ-Vol-03 default 100.0 GB 100.0 GB 0 online read-write 0
EQ-Vol-04 default 100.0 GB 100.0 GB 0 online read-write 0
Tools
No outstanding alarms
Trusted sites
100%

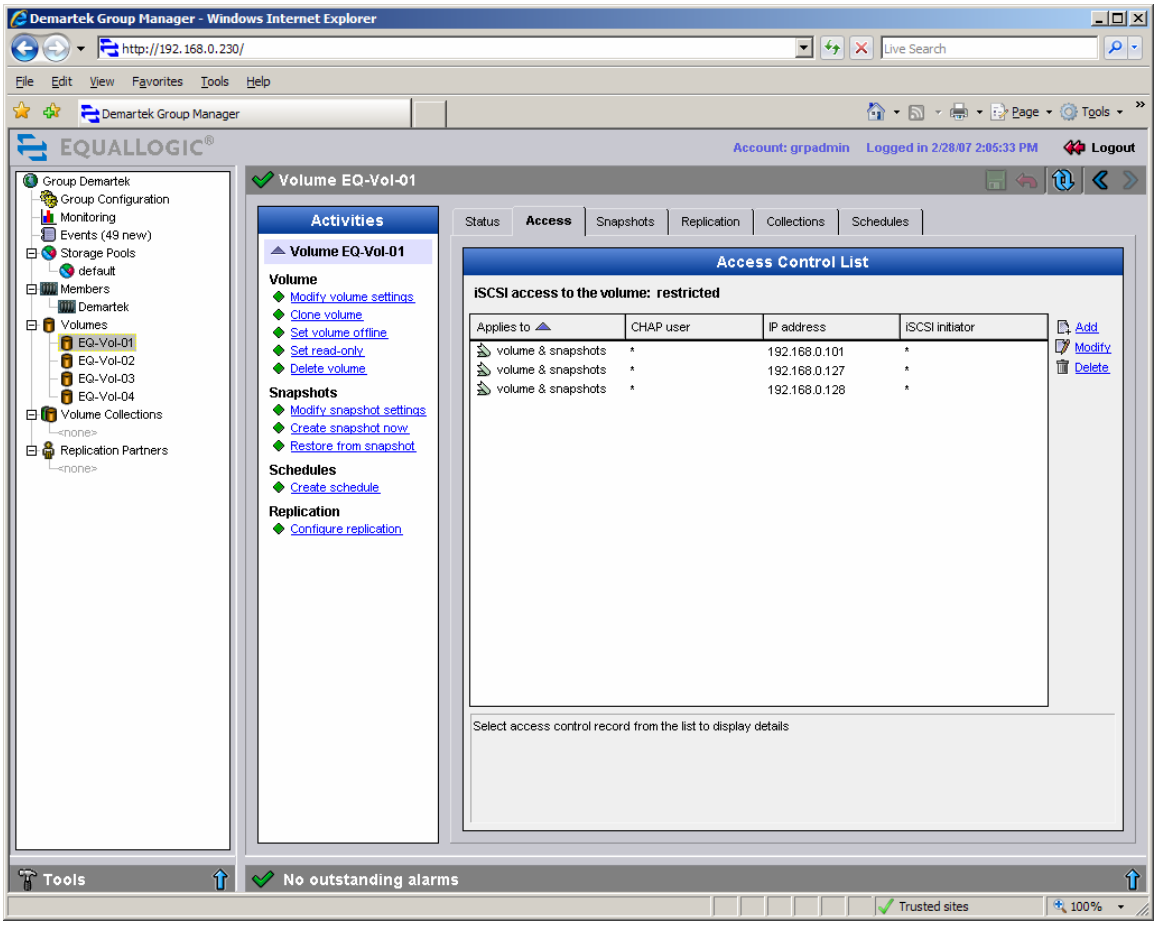
Demartek Group Manager - Windows Internet Explorer
http://192.168.0.230/
Live Search
File Edit View Favorites Tools Help
Demartek Group Manager
Page Tools
EQUALLOGIC®
Account: grpadmin Logged in 2/28/07 2:05:33 PM
Logout
Group Demartek
Group Configuration
Monitoring
Events (49 new)
Storage Pools
default
Members
Demartek
Volumes
EQ-Vol-01
EQ-Vol-02
EQ-Vol-03
EQ-Vol-04
Volume Collections
<none>
Replication Partners
<none>
Volume EQ-Vol-01
Activities
Volume EQ-Vol-01
Volume
Modify volume settings
Clone volume
Set volume offline
Set read-only
Delete volume
Snapshots
Modify snapshot settings
Create snapshot now
Restore from snapshot
Schedules
Create schedule
Replication
Configure replication
Status Access Snapshots Replication Collections Schedules
Access Control List
iSCSI access to the volume: restricted
Applies to CHAP user IP address iSCSI initiator
volume & snapshots * 192.168.0.101 *
volume & snapshots * 192.168.0.127 *
volume & snapshots * 192.168.0.128 *
Add
Modify
Delete
Select access control record from the list to display details
Tools
No outstanding alarms
Trusted sites
100%

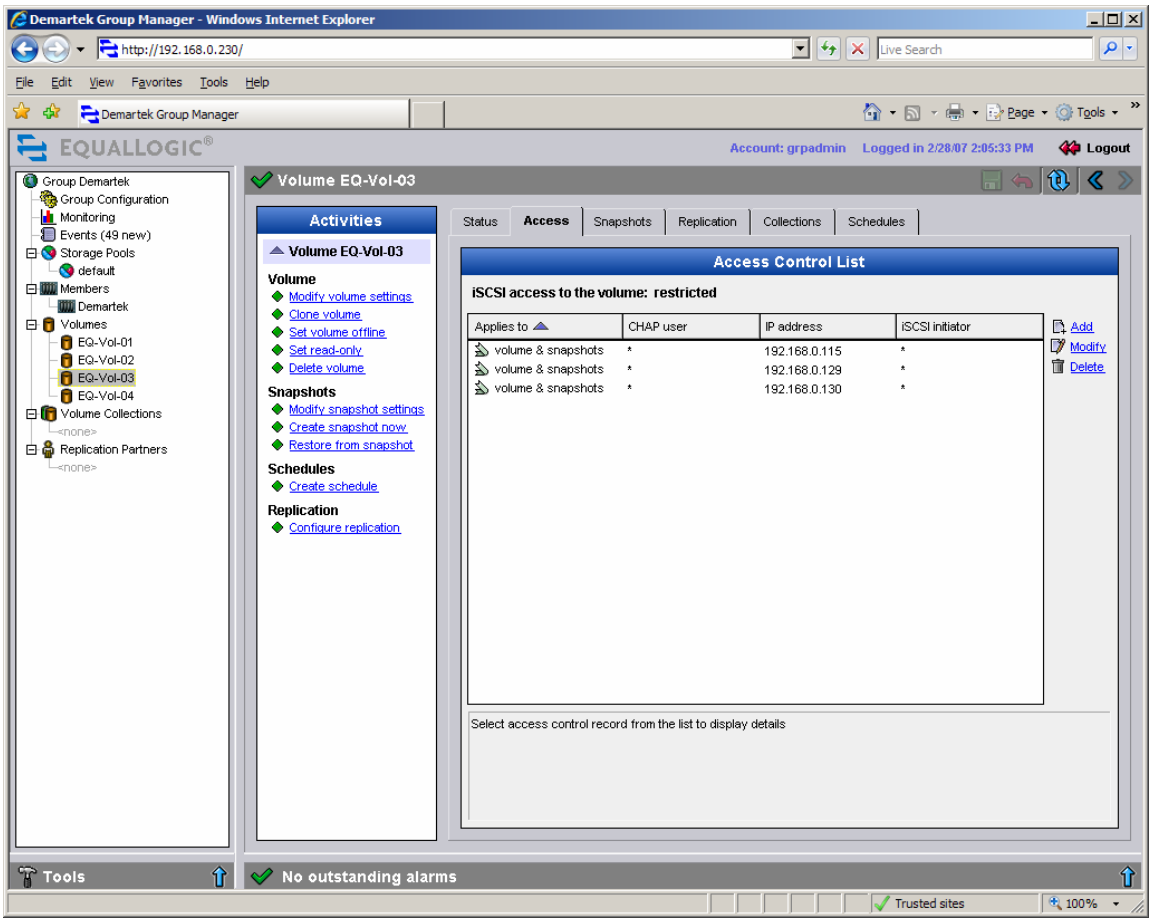

4. Make Volumes Ready for Use

The Volumes are ready for use as soon as the LUNs are created in the previous step.

5. Create iSCSI Targets

The targets are created when the iSCSI initiator addresses are supplied in the LUN creation step above. Access can be restricted to iSCSI initiators by their CHAP credentials, IP address, IQN or any combination of those three parameters.

6. Create Multi-path I/O for iSCSI Targets (optional)

The EqualLogic system will use multi-path I/O if at least two of its network interfaces have been enabled, and the iSCSI initiator is configured to use multi-path I/O. The EqualLogic system will automatically choose the paths to use for each I/O session. The EqualLogic system may vary the paths chosen between I/O sessions. When the iSCSI initiator is configuring multiple paths to the EqualLogic system, the iSCSI initiator will only see the “group” IP address. The EqualLogic system will handle the multi-pathing behind the “group” address.

7. Configure Security for iSCSI Targets (optional)

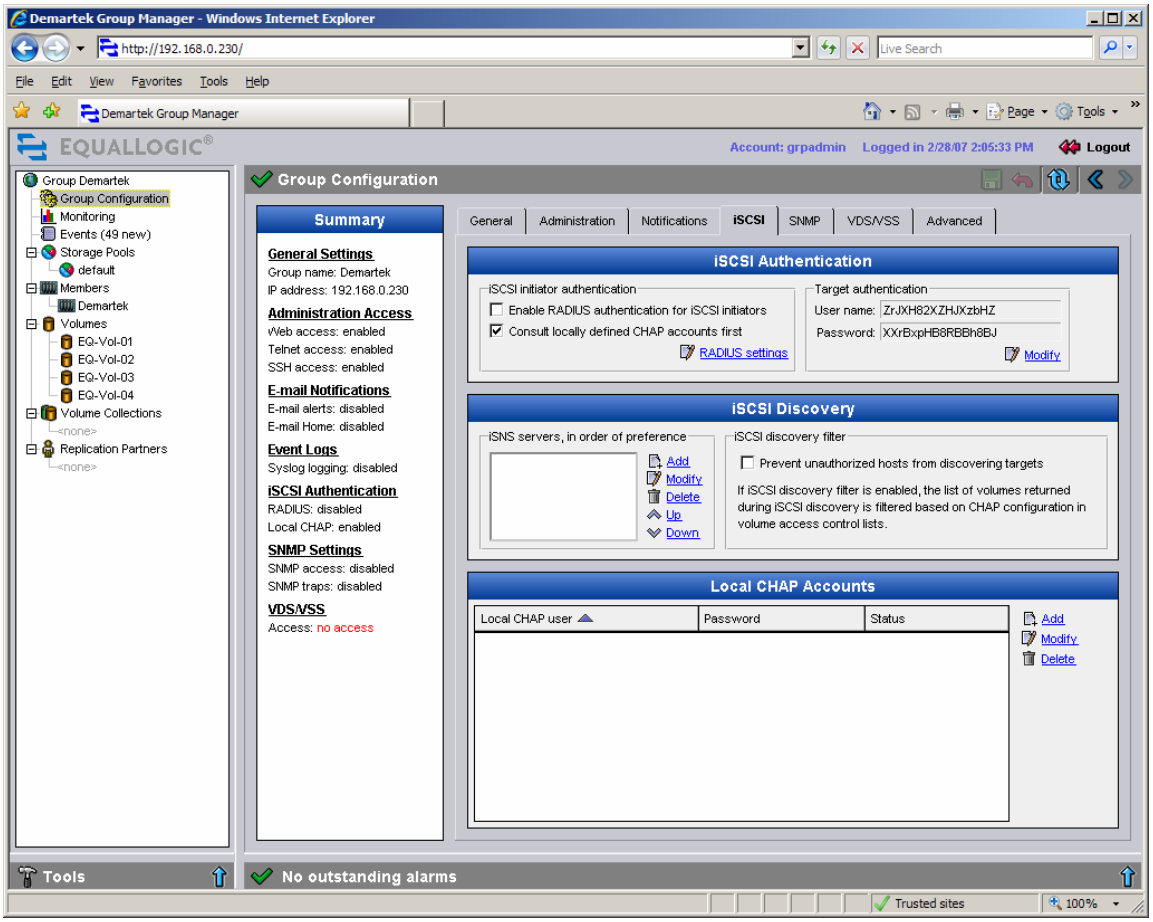

8. Make iSCSI Targets Ready for Use for iSCSI Initiators

No further steps are required to make the iSCSI targets ready to be accessed by the iSCSI initiators.

Initiator Configuration Steps

1. Configure Multi-path I/O from Application Host

The management console provides the option to activate the other two network interfaces on the array. This was done by enabling and modifying the settings of the two network interfaces that were not activated in the earlier steps.

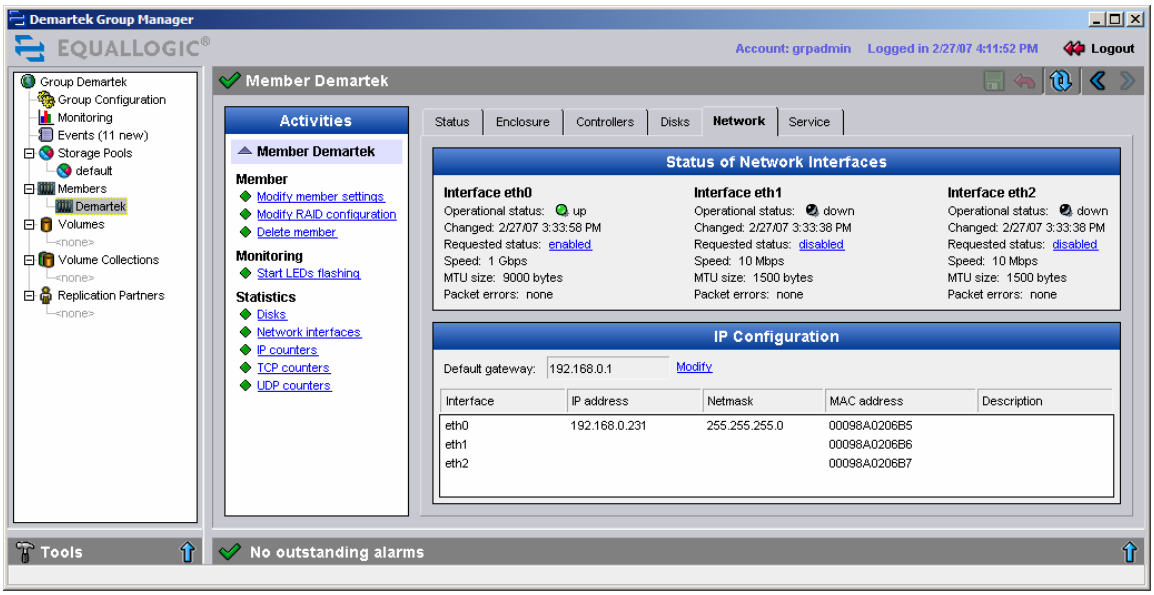

Each of the unassigned network interfaces, *eth1* and *eth2*, were assigned an address and activated. All three network interfaces must use the same subnet mask and default gateway.

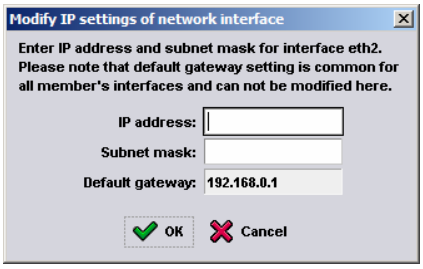

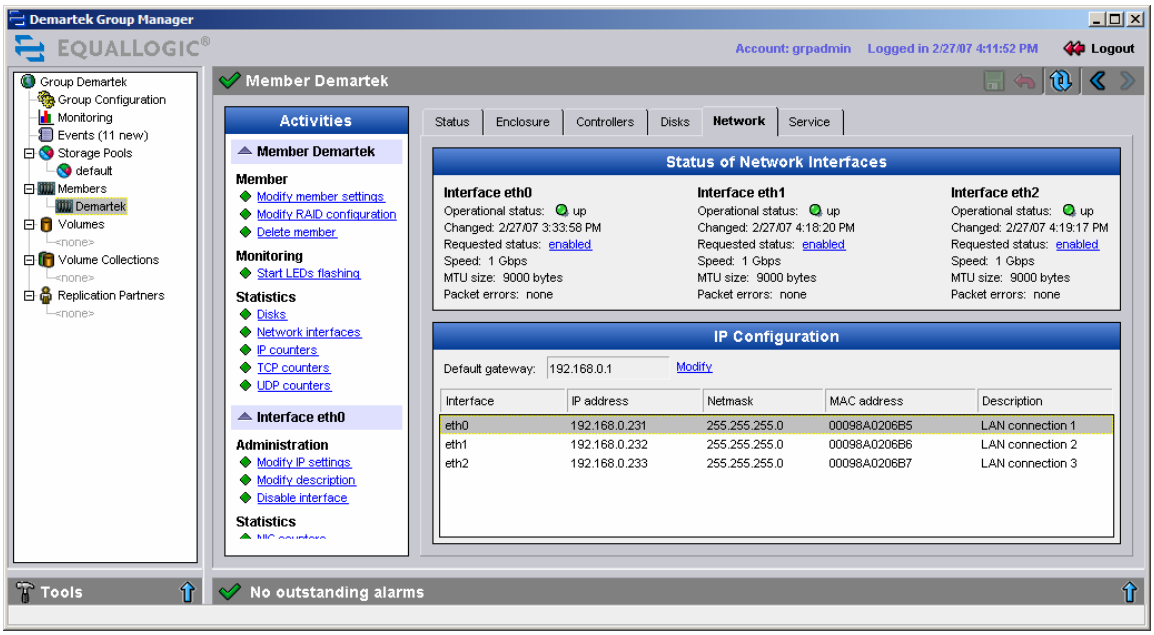

The EqualLogic MPIO DSM (Device Specific Module) must also be installed onto the iSCSI initiator hosts. The DSM is available on the Host Integration Tools CDROM or from EqualLogic's website.

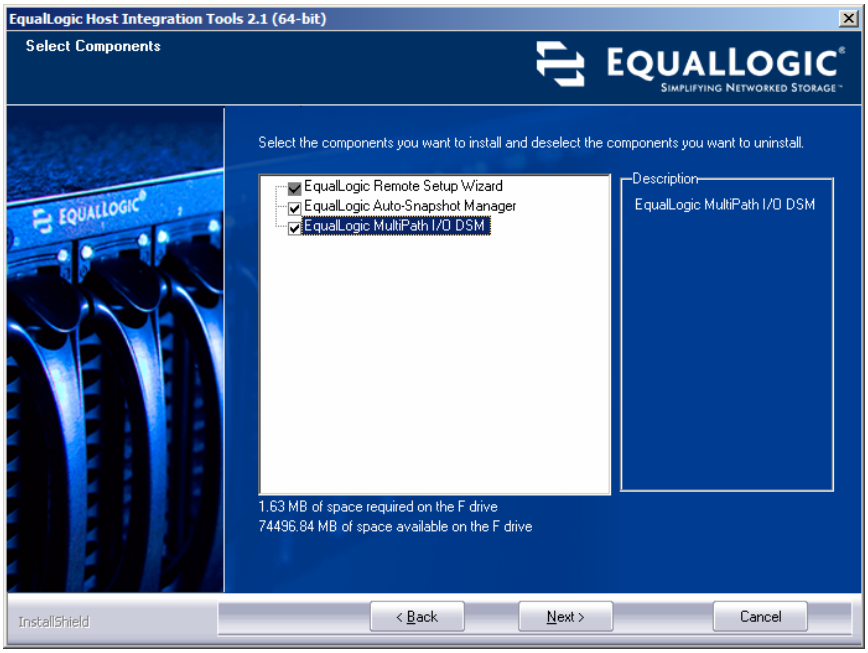


EqualLogic Host Integration Tools 2.1 (64-bit)
EQUALLOGIC
SIMPLIFYING NETWORKED STORAGE
Maintenance Complete
EqualLogic Host Integration Tools maintenance completed successfully.
Click Finish to exit.
InstallShield
< Back
Finish
Cancel

Basic Performance Results

The following performance data is not intended to be viewed as a comprehensive performance benchmark, but to provide a general performance overview for the EqualLogic PS3800XV Solution.

Selected performance results are shown below, using a standard server-class network adapter, with receive-side scaling on the host. This configuration used two paths from one host, two I/O workers, simultaneously accessing two target volumes and a queue depth of 5. Each volume shared round-robin access across both paths.

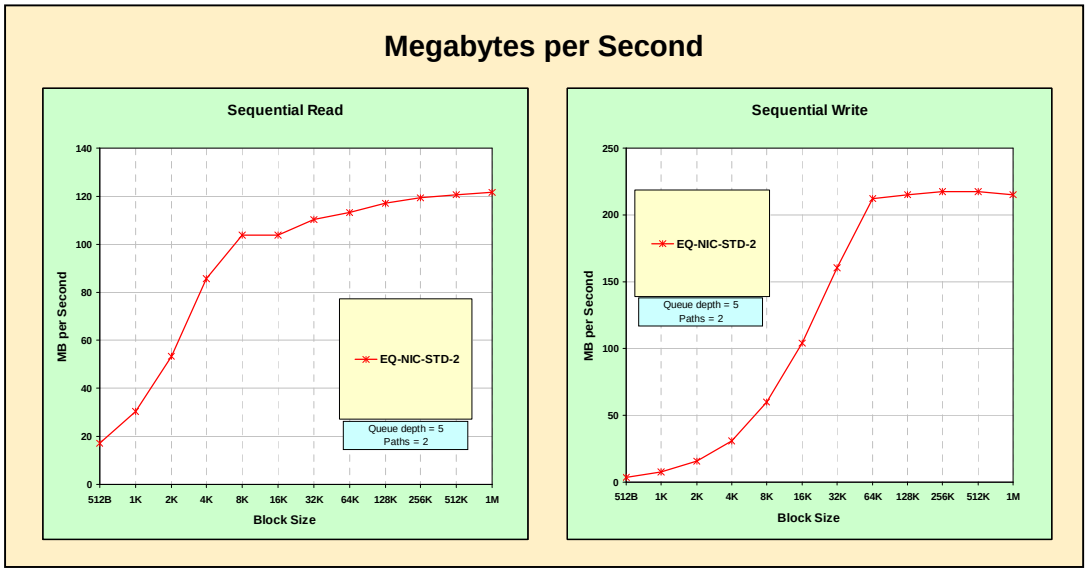

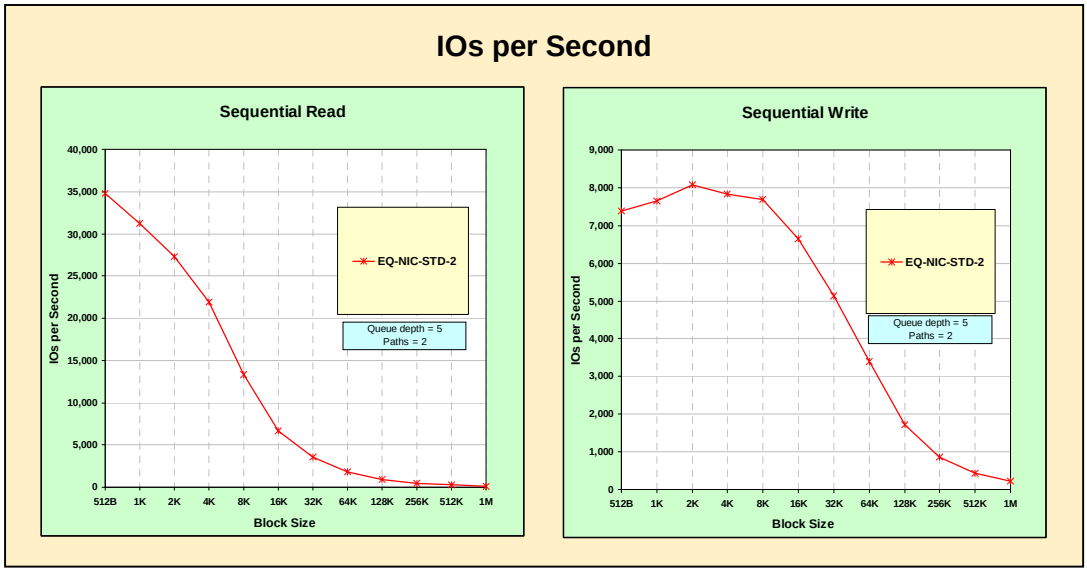

HDS TagmaStore™ AMS1000

The HDS TagmaStore AMS1000 is an iSCSI target solution that includes a dual-controller disk subsystem with SATA and/or Fibre Channel disk drives with raw capacities up to more than 200 TB. It includes cache binding, cache partitioning and a variety of hardware copy and management functions. In the iSCSI configuration it supports RAID 1, RAID 1+0, RAID-5 and RAID-6. The AMS1000 supports iSCSI, Fibre Channel and NAS protocols with the additional benefit of supporting two protocols simultaneously.

Hitachi's Adaptable Modular Storage is enterprise class storage designed and priced to meet the needs of small and medium businesses. The AMS line supports FC and iSCSI protocols and consists of three models, AMS200, AMS500 and AMS1000 and cost effectively support storage needs from under one terabyte to over 300 terabytes with the AMS200 and AMS500 being upgradeable in the rack. The AMS models have many features unique to the modular market including: RAID 6 (dual parity for highest availability); modular volume migration (non disruptive movement of volumes within the storage array); cache partitioning (allotment of cache memory size to specific applications) as well as support for both Fibre Channel and Serial ATA disk drives. When combined with Plug and play kits for Microsoft Simple SAN, Windows servers attached to AMS storage models can easily manage their storage as well as their entire storage area network.

Target Configuration Steps

1. Configure Network Settings for iSCSI Target Ports

The AMS1000 is initially configured for access by the Storage Navigator Modular software. Change the Storage Navigator Modular setting from "Operation Mode" to "Maintenance Mode". Select the desired AMS1000 array. From the Array System viewer select "Tools/Configuration Settings/iSCSI tab" to change the network settings for the iSCSI ports.

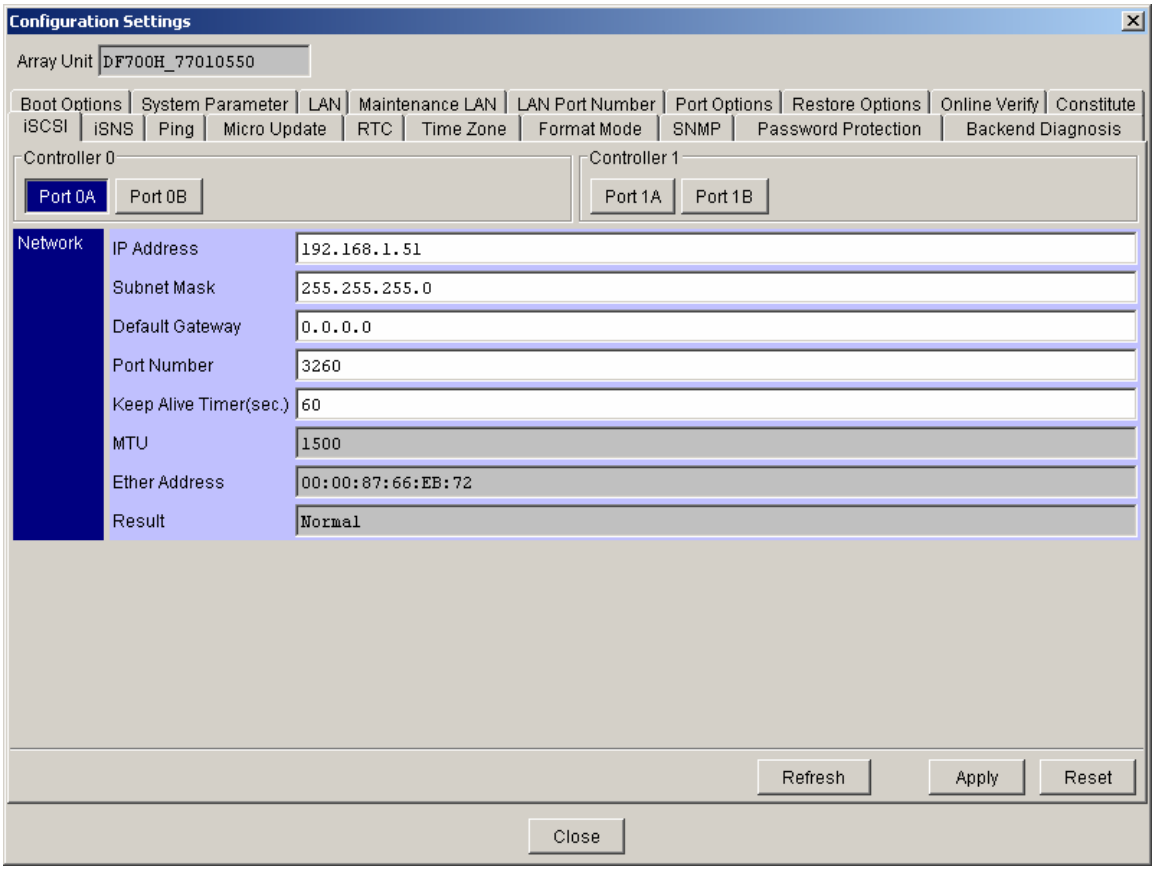

2. Launch Management Console

The AMS1000 management console, also known as Storage Navigator Modular, provides access to all array management functions.

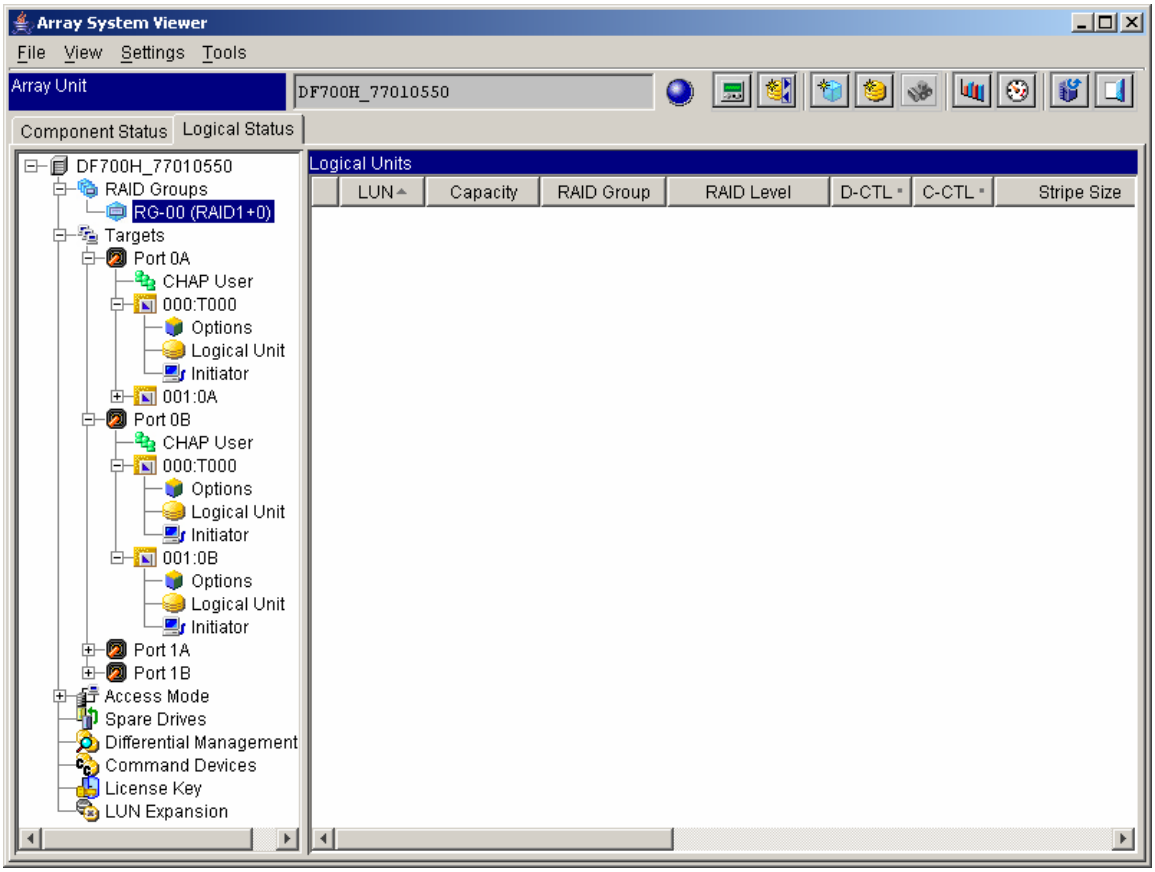

3. Create LUNs on Disk Array

To create a LUN, the “Create Logical Unit” option is selected. This particular array was previously configured as a RAID 1+0 by the hardware field engineer.

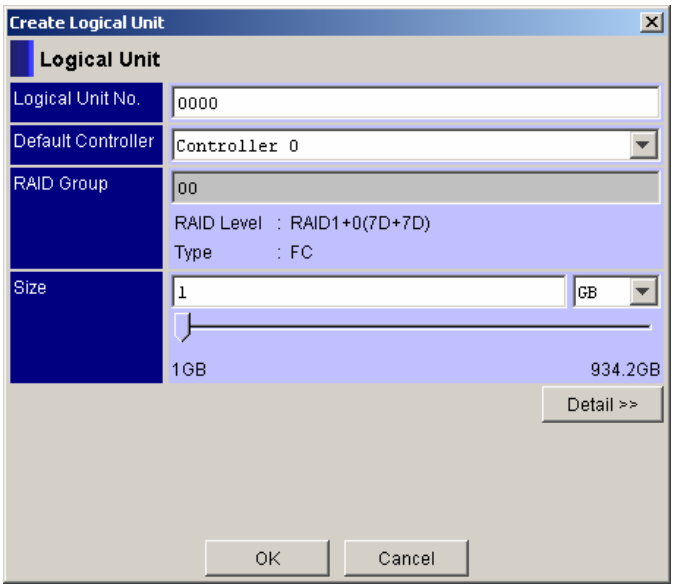

Note – It is important at this point to note that the storage solution LUN size should not be confused with the size of the iSCSI target. The iSCSI target will be configured in a later step and is associated with the storage needed for a particular application on the host server. It is recommended that the LUN size on the storage hardware be as large as reasonably possible to allow the storage subsystem to optimize the use of the physical disks underlying the LUN that is created. In this case, as shown below, we are choosing to create one LUN at the maximum size available for this hardware. The iSCSI targets created later will fit into this one LUN, based on the needs of the host application.

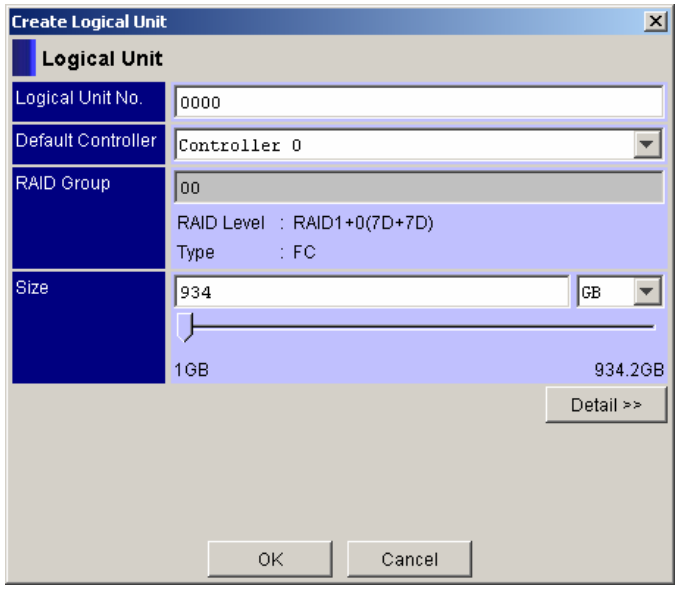

After the confirmation window appears and is accepted, the LUN is created.

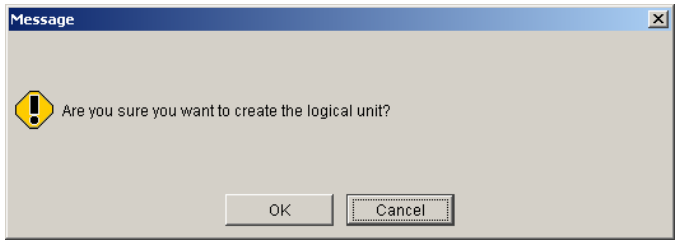

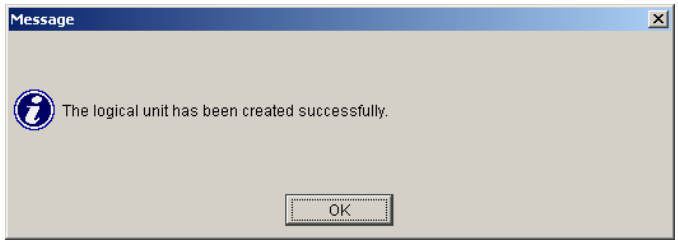

The management console shows the configuration.

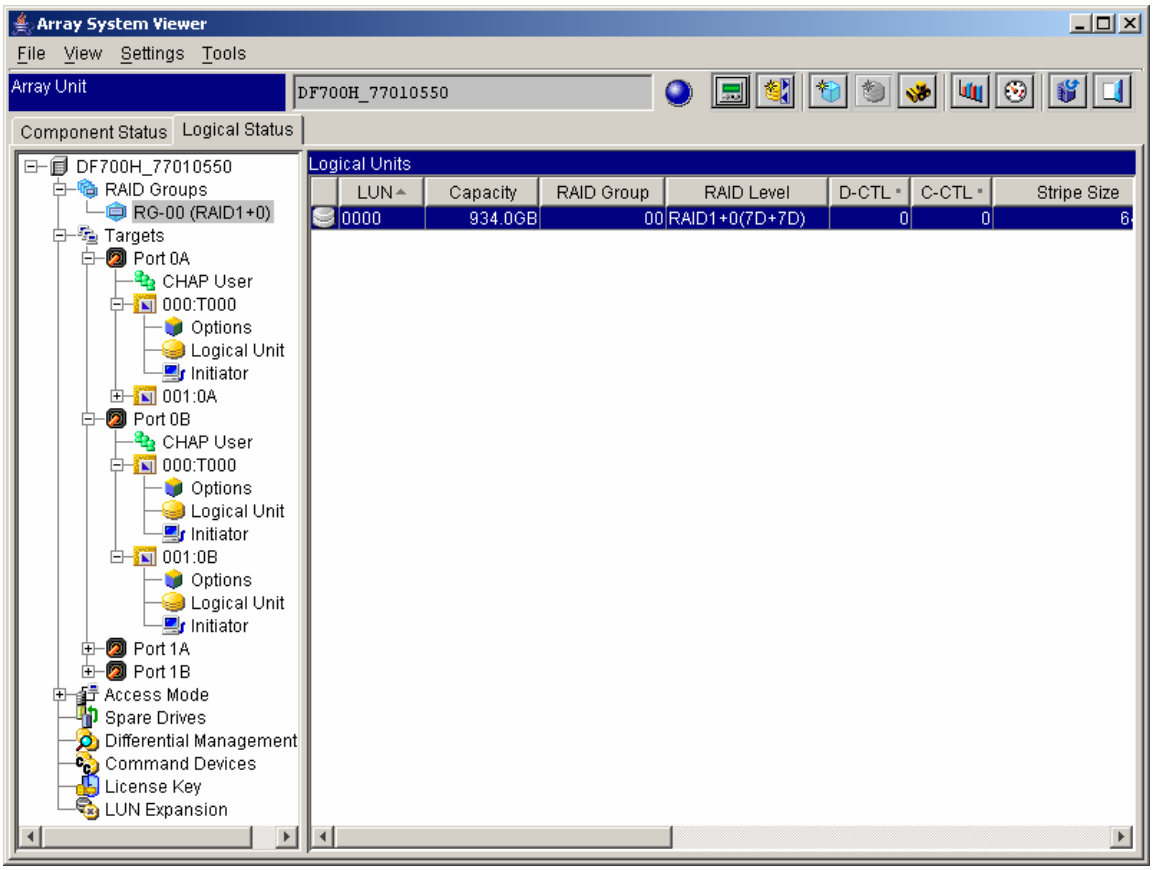

4. Make LUNs Ready for Use

The LUN must be formatted before it can be used by the hosts requiring storage. The format process runs in the background in the array and the LUN can be presented to hosts while the format operation is taking place.

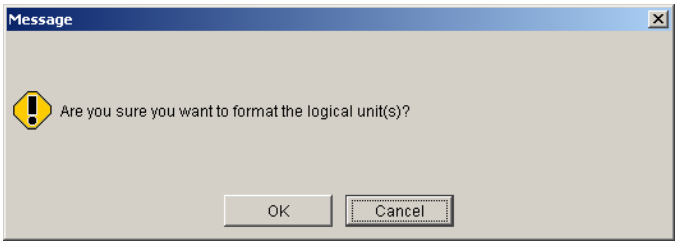

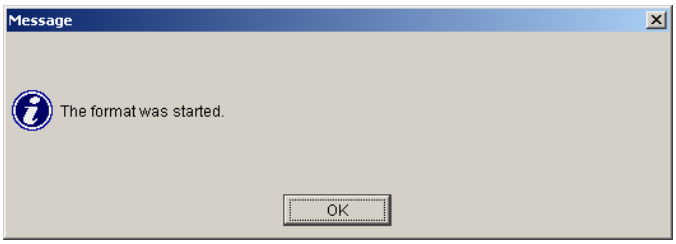

5. Create iSCSI Targets

The LUN must be mapped to an iSCSI target before an iSCSI initiator can access it. This LUN will be mapped to two iSCSI targets, and then multi-path information will be applied at a later step.

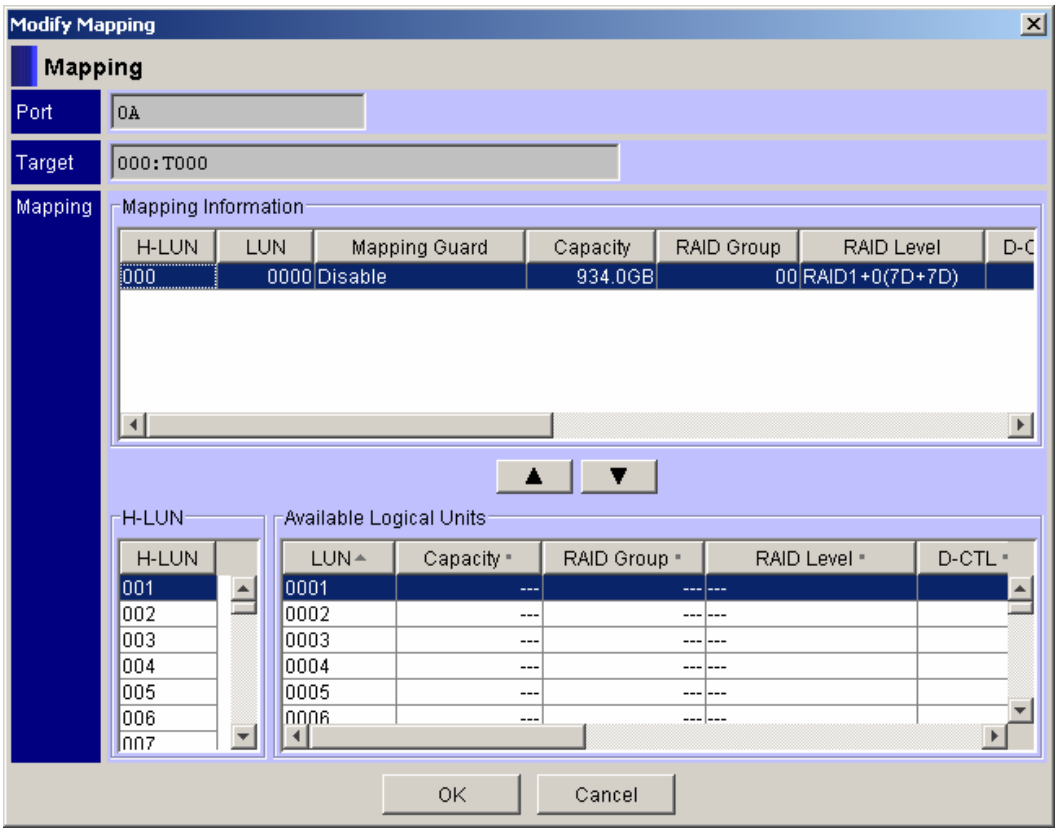

The LUN is selected and it appears in the lower list. This LUN is mapped to target 000:T000 which belongs to port 0A. The port number will be used later as the internal target identifier.

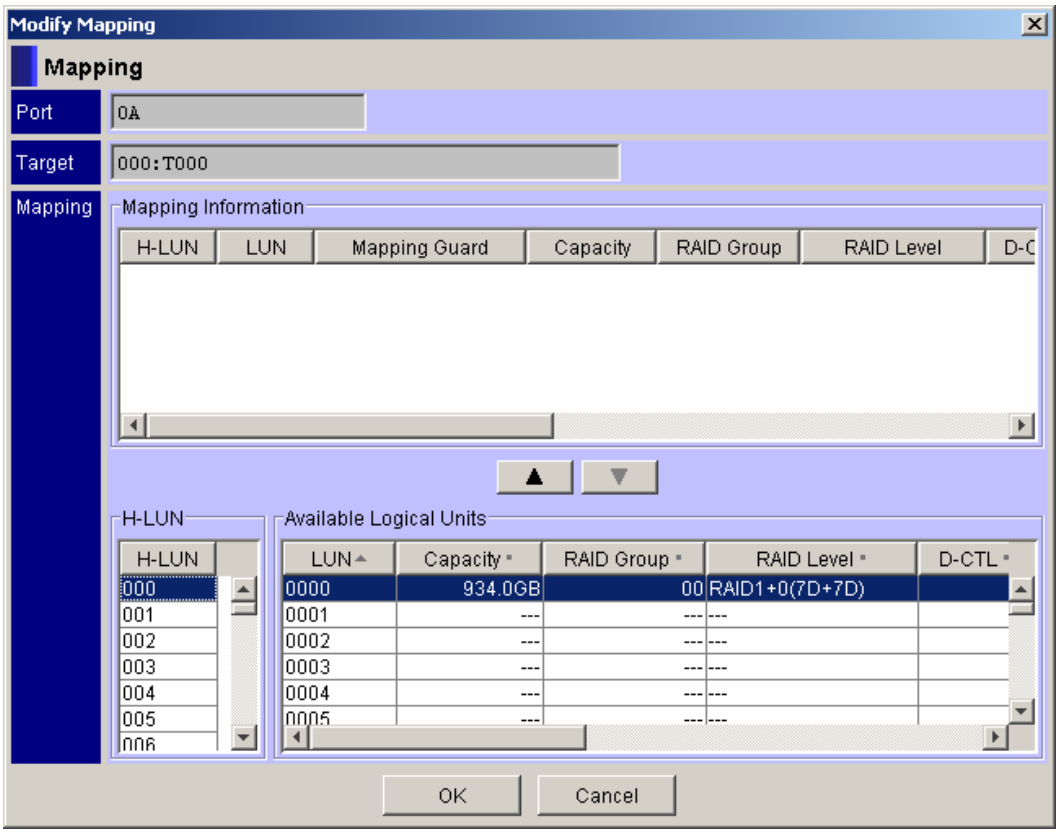

6. Create Multi-path I/O for iSCSI Targets (optional)

The multi-path functions are handled by the Hitachi Dynamic Link Manager (HDLM) MPIO DSM that is installed on the application host server.

7. Configure Security for iSCSI Targets (optional)

The iSCSI target properties include the CHAP information, which can be set by choosing the authentication method drop-down box.

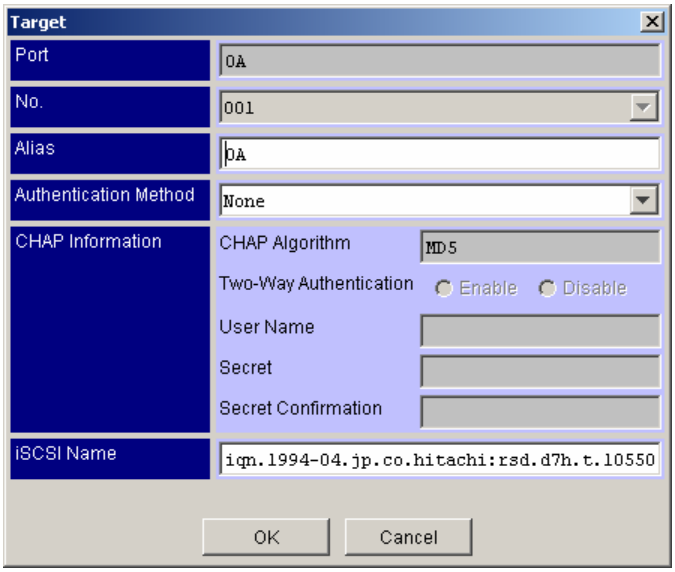

8. Make iSCSI Targets Ready for Use for iSCSI Initiators

The iSCSI initiator is assigned to the iSCSI target. Because multi-path has been selected, the same iSCSI initiator is assigned to targets 0A and 0B.

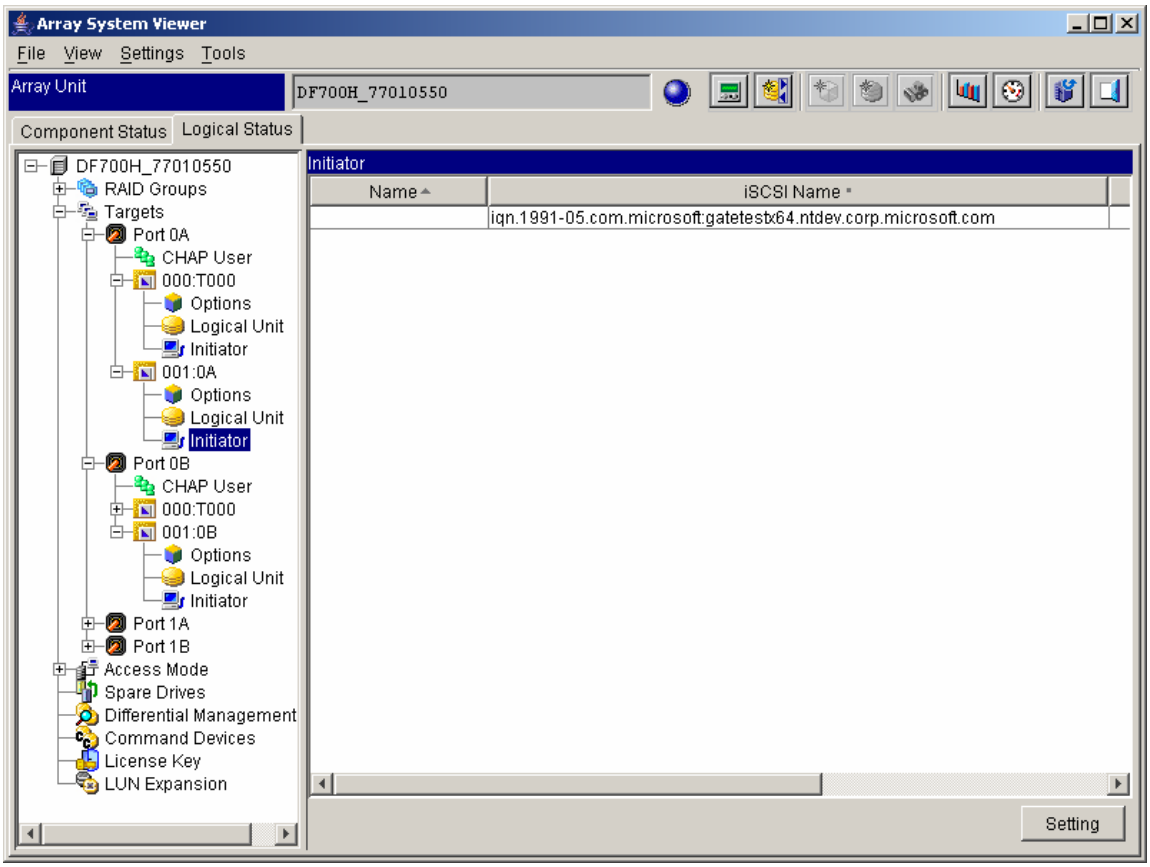

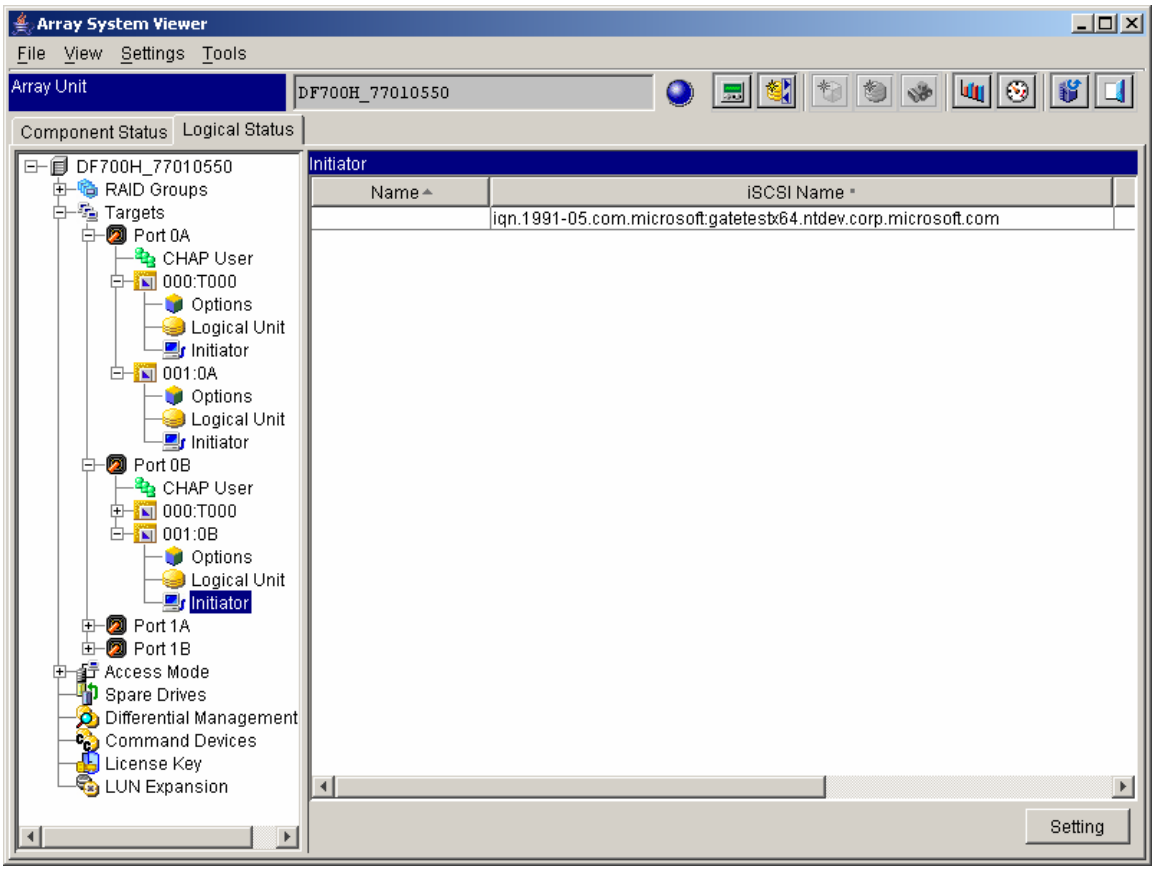

The target is now ready for use.

Initiator Configuration Steps

1. Configure Multi-path I/O from Application Host

Hitachi uses the Hitachi Dynamic Link Manager (HDLM) to manage multiple-path I/O. HDLM is implemented as a Microsoft MPIO DSM which is onto the application host server.

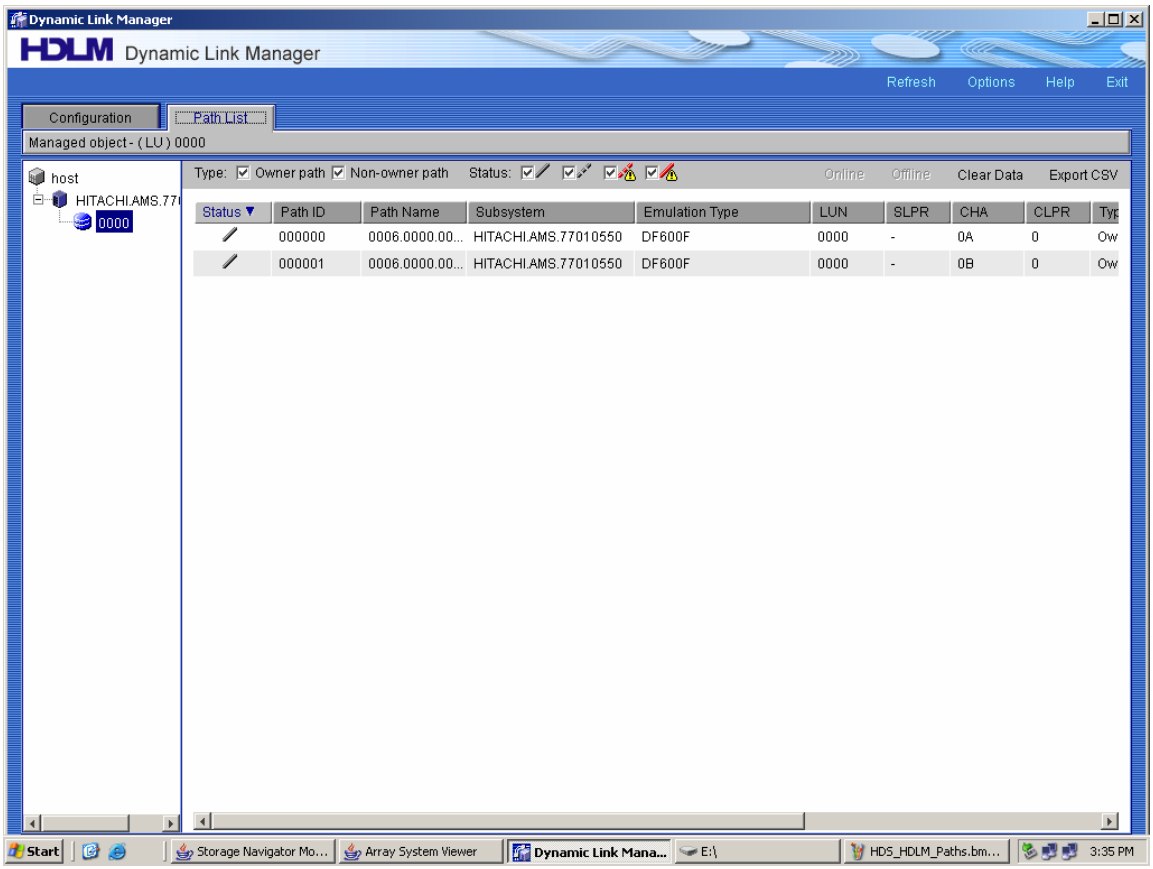

The two paths are assigned to the same iSCSI initiator in the HDLM software so that the iSCSI initiator software can take advantage of the multi-path options.

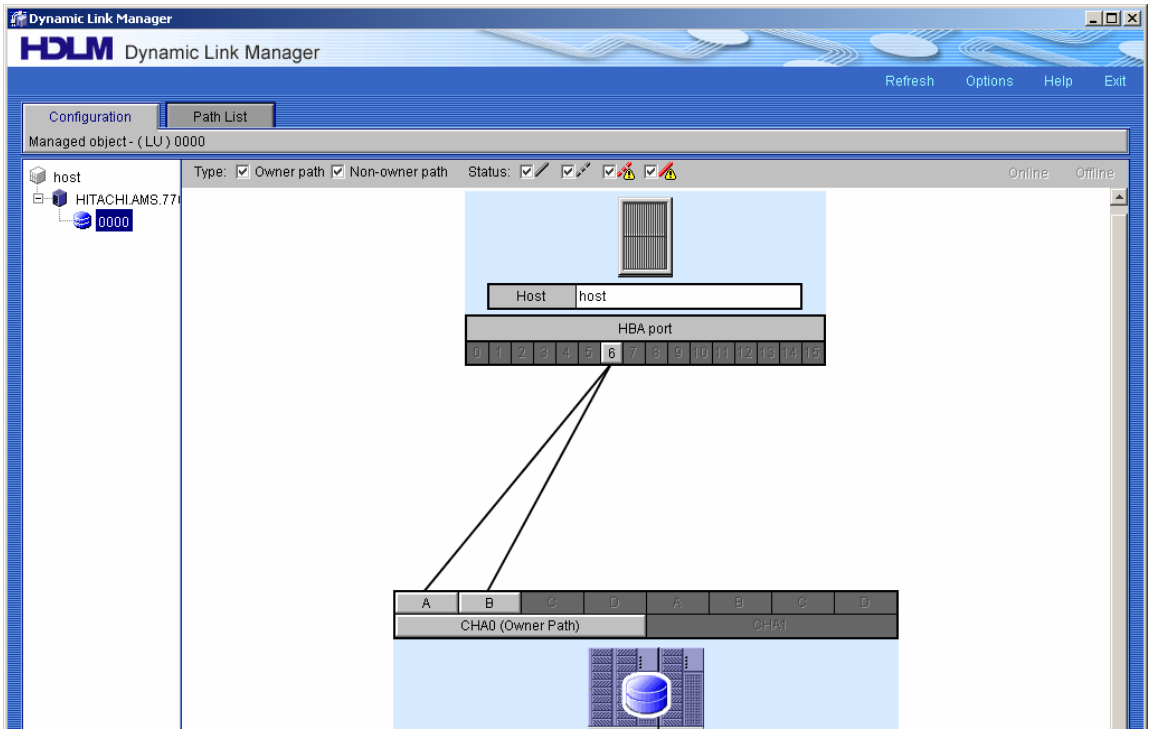

The Microsoft iSCSI initiator software is then configured for multi-path in the standard fashion.

Basic Performance Results

The following performance data is not intended to be viewed as a comprehensive performance benchmark, but to provide a general performance overview for the HDS TagmaStore AMS1000 Solution. Performance will vary depending on drive speeds, number of drives, applications and many other factors.

HDS has a significant body of performance information available that will accurately represent the performance ability of the AMS 1000. Please contact HDS for more detailed performance data.

Selected performance results are shown below, using a standard server-class network adapter, with receive-side scaling on the host. This configuration used two paths from one host, four I/O workers, accessing four target volumes and a queue depth of 10.

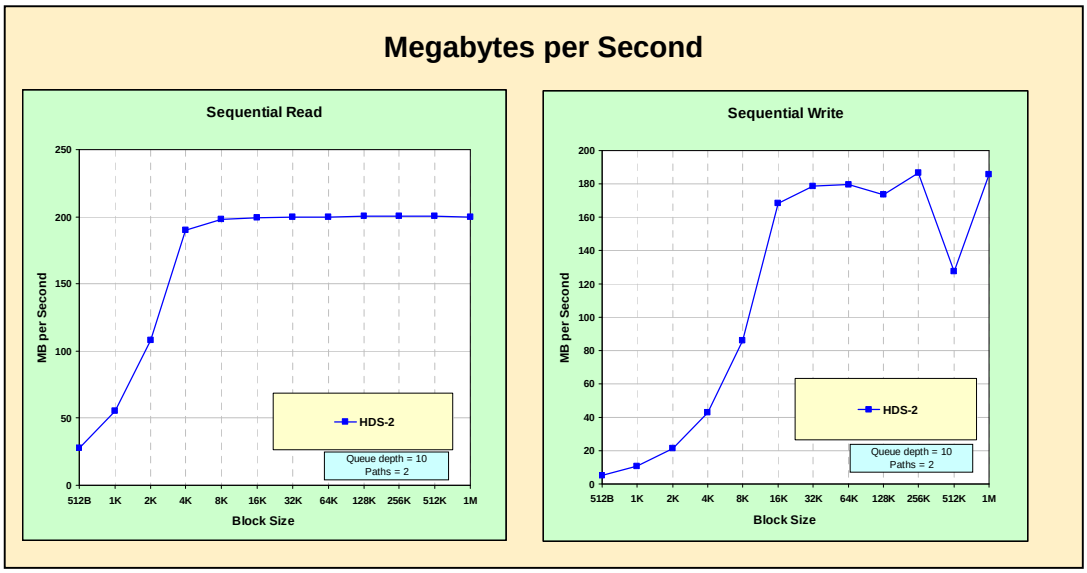

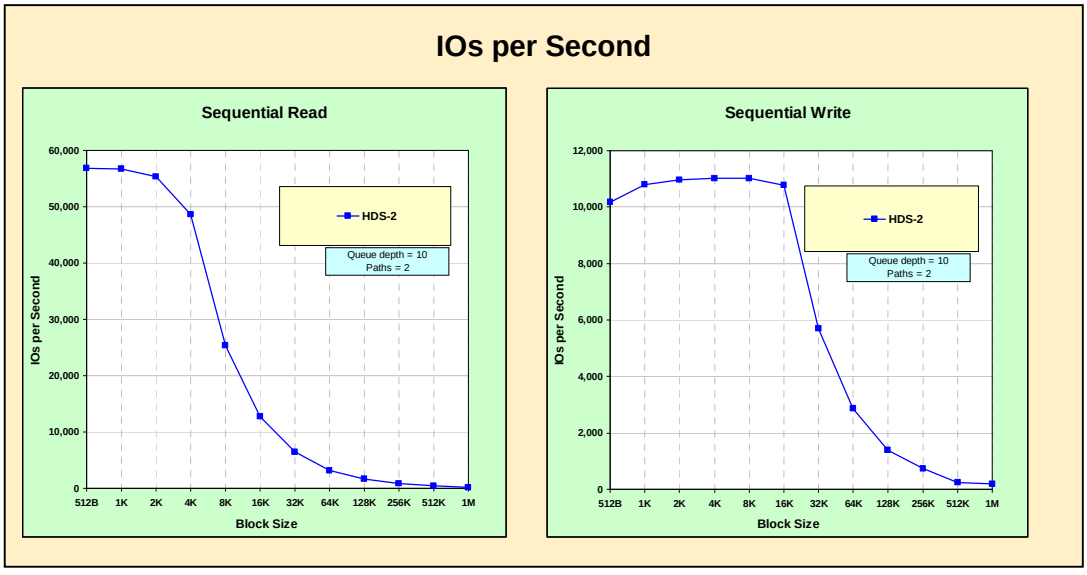

HP StorageWorks 1200 All-in-One Storage System

The HP StorageWorks 1200 All-in-One Storage System (AiO1200) is based on Microsoft Windows Storage Server 2003 R2 (WSS) and the Microsoft iSCSI software target. It is a unified NAS device and iSCSI target solution which also includes integrated data protection software and a management console designed for IT generalists who may be new to storage configuration. The AiO1200 is built on the HP ProLiant hardware platform and has a disk subsystem of 12 internal SATA or SAS disk drives available in various storage capacities. The HP All-in-One Storage System is also available in 4 and 6 drive configurations all of which support external storage expansion to direct attach storage enclosures. The version tested for this report included 12 drives and 3.6 terabytes (TB) of raw storage capacity.

The HP All-in-One Storage Manager (ASM) is a unique toolset which is included in the AiO and is designed to reduce the time, expense, and expertise required to deploy and manage storage resources. ASM hides much of the complexity traditionally associated with storage configuration and presents storage in an application-centric context rather than a storage-centric view. The toolset integrates with several Microsoft applications and includes data migration tools for Exchange Server 2003/2007 and SQL Server 2000/2005. The following configuration steps will illustrate how storage can be configured using ASM or how other system-level tools can be used for more granular control.

Target Configuration Steps

1. Configure Network Settings for iSCSI Target Device

The main HP All-in-One management console includes a “Rapid Startup Wizard” which automatically launches at first boot and is used to simplify initial system configuration.

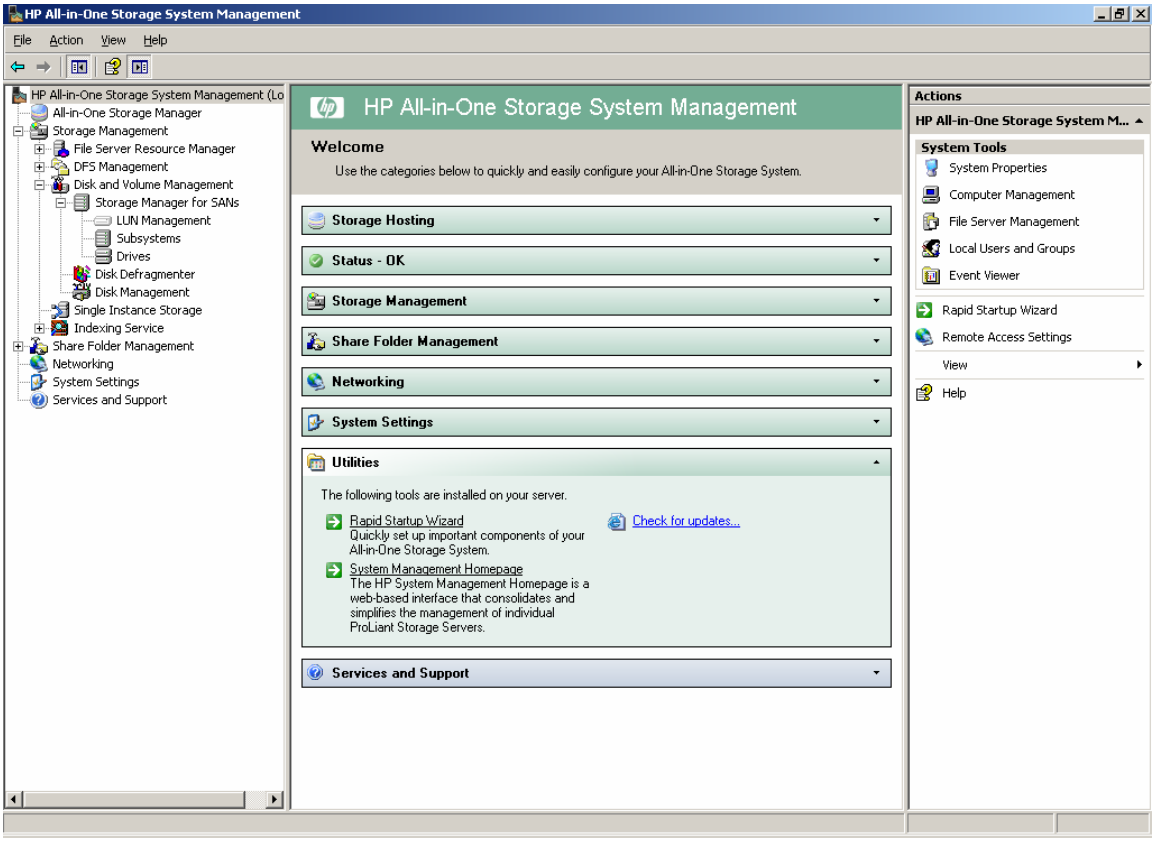

The Rapid Startup Wizard includes a place to set the network configuration. This system was configured to use DHCP on its three Ethernet network interfaces.

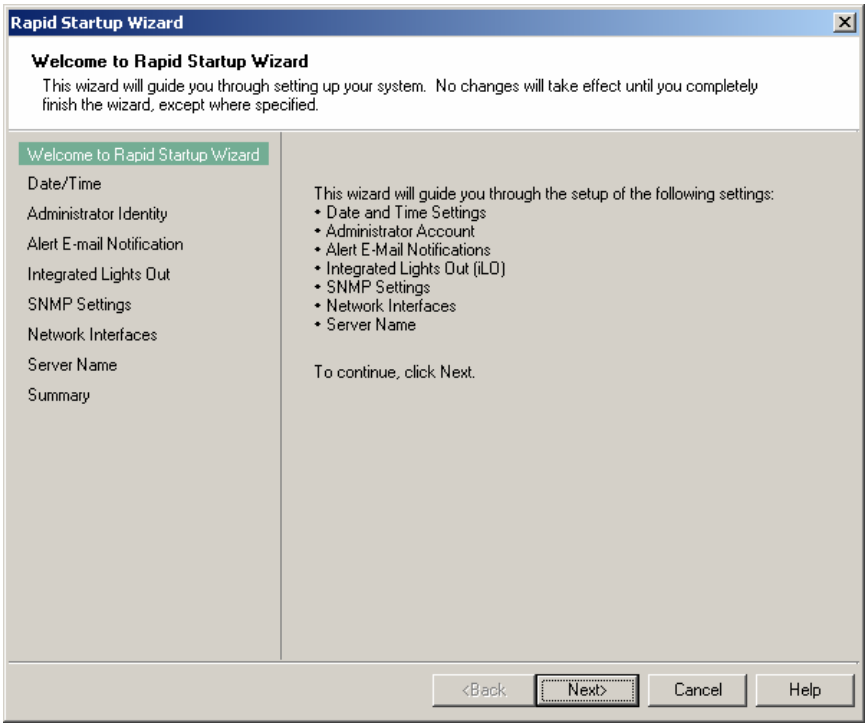

2. Launch Management Console

All the storage management functions for the HP StorageWorks 1200 All-in-One Storage System are performed from the management console, shown below.

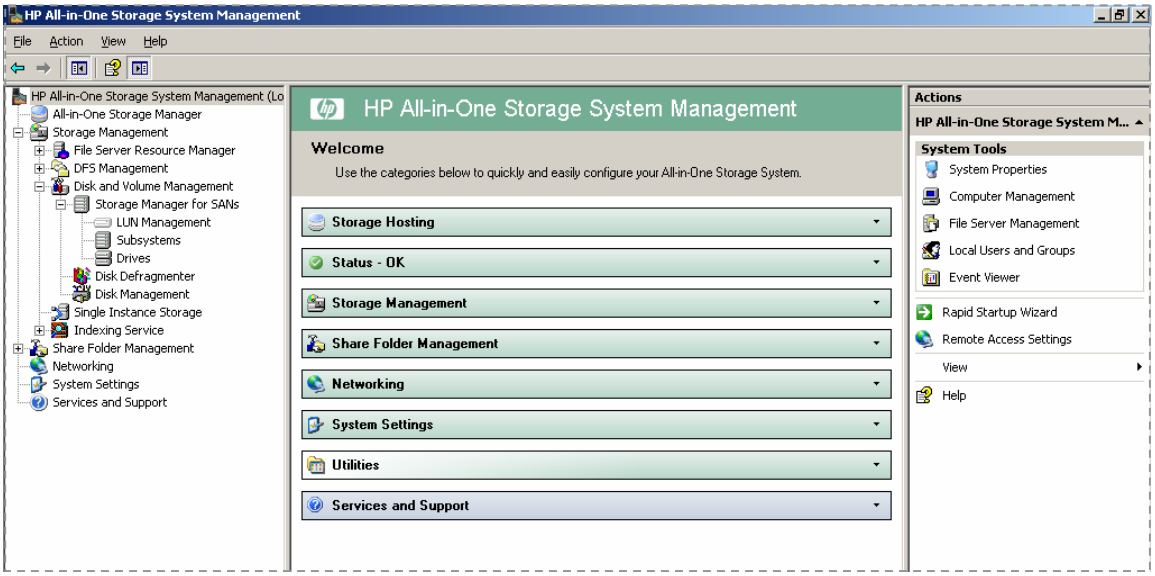

In the center section of the management console, several categories of management functions are available.

3. Guided configuration: Link to the All-in-One Storage Manager (ASM)

ASM is an 'application-centric' management interface into the file shares and iSCSI target storage on the AiO and shows capacity usage in terms of managed data areas. ASM exists as the first node in the management console tree. The main ASM management screen is shown below.

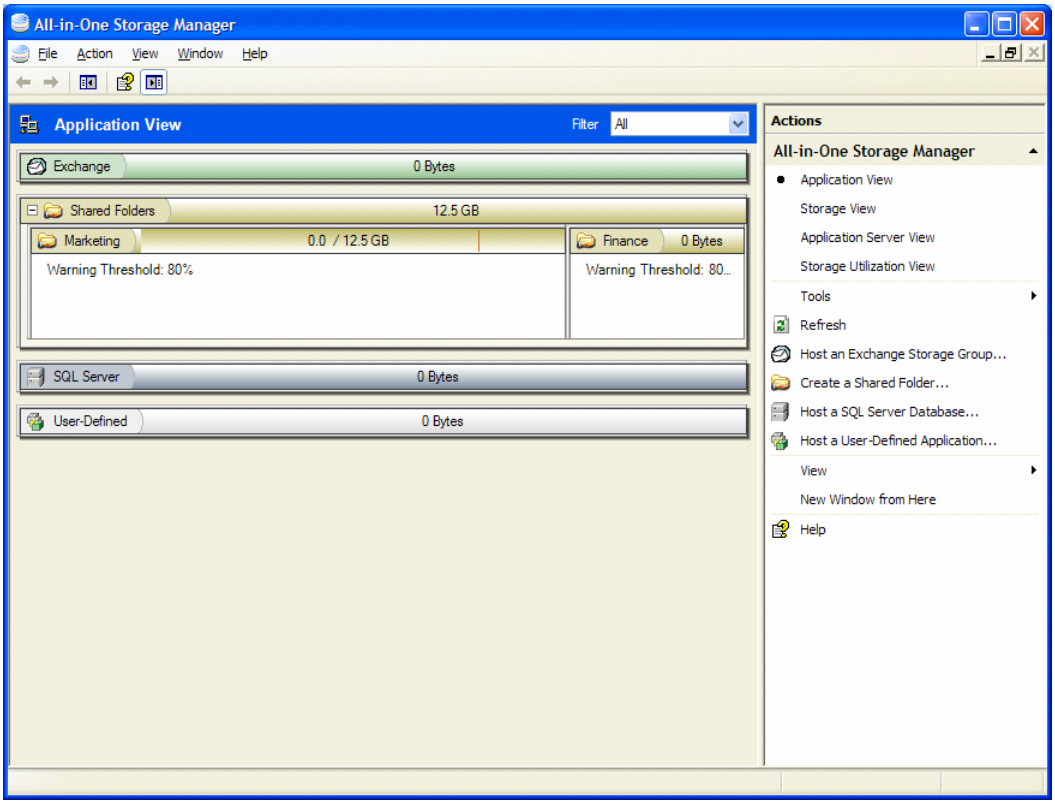

The interface includes several wizards which automate configuration, provisioning, and data migration. These include

- Host an Exchange Storage Group
- Create a Shared Folder
- Host a SQL Server Database
- Host a User-Defined Application

Prior to launching any of the interface wizards, administrators will need to install an ASM agent onto their host server (Windows Server 2003 32-bit or 64-bit). This agent packages the Microsoft iSCSI initiator, an ASM Service which runs on the host, as well as some Microsoft .NET components to enable communication between the host server and the AiO. The agent installation process also sets up application specific security parameters for SQL Server and Exchange.

As an example of the guided configuration capabilities of the AiO we will walk you through the ‘Host an Exchange Storage Group” tool.

Clicking the link on the right side of the interface launches the wizard. The first input screen asks the administrator to specify the Name or IP address of the Microsoft Exchange server where they have installed the ASM Agent.

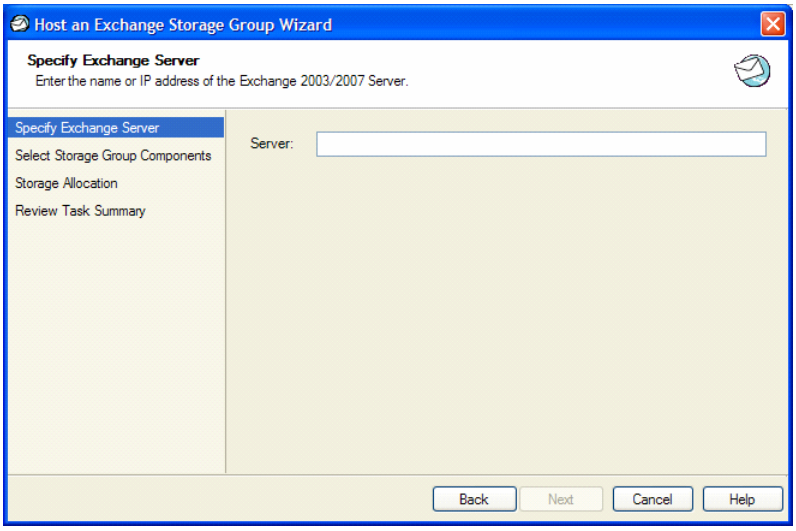

Clicking NEXT enables ASM to communicate with the service on the host server and brings up the 'Select Storage Group Components' screen. The service uses the Microsoft Exchange API to gather data about the installation and return it to the AiO administrator. Specifically, the administrator is prompted to choose the Storage Group(s) which need to be moved to the iSCSI target storage.

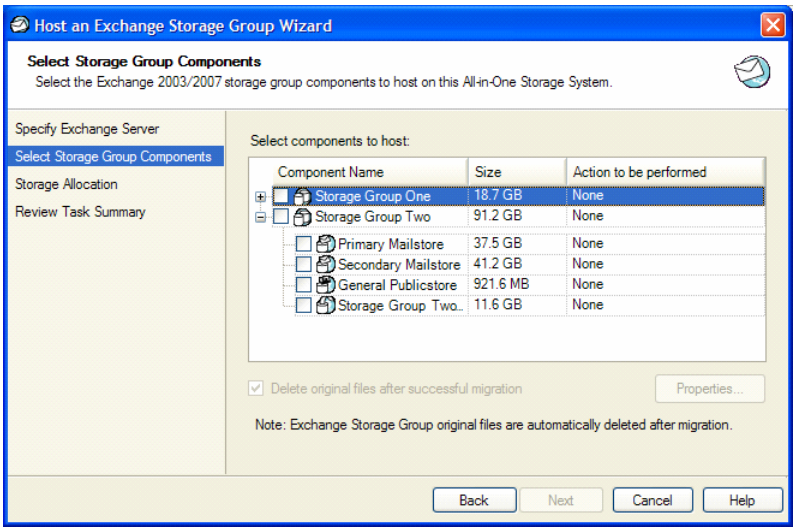

Clicking NEXT brings up the Storage Allocation screen. Here, ASM has recommended a set of storage settings based on Microsoft 'best practices' for Exchange; these include RAID type, stripe settings, exclusive disk use, warning threshold, etc. Administrators can accept these settings or click the Advanced button to manually override. They can also choose between HDD technologies (SATA and/or SAS) if both disk types reside on the AiO or any direct attached storage enclosures.

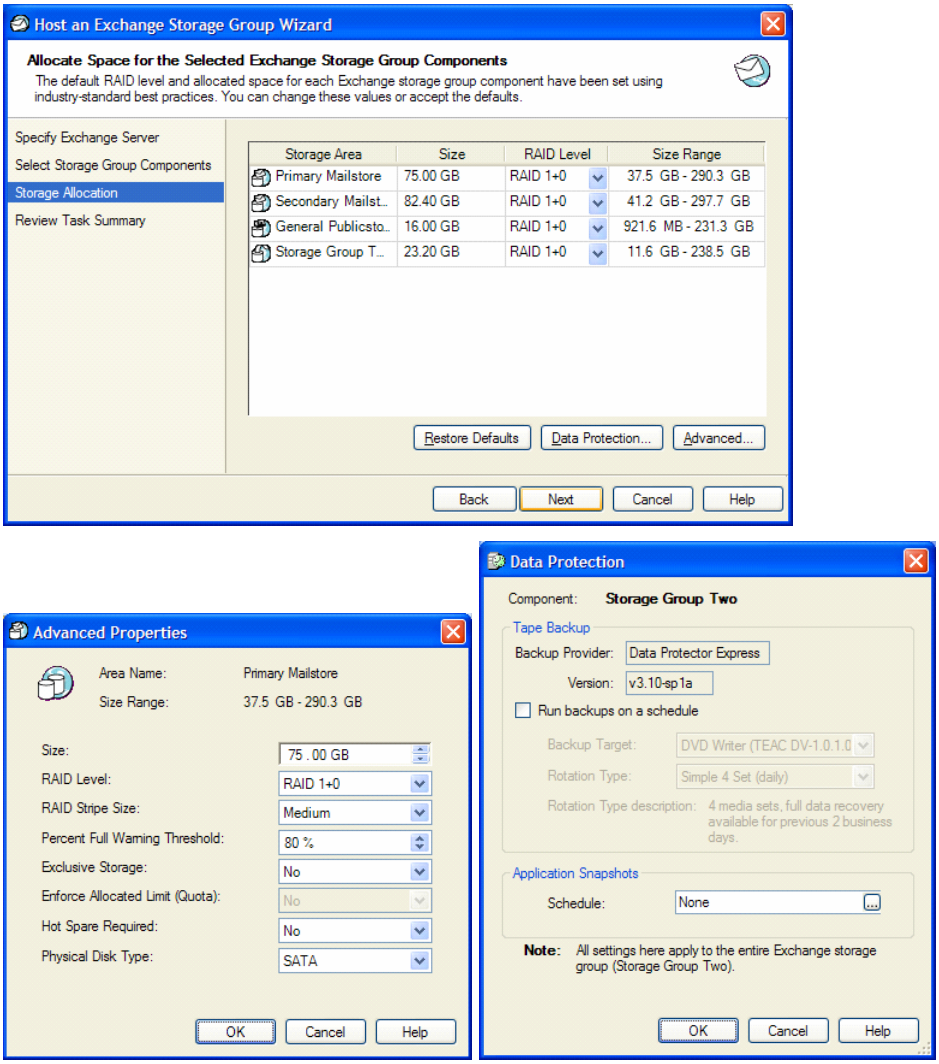

Clicking on the Data Protection button brings up a screen that allows administrators to configure VSS snapshots for this iSCSI Target storage and to configure a traditional backup job using the HP Data Protector Express software which is included on the AiO.

Clicking NEXT will bring up a task summary of the steps that ASM will automatically perform. Tasks can be run immediately or scheduled for later; in the case of data migration the process will tell the Exchange server to temporary halt the service so that data can be migrated to the newly created iSCSI LUNs and will then restart the service.

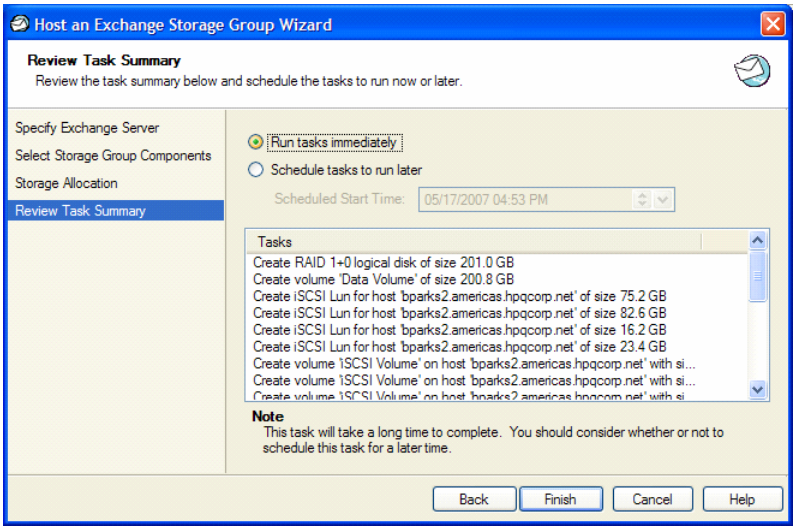

Once the job runs the system will automatically

- Create the logical drive on the physical disks
- Create the data volume
- Create iSCSI LUN and connect the host to the target
- Create and format the iSCSI volume
- Set directory quotas and alerts on the iSCSI target storage
- Migrate storage group components
- Create the backup job

Once the job is executed, storage tasks such as expansion and protection of the Storage Group can be managed from the main ASM interface as shown below.

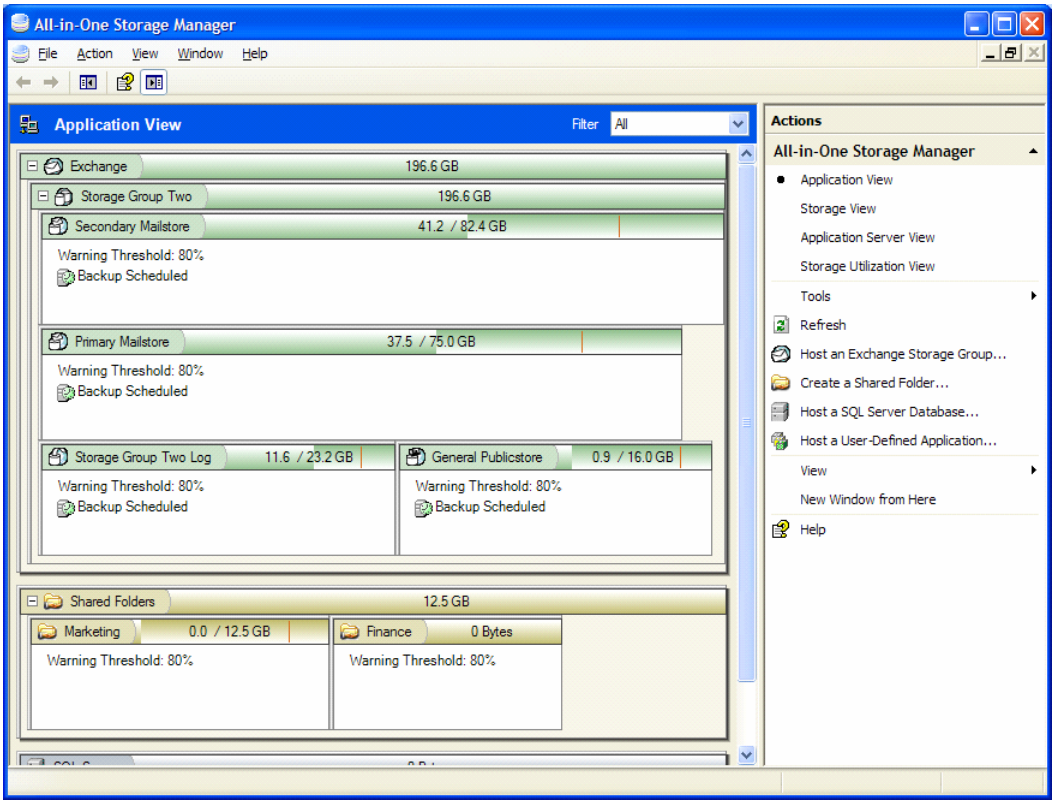

NOTE: Remaining steps 4 through 8 illustrate how to manually configure iSCSI target storage outside of the All-in-One Storage Manager. If using ASM, steps 4 through 9 are performed automatically by the various ASM application tools.

Manual configuration if not using ASM (Steps 4 through 9):

4. Create LUNs on Disk Array

To create the LUNs on the Disk Array, the administrator selects the “Array Configuration Utility” scenario from the “Storage Management” section of the management console. This launches the array configuration utility.

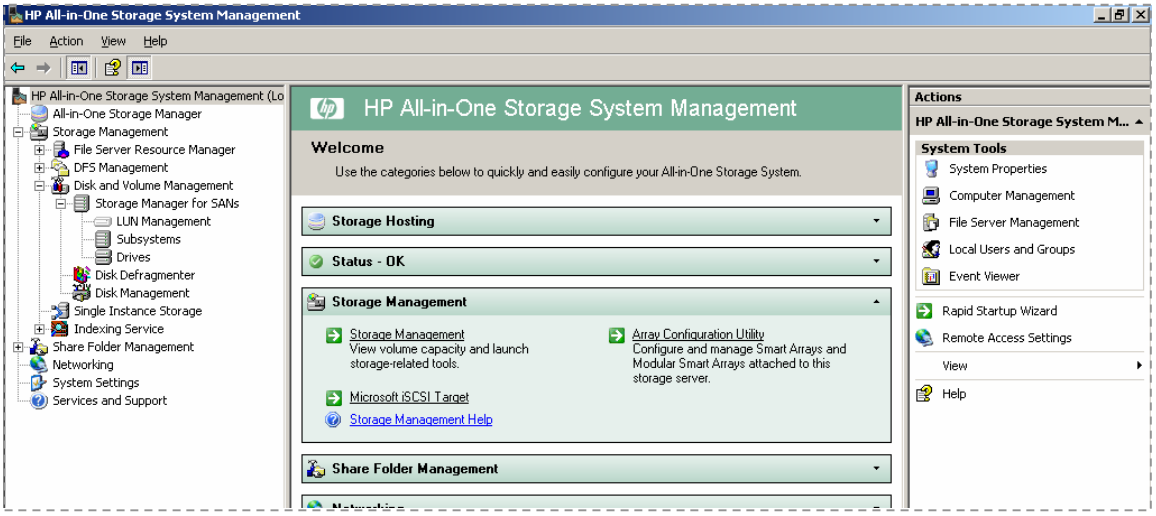

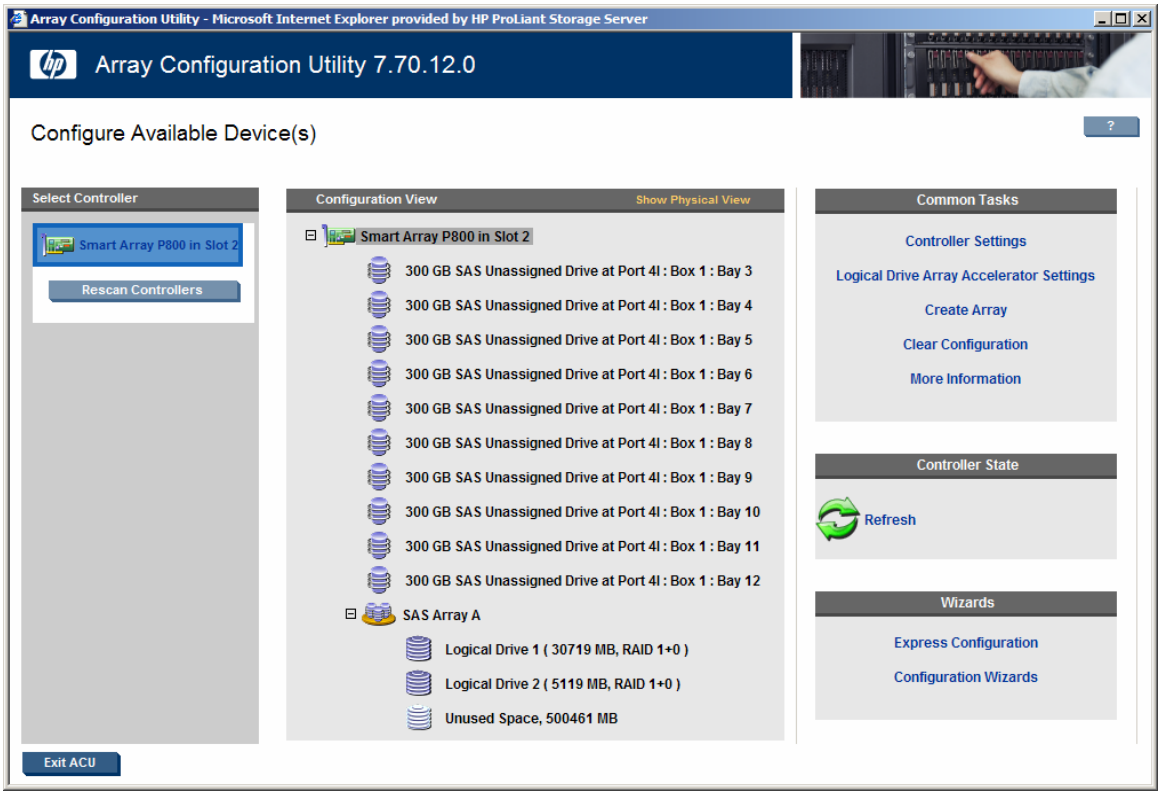

Highlighting the “Create Array” menu item allows the creation of a data volume across the unassigned drives in the system. All ten of the unassigned drives will be used for the iSCSI data volumes.

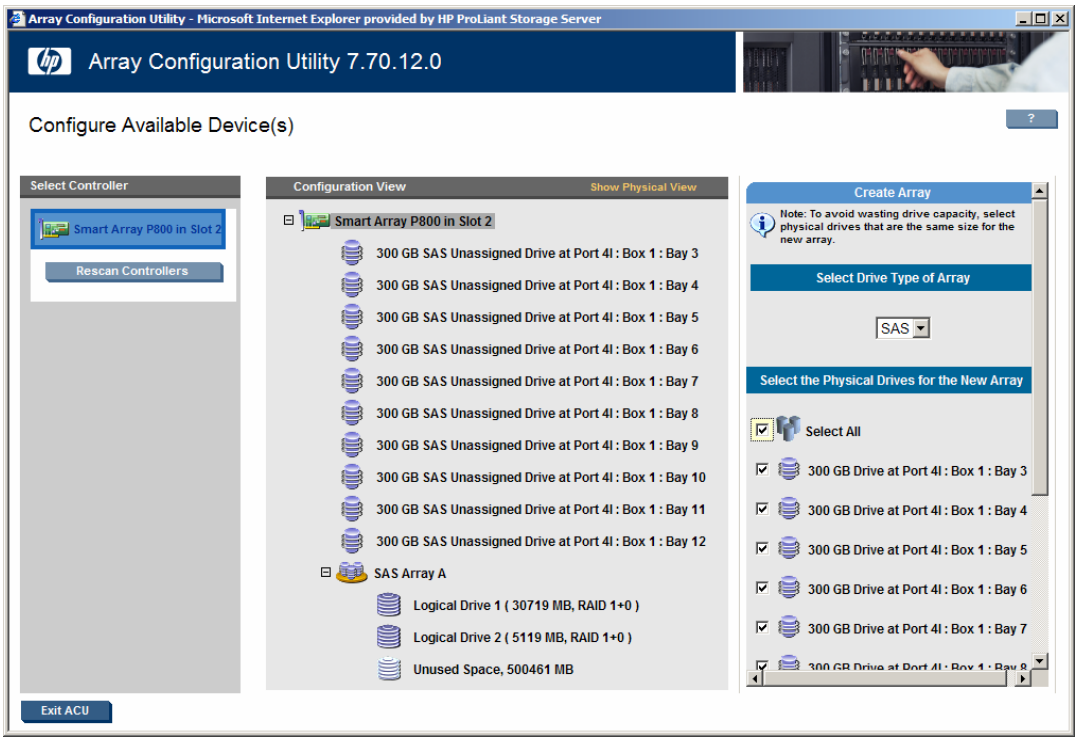

Clicking “Create Logical Drive” displays the options for Logical Drive creation, including RAID type and size. We create a logical drive of the type RAID 1+0 and use all the available space.

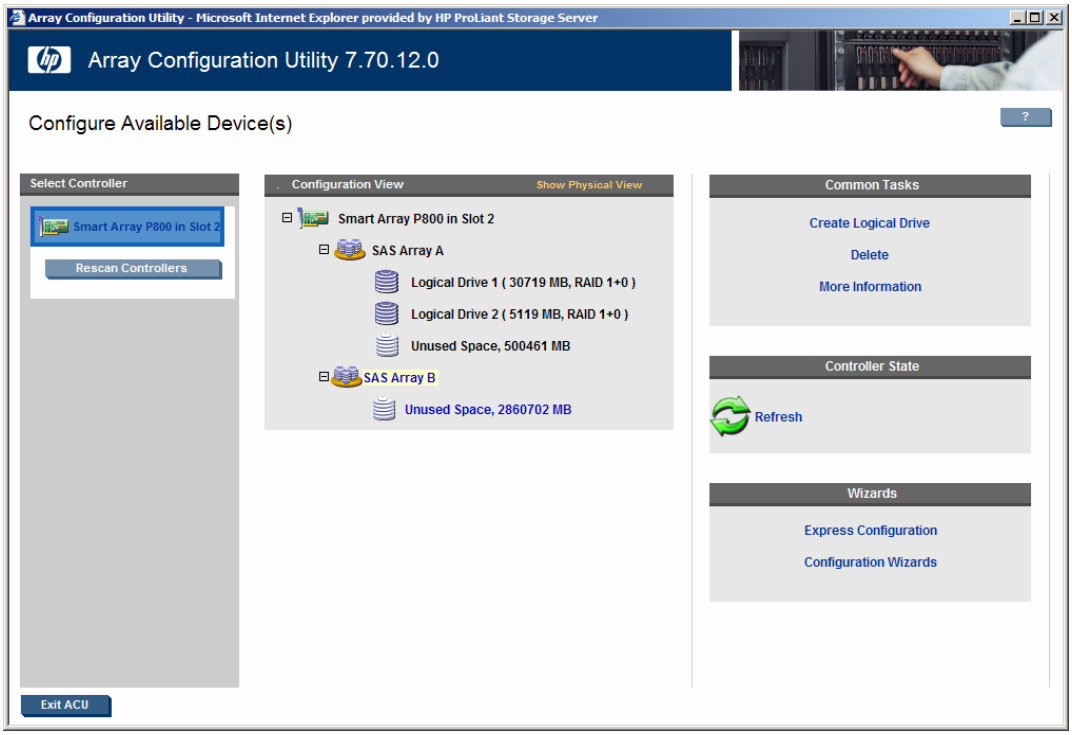

Note – It is important at this point to note that the storage solution LUN size should not be confused with the size of the iSCSI target. The iSCSI target will be configured in a

later step and is associated with the storage needed for a particular application on the host server. It is recommended that the LUN size on the storage hardware be as large as reasonably possible to allow the storage subsystem to optimize the use of the physical disks underlying the LUN that is created. In this case, as shown below, we are choosing to create one LUN at the maximum size available for this hardware. The iSCSI targets created later will fit into this one LUN, based on the needs of the host application.

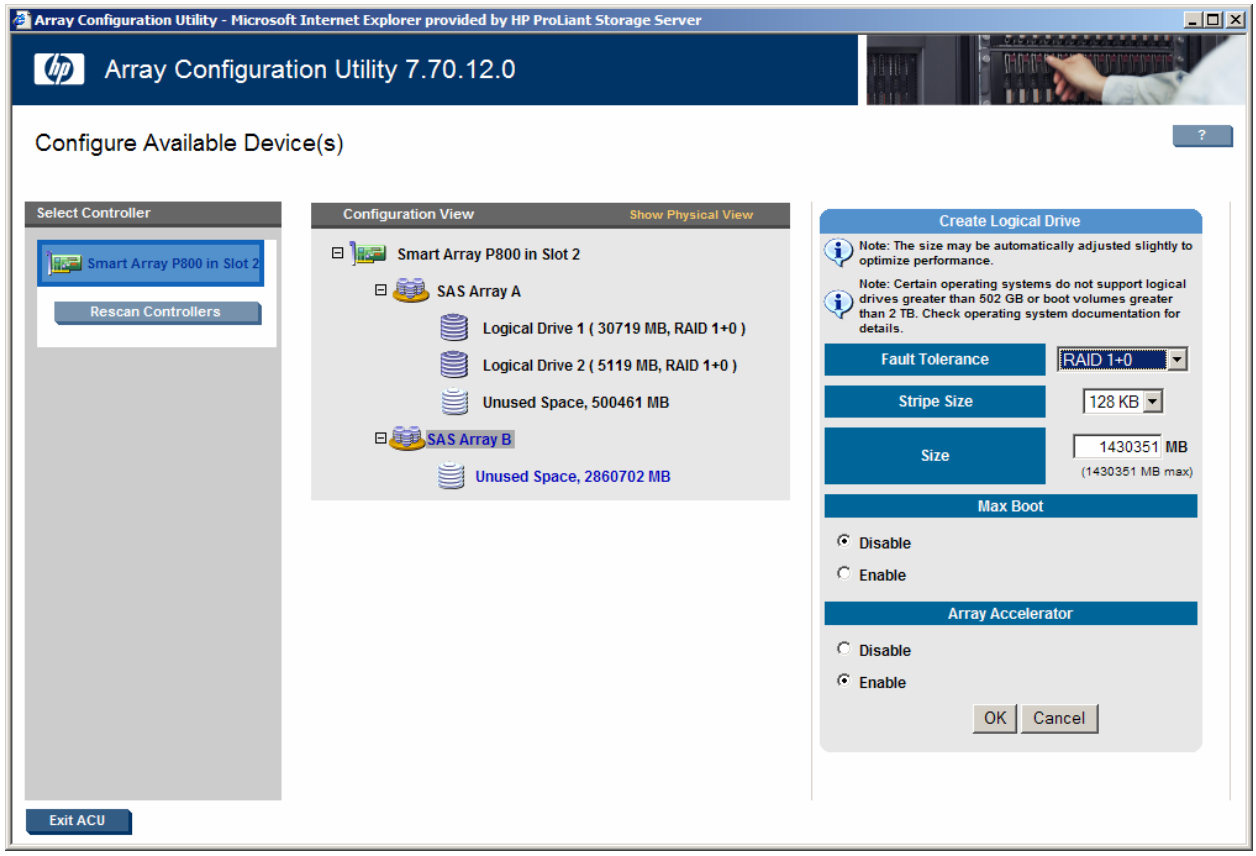

This configuration is saved.

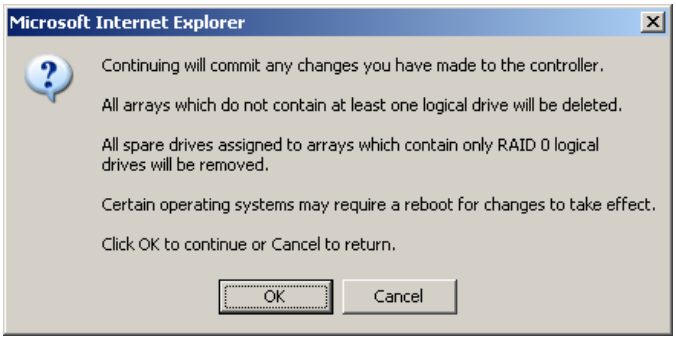

The final configuration is displayed.

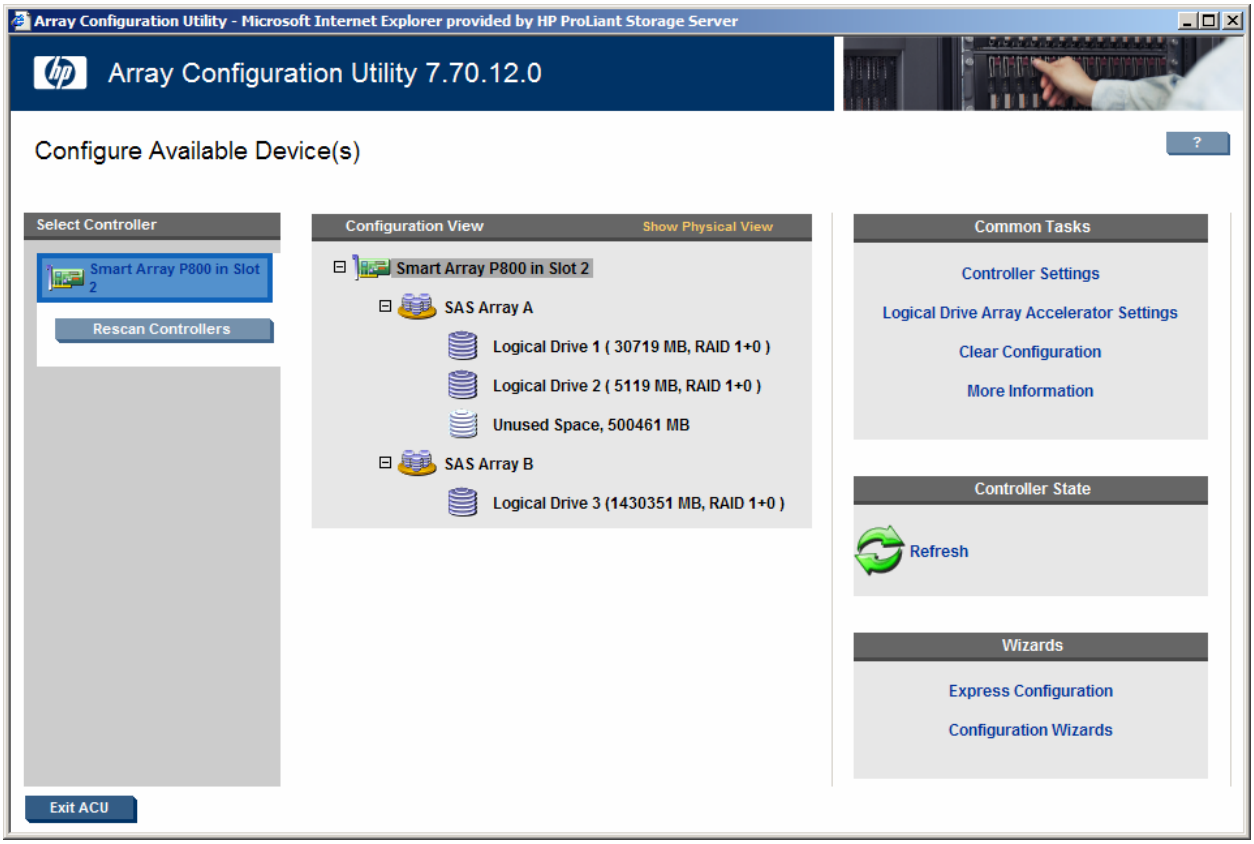

5. Make LUNs Ready for Use

Because this storage solution is running on a Microsoft Windows-based platform, the remaining steps would be familiar to a Windows administrator. These include assigning a drive letter for the internal server, providing a volume name, etc. These are prompted by using the standard Disk Management function from the management console.

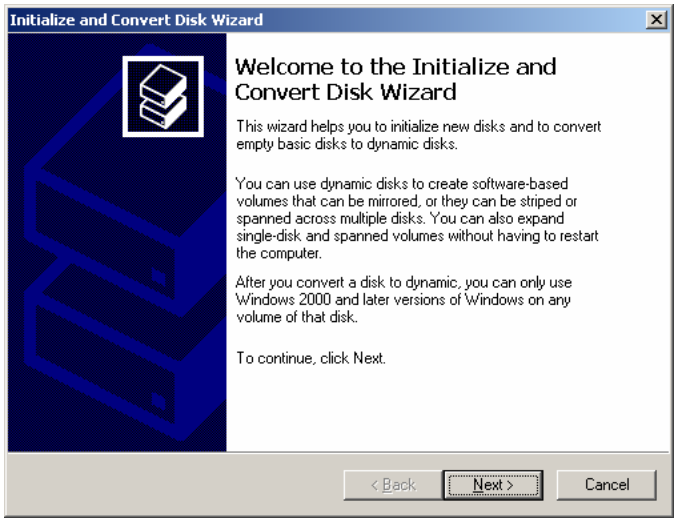

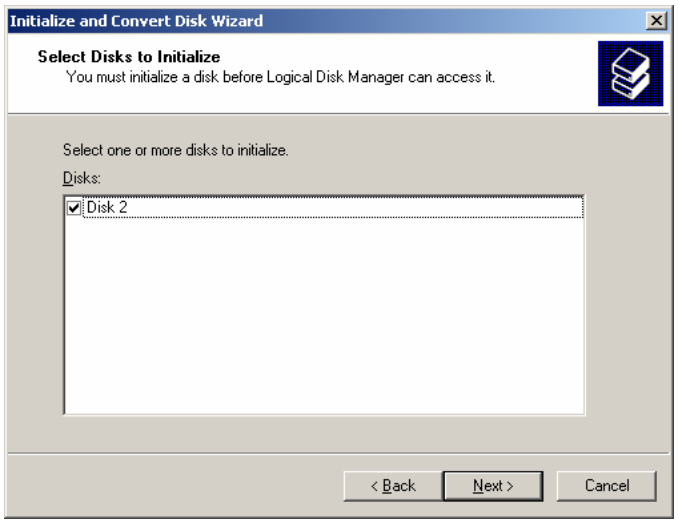
Initialize and Convert Disk Wizard
Select Disks to Initialize
You must initialize a disk before Logical Disk Manager can access it.
Select one or more disks to initialize.
Disks:
Disk 2
< Back
Next >
Cancel

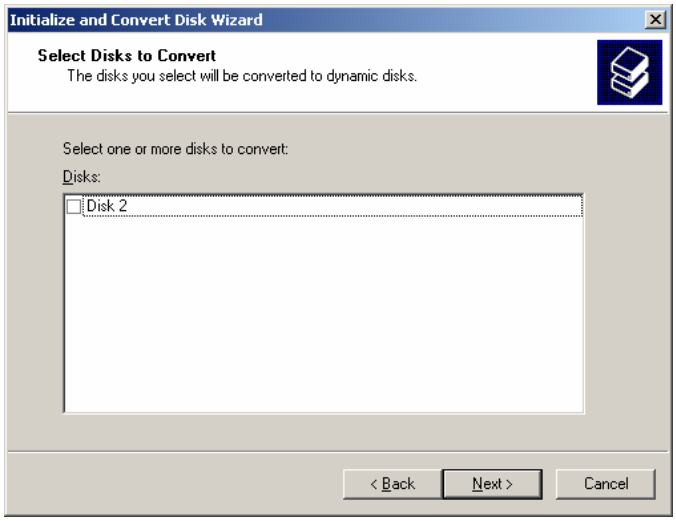
Initialize and Convert Disk Wizard
Select Disks to Convert
The disks you select will be converted to dynamic disks.
Select one or more disks to convert:
Disks:
Disk 2
< Back
Next >
Cancel

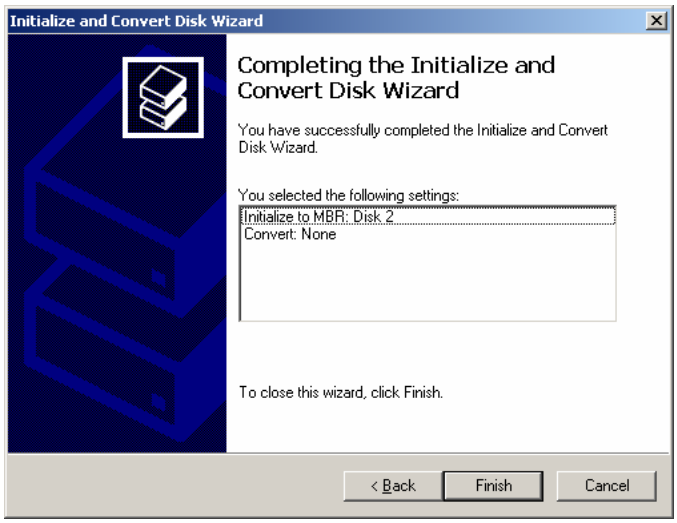
Initialize and Convert Disk Wizard
Completing the Initialize and Convert Disk Wizard
You have successfully completed the Initialize and Convert Disk Wizard.
You selected the following settings:
Initialize to MBR: Disk 2
Convert: None
To close this wizard, click Finish.
< Back
Finish
Cancel

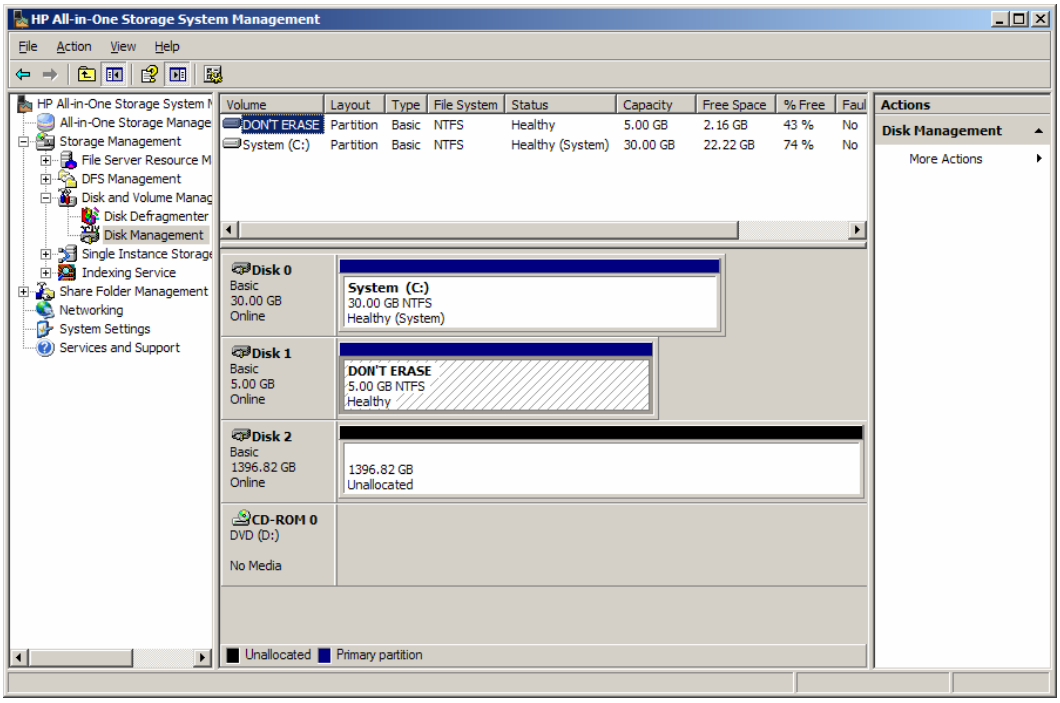

The LUN has now been created. The next step is to create a partition on the LUN and format it in the normal manner that Windows administrators would expect.

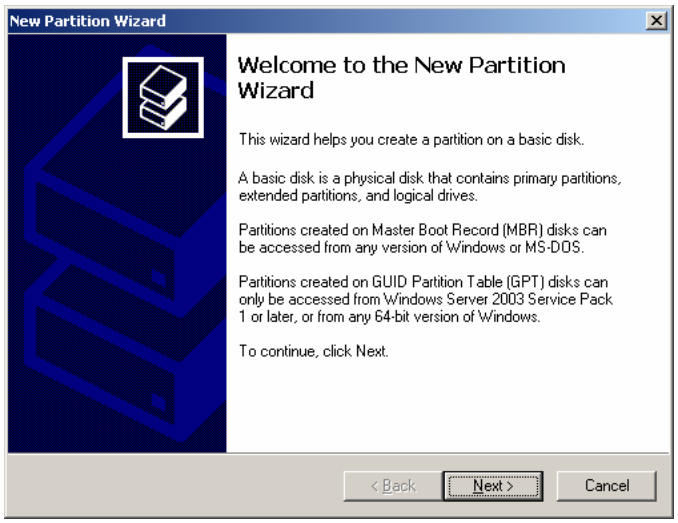

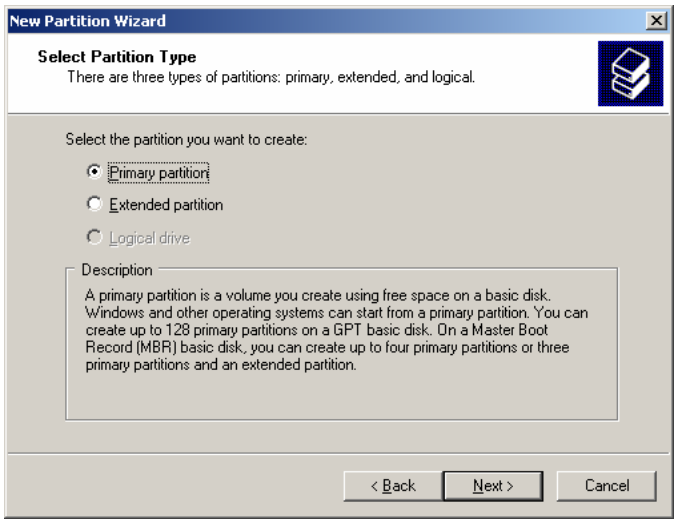

We choose the maximum available size for the new partition. We want Windows to manage the entire partition as one entity.

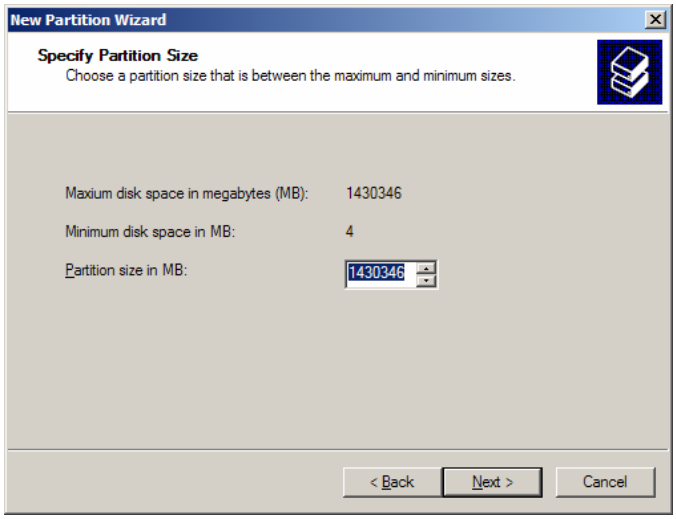

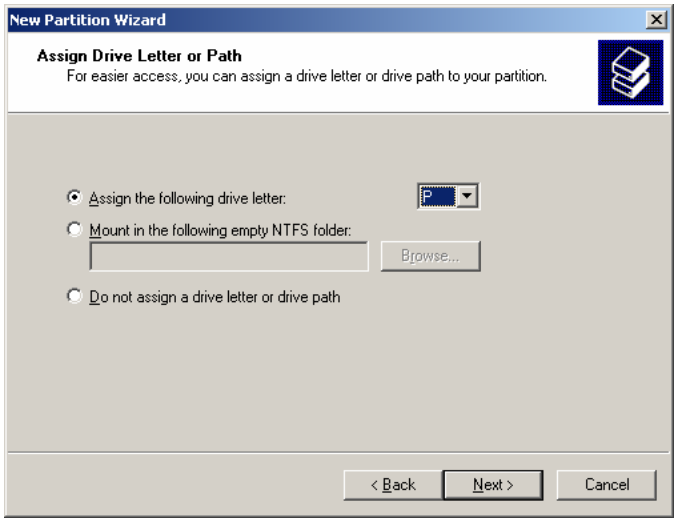

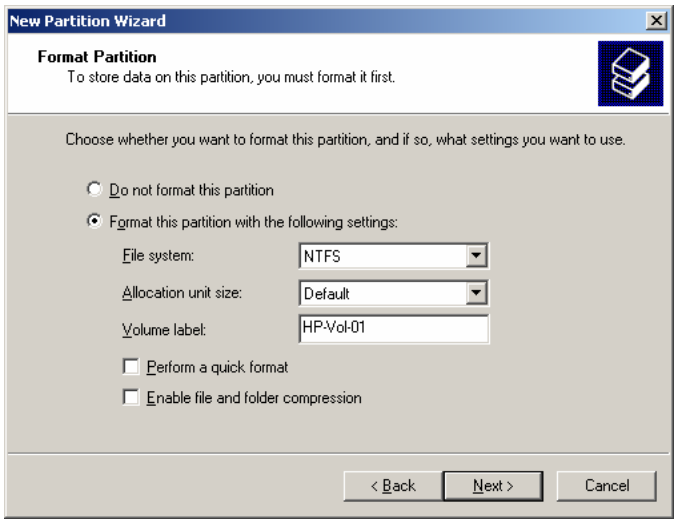
New Partition Wizard
Format Partition
To store data on this partition, you must format it first.
Choose whether you want to format this partition, and if so, what settings you want to use.
Do not format this partition
Format this partition with the following settings:
File system:
NTFS
Allocation unit size:
Default
Volume label:
HP-Vol-01
Perform a quick format
Enable file and folder compression
< Back
Next >
Cancel

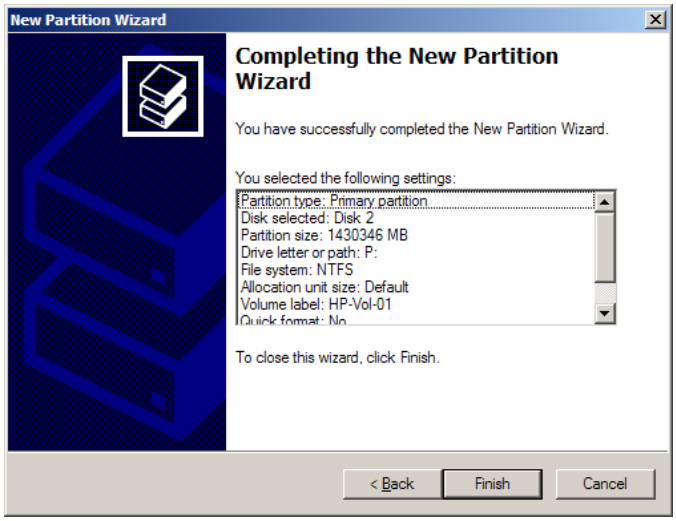
New Partition Wizard
Completing the New Partition Wizard
You have successfully completed the New Partition Wizard.
You selected the following settings:
Partition type: Primary partition
Disk selected: Disk 2
Partition size: 1430346 MB
Drive letter or path: P:
File system: NTFS
Allocation unit size: Default
Volume label: HP-Vol-01
To close this wizard, click Finish.
< Back
Finish
Cancel

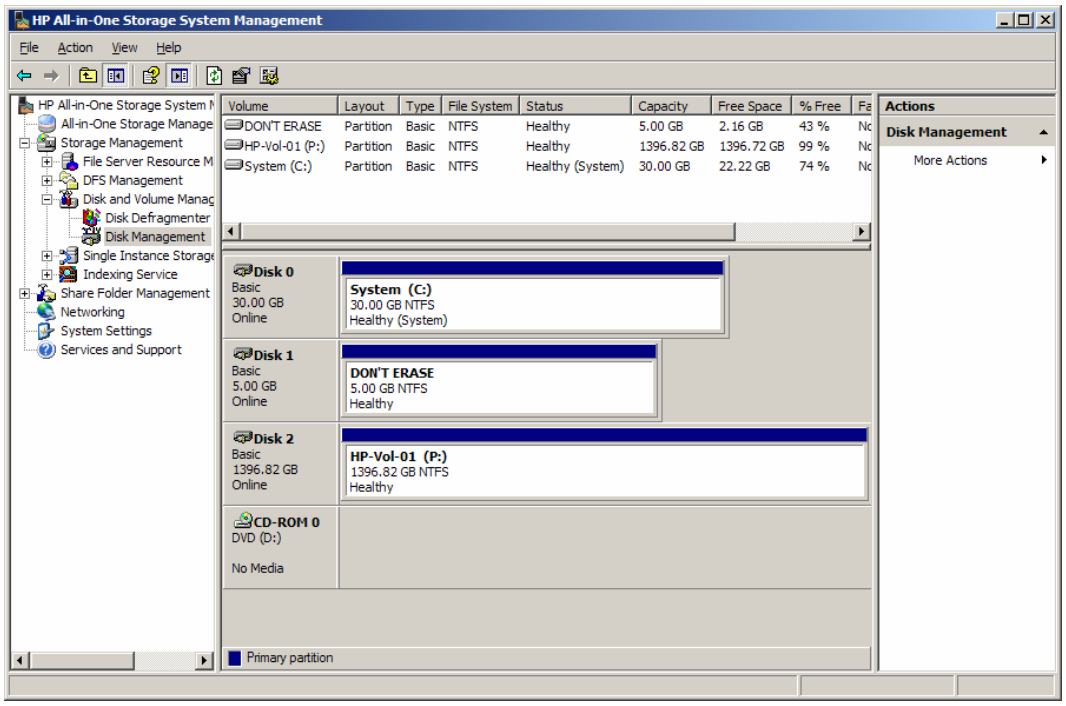

The LUN is now ready for use. The next step will create iSCSI targets and associate them with this newly-created LUN.

6. Create iSCSI Targets

Clicking on the Microsoft iSCSI target function in the Storage Management section of the management console begins the iSCSI target creation process.

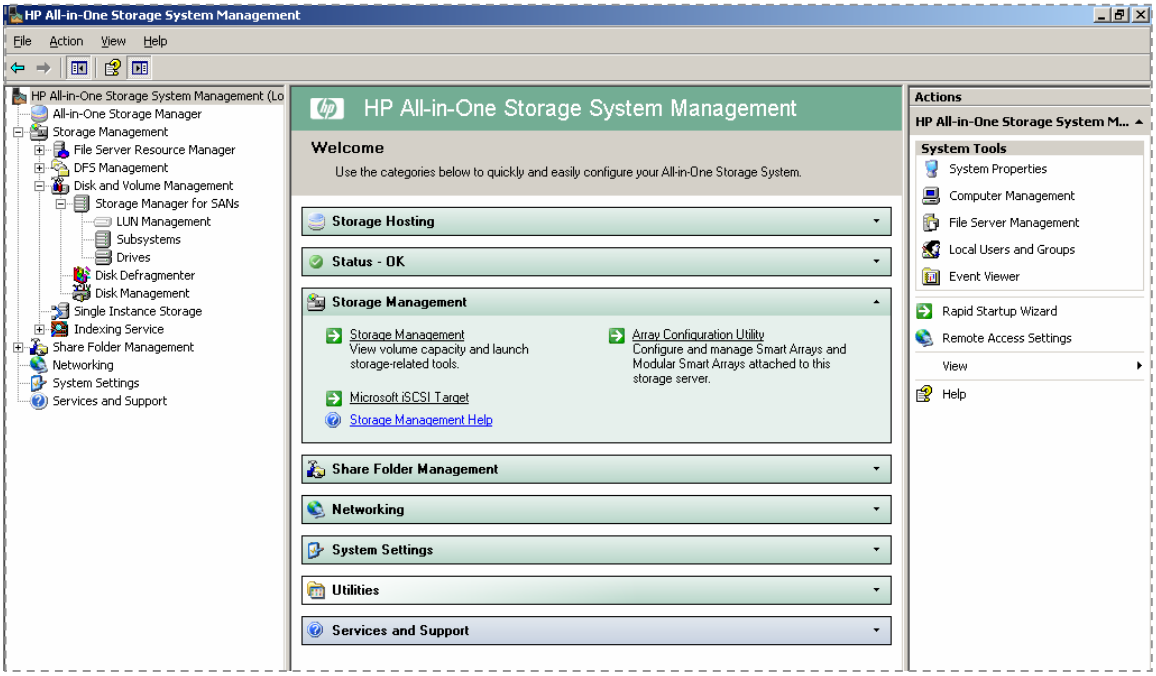

The iSCSI target window opens, and the iSCSI target creation wizard is started by using the right mouse click button.

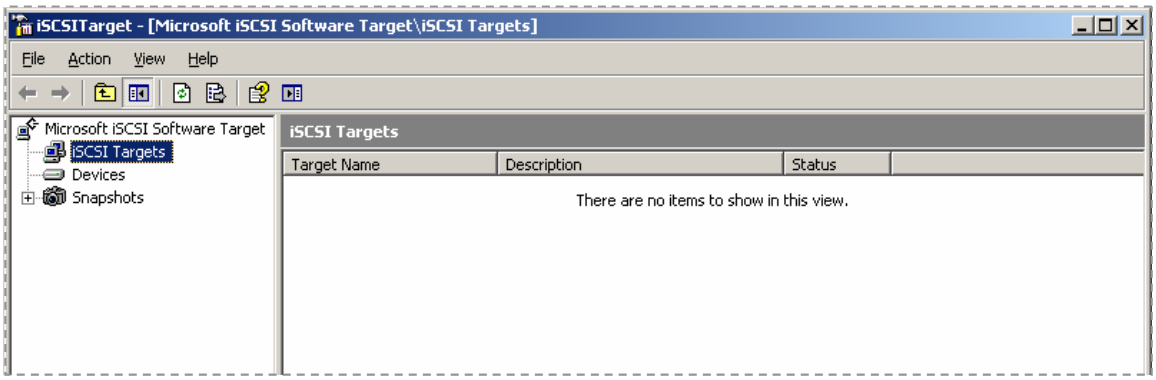

We will create a total of four targets, two associated with host number one and two associated with host number two. The individual steps to create the first target are shown below.

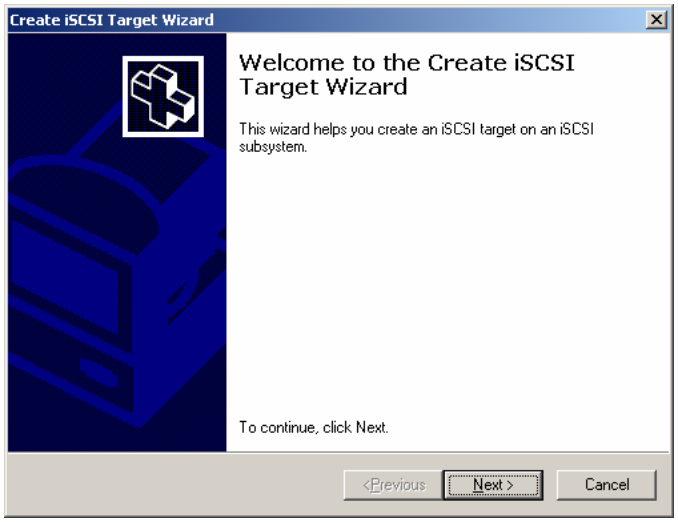

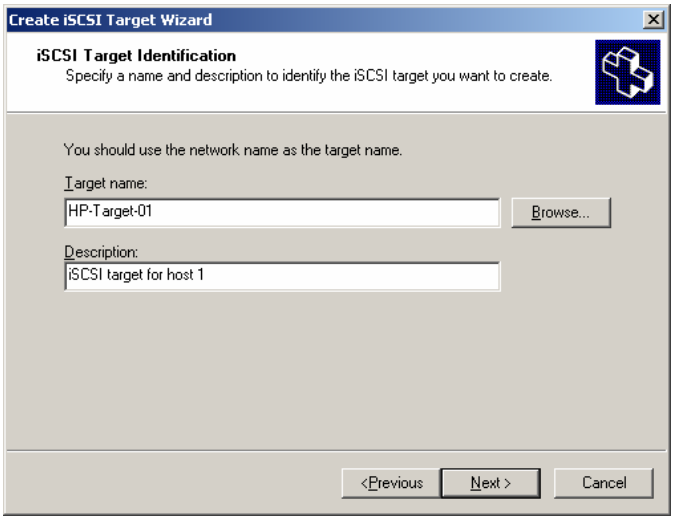

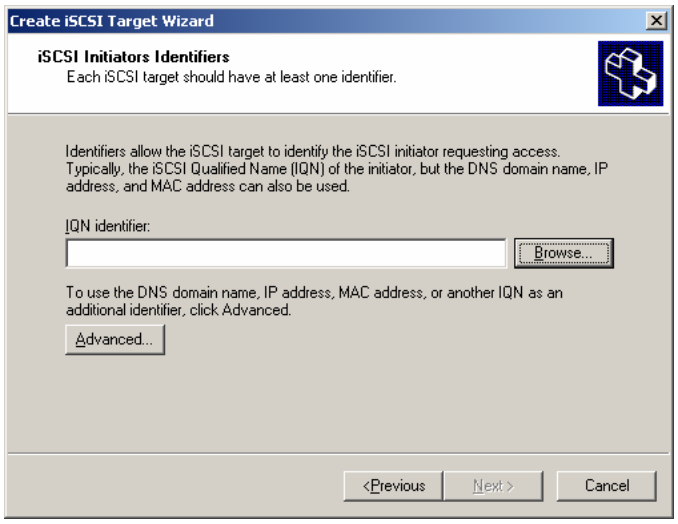

Because there have been no previous iSCSI initiators assigned to this target, the initiators must be specifically identified.

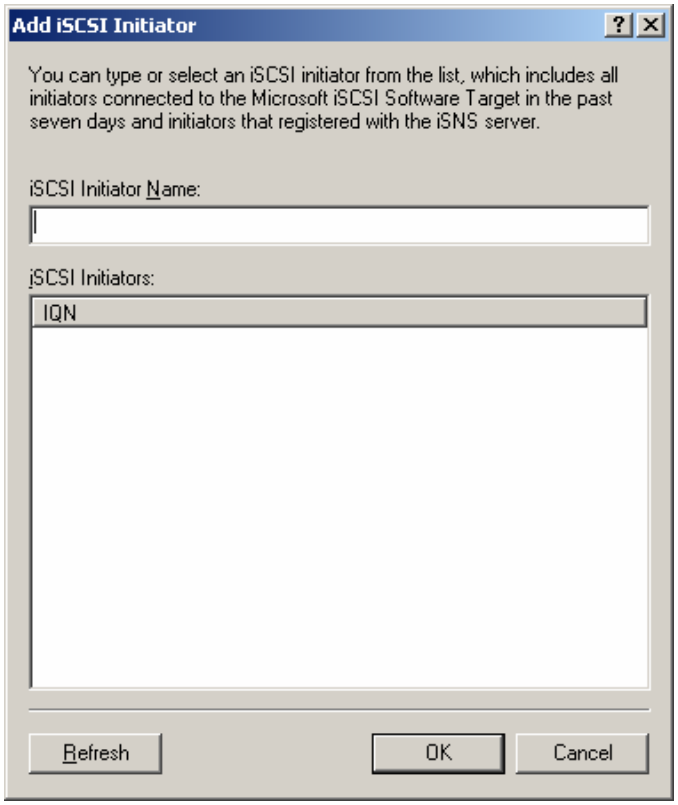

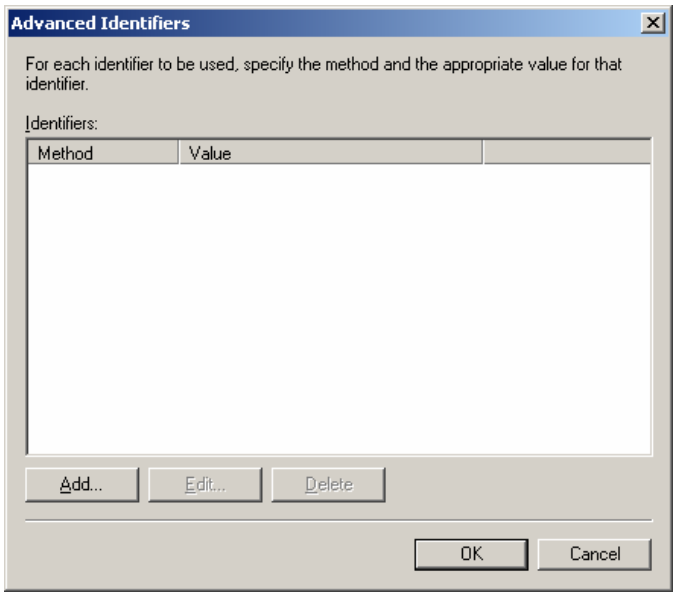

The initiators will be identified by their IP address.

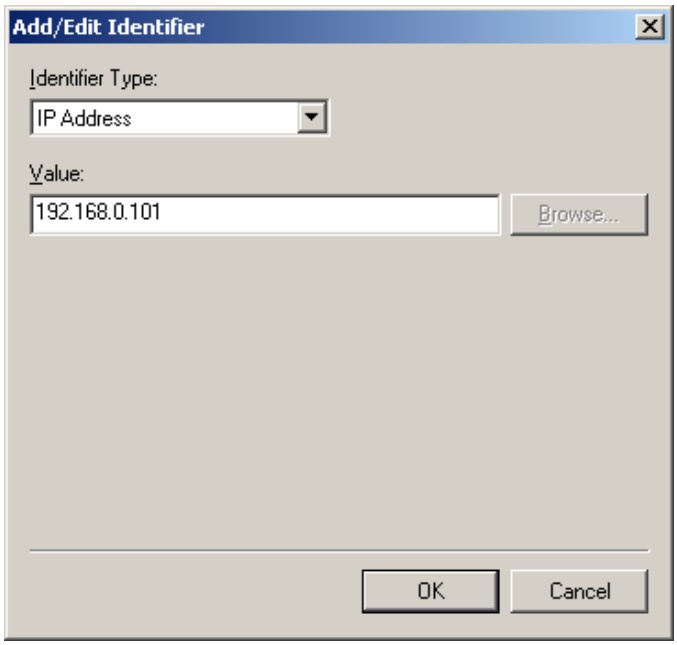

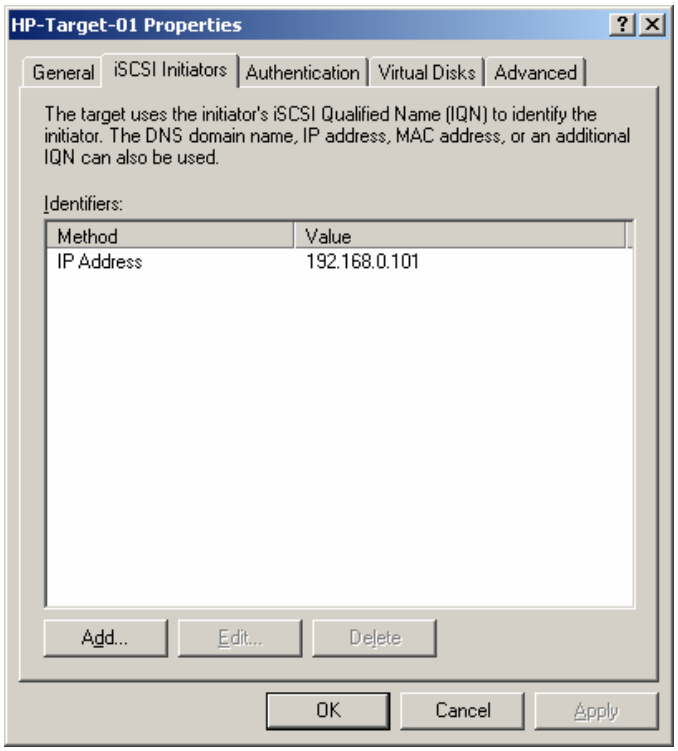
HP-Target-01 Properties
General
iSCSI Initiators
Authentication
Virtual Disks
Advanced
The target uses the initiator's iSCSI Qualified Name (IQN) to identify the initiator. The DNS domain name, IP address, MAC address, or an additional IQN can also be used.
Identifiers:
Method
Value
IP Address
192.168.0.101
Add...
Edit...
Delete
OK
Cancel
Apply

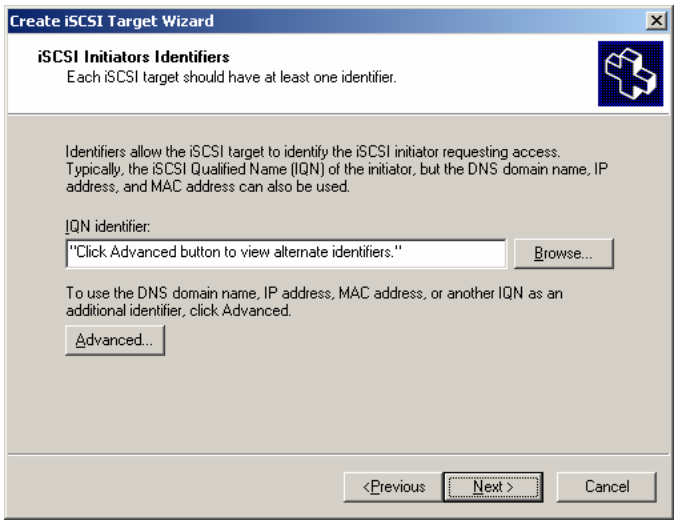
Create iSCSI Target Wizard
iSCSI Initiators Identifiers
Each iSCSI target should have at least one identifier.
Identifiers allow the iSCSI target to identify the iSCSI initiator requesting access. Typically, the iSCSI Qualified Name (IQN) of the initiator, but the DNS domain name, IP address, and MAC address can also be used.
IQN identifier:
"Click Advanced button to view alternate identifiers."
Browse...
To use the DNS domain name, IP address, MAC address, or another IQN as an additional identifier, click Advanced.
Advanced...
<Previous
Next>
Cancel

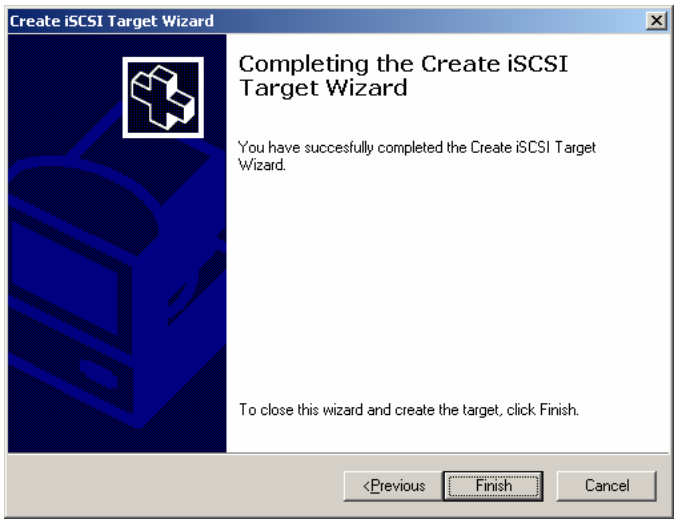

The target now appears in the target list of the iSCSI target management window.

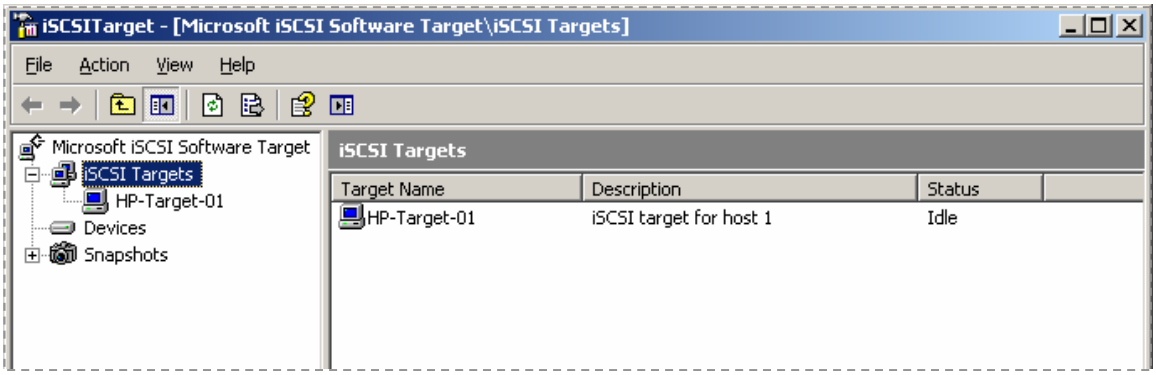

The management console also shows the underlying devices available for the iSCSI targets. The storage that will be used by the iSCSI initiators (application hosts) will be defined in a later step when the virtual disks are created.

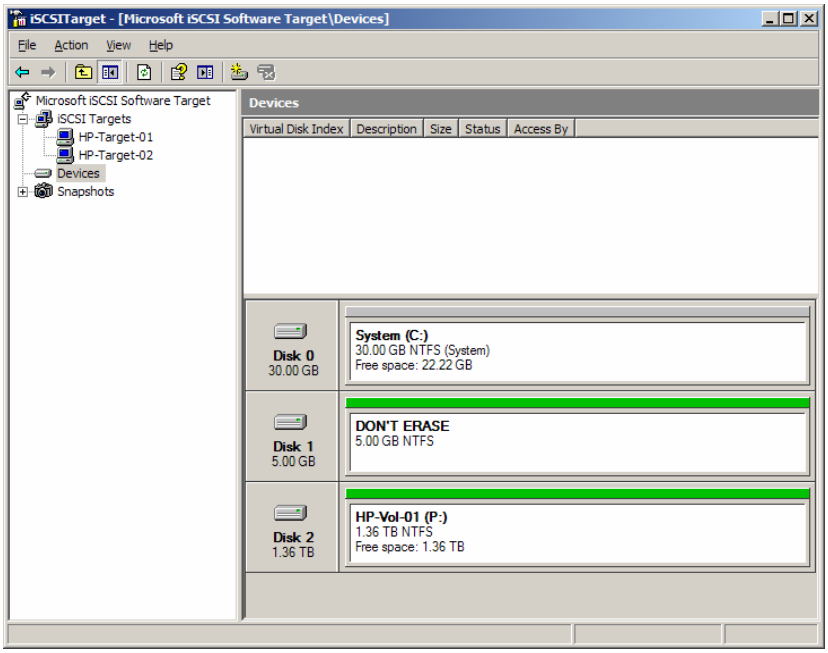

7. Create Multi-path I/O for iSCSI Targets (optional)

For a Microsoft-based target solution, multiple paths to the target device are created by providing multiple iSCSI initiator identifiers for the same target. In the example above, we add a second and third IP address that is to be associated with the same iSCSI initiator, so that there are three IP addresses that can access the target. These three addresses are associated with three Ethernet ports on the same host server.

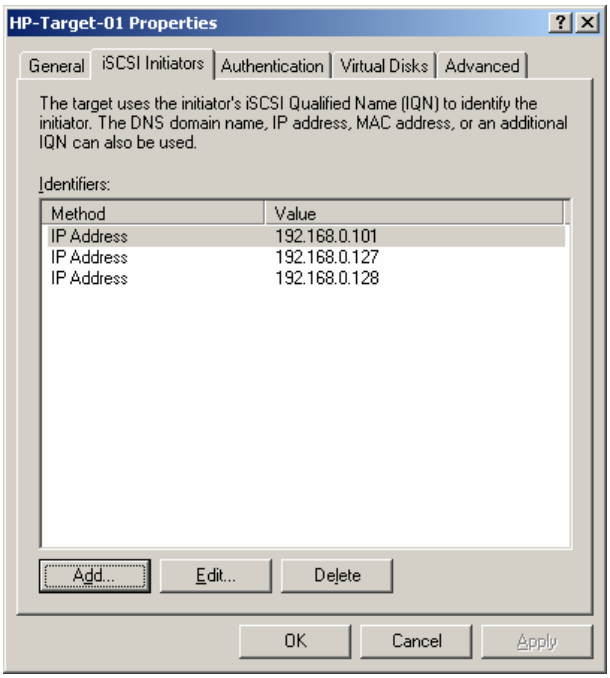

8. Configure Security for iSCSI Targets (optional)

If security is desired for the iSCSI communication between the initiator and the target, the CHAP secret can be entered in the Authentication tab.

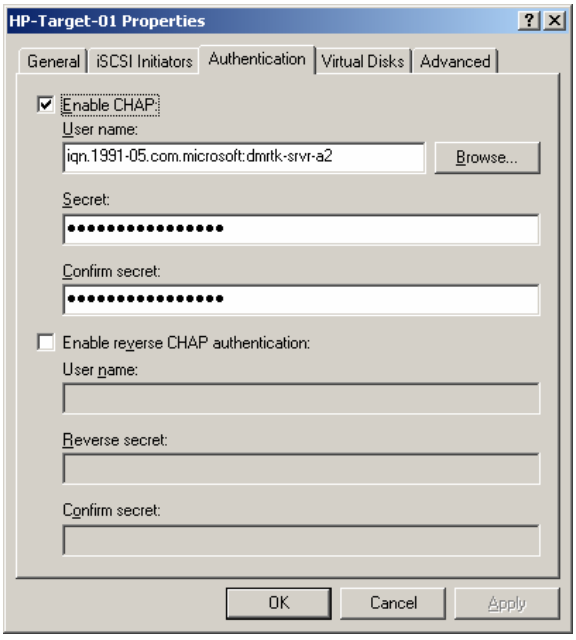

9. Make iSCSI Targets Ready for Use for iSCSI Initiators

Virtual disks need to be created on the iSCSI targets for Microsoft-based iSCSI target solutions. These virtual disks represent the storage volumes that the iSCSI initiators will use. The maximum capacity represented by all the virtual disks on a given iSCSI target on a Microsoft-based iSCSI target solution is two terabytes (2 TB) per target. In this example, we create a 100 GB and 200 GB virtual disk for each of the two iSCSI targets. These virtual disks will be viewed as volumes by the iSCSI initiators over the TCP/IP network.

By right-clicking on the target name, the “Create Virtual Disk” wizard is launched.

The virtual disk is created on the internal disk volume that is available to the iSCSI target. In this case, it is the “P” volume.

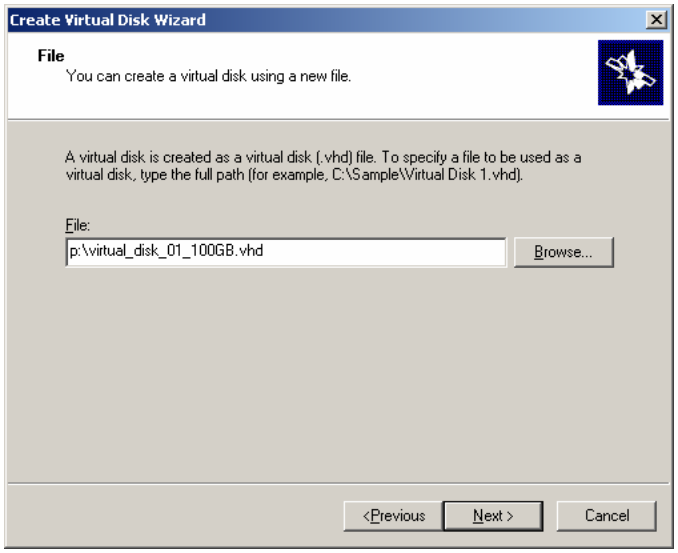

The size of the virtual disk depends on the needs of the application on the host server. For this virtual disk, we choose a size of 100 GB from the available 1.34 TB on this volume.

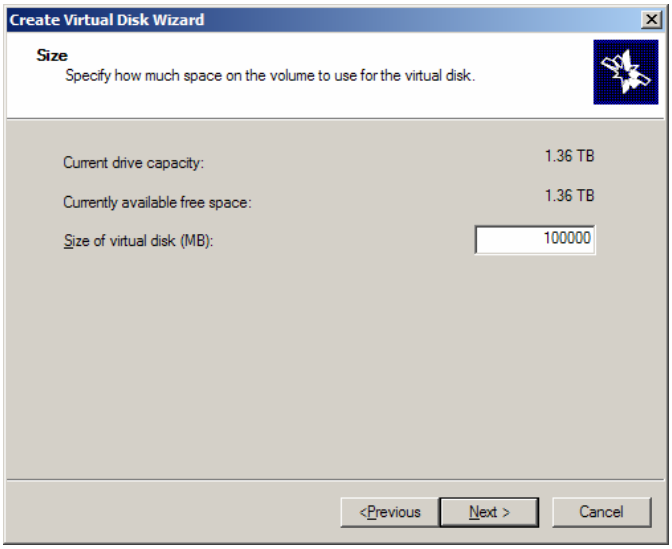

A description, although optional, is useful for better management.

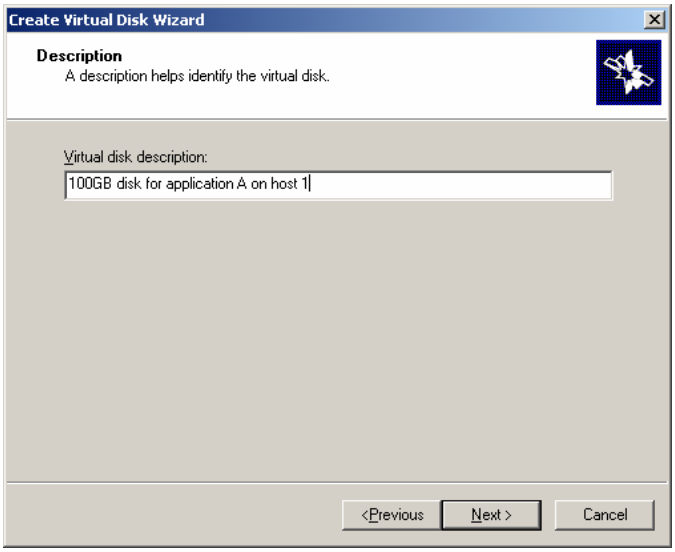

This virtual disk must be associated with an iSCSI target in order for the application host to use this storage as an iSCSI storage volume.

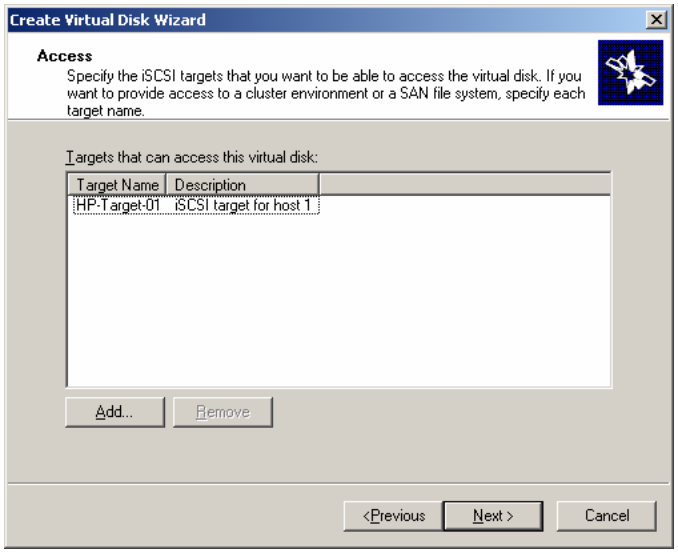

This completes the virtual disk configuration.

We repeat this process to create three more virtual disks. After configuring the virtual disks, the management console shows the virtual disks associated with each iSCSI target.

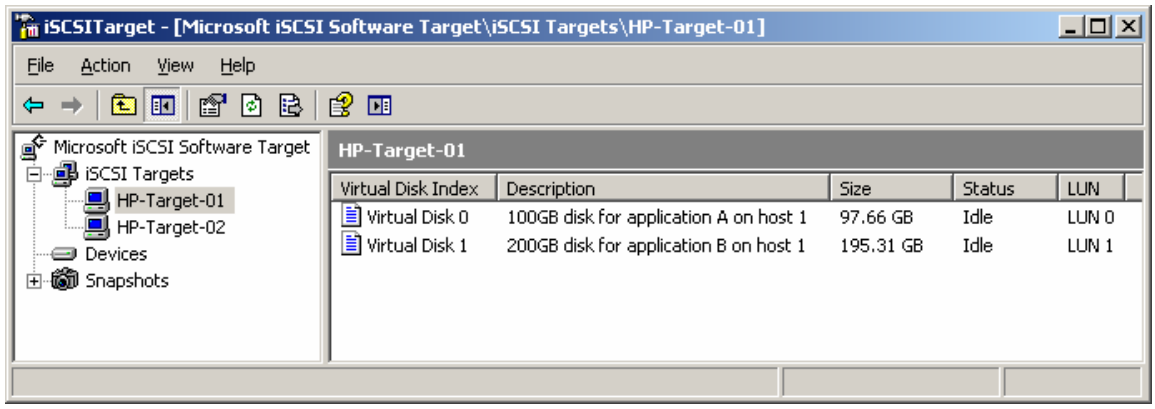
iSCSITarget - [Microsoft iSCSI Software Target\iSCSI Targets\HP-Target-01]
File Action View Help
Microsoft iSCSI Software Target
iSCSI Targets
HP-Target-01
HP-Target-02
Devices
Snapshots
HP-Target-01
Virtual Disk Index | Description | Size | Status | LUN
Virtual Disk 0 | 100GB disk for application A on host 1 | 97.66 GB | Idle | LUN 0
Virtual Disk 1 | 200GB disk for application B on host 1 | 195.31 GB | Idle | LUN 1

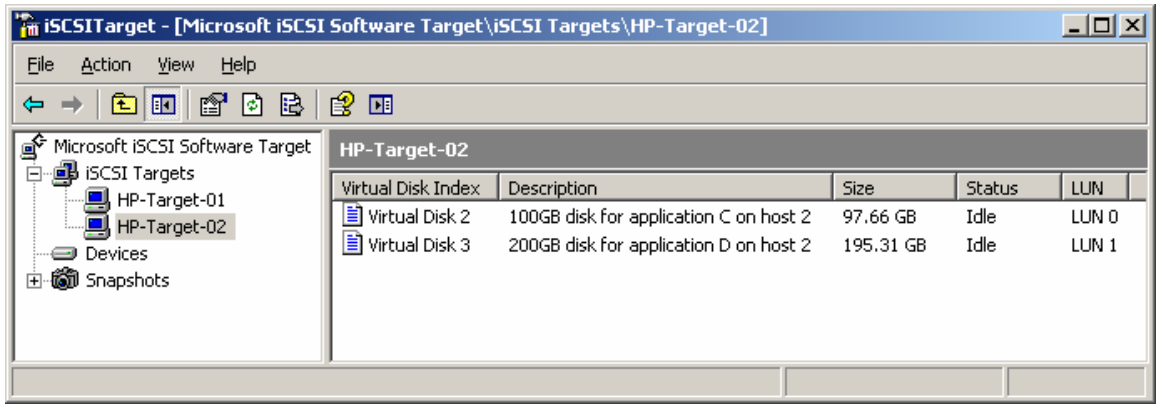
iSCSITarget - [Microsoft iSCSI Software Target\iSCSI Targets\HP-Target-02]
File Action View Help
Microsoft iSCSI Software Target
iSCSI Targets
HP-Target-01
HP-Target-02
Devices
Snapshots
HP-Target-02
Virtual Disk Index | Description | Size | Status | LUN
Virtual Disk 2 | 100GB disk for application C on host 2 | 97.66 GB | Idle | LUN 0
Virtual Disk 3 | 200GB disk for application D on host 2 | 195.31 GB | Idle | LUN 1

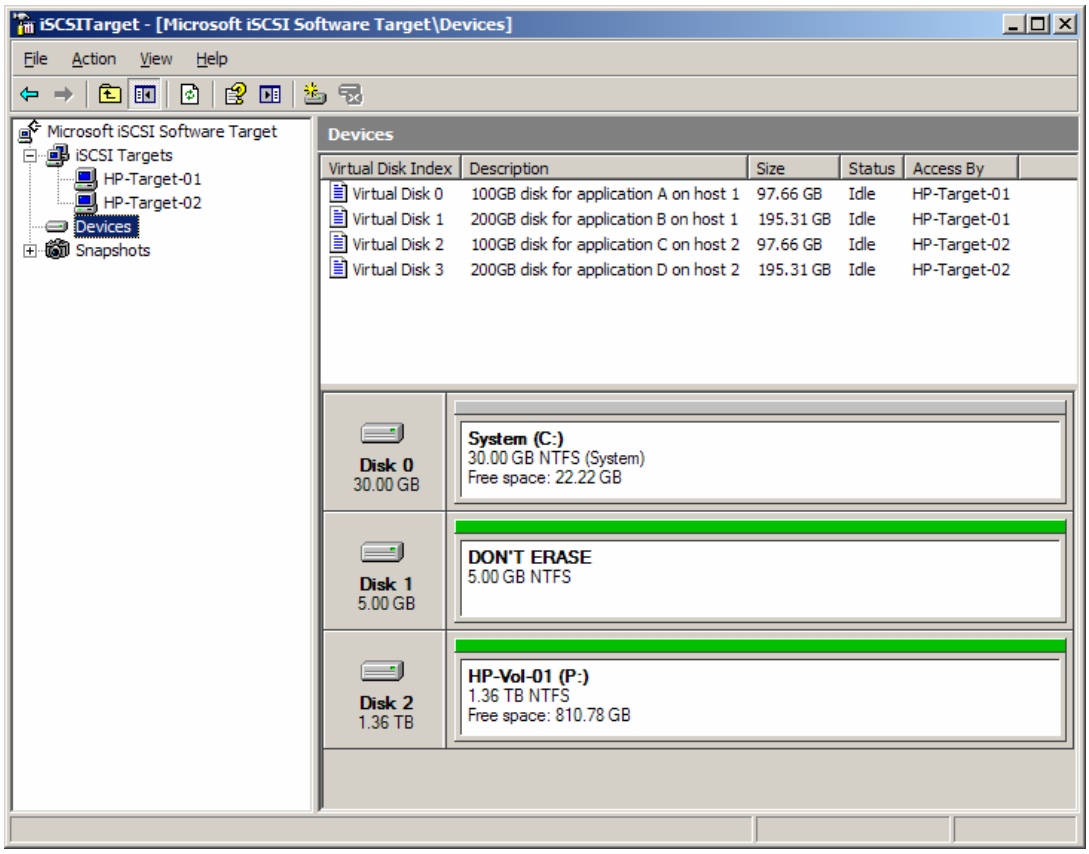

The target side configuration is now complete.

Initiator Configuration Steps

Configure Multi-path I/O from Application Host

To configure multi-path iSCSI I/O for the initiator that uses the HP All-in-One 1200 iSCSI targets, follow the directions for Microsoft Multi-path I/O from the Deployment section of this document above.

Basic Performance Results

The following performance data is not intended to be viewed as a comprehensive performance benchmark, but to provide a general performance overview for the HP StorageWorks All-in-One1200 Storage System.

Selected performance results are shown below, using a standard server-class network adapter, with receive-side scaling on the host. This configuration used two paths from one host, two I/O workers, simultaneously accessing two target volumes and a queue depth of 10. Each target volume used a dedicated path, with no load-balancing across the paths.

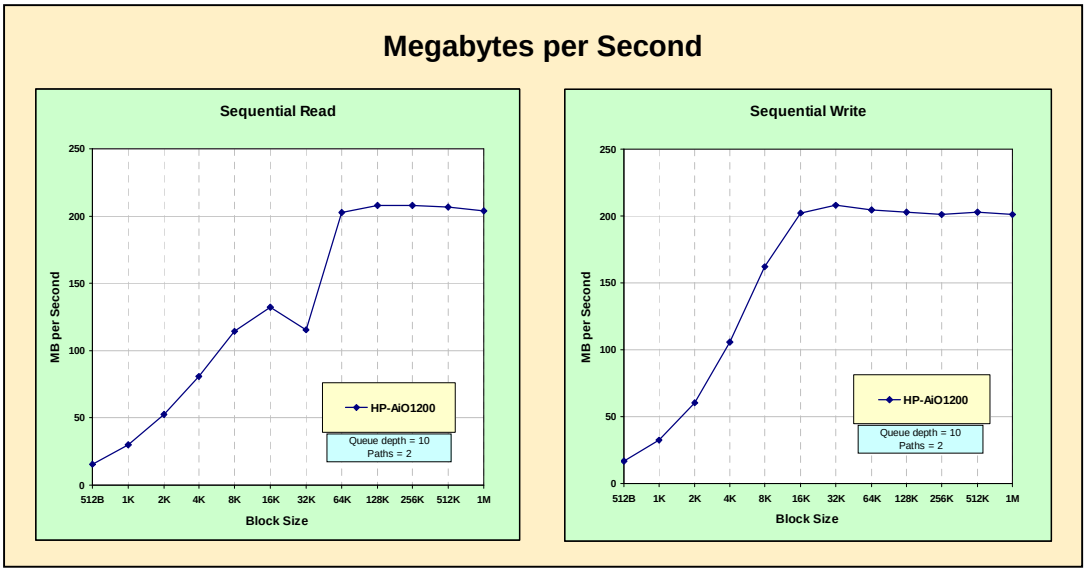

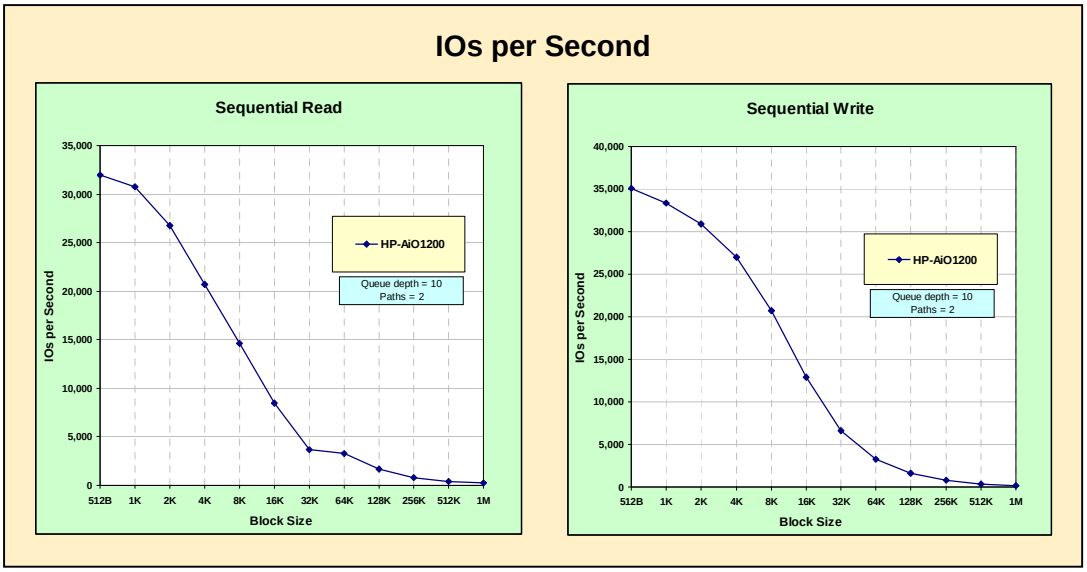

LeftHand Networks® SAN/iQ®

The LeftHand Networks SAN/iQ storage system is an iSCSI target solution that includes three (3) HP ProLiant DL320s systems and the LeftHand Networks SAN/iQ software. It includes 10K or 15K RPM SAS disk drives totaling up to 3.6 TB of raw capacity per module, configured as RAID 10. In this case, the three ProLiant DL320s servers are clustered together to create a "virtual storage array" consisting of 5.4 TB of total usable storage. In addition to RAID10 at the disk level, LeftHand offers "network RAID" to provide an additional layer of protection for individual LUNs which guards against network or any other hardware failure.

SAN/iQ includes many additional management features as part of the basic SAN offering. These features include snapshots, volume branching, thin provisioning, offsite DR snapshots, iSCSI load balancing, block-level load balancing, and automated capacity management. Customers can expand the storage cluster at any time by adding additional units to the cluster. This not only increases capacity, but SAN/iQ automatically re-load balances all the existing LUNs across the new configuration, increasing the performance of the SAN as well.

Another capability is LeftHand's "multi-site SAN" in which customers physically locate half their cluster in one location and the other half at a different location, such as another building or floor. In this case, the SAN is now "fault-tolerant" in that a location disaster will not interrupt service. This capability is included in the base offering, and does not require any additional administration to set up and manage.

Target Configuration Steps

1. Configure Network Settings for iSCSI Target Device

To install the LeftHand Networks SAN/iQ from factory settings, a computer must be connected via the supplied serial cable to the LeftHand NSM. The first Ethernet port must be given an IP address. Later, a virtual IP address will be assigned using the management console that will be the address that the clients use to access the target volumes and to manage the storage cluster.

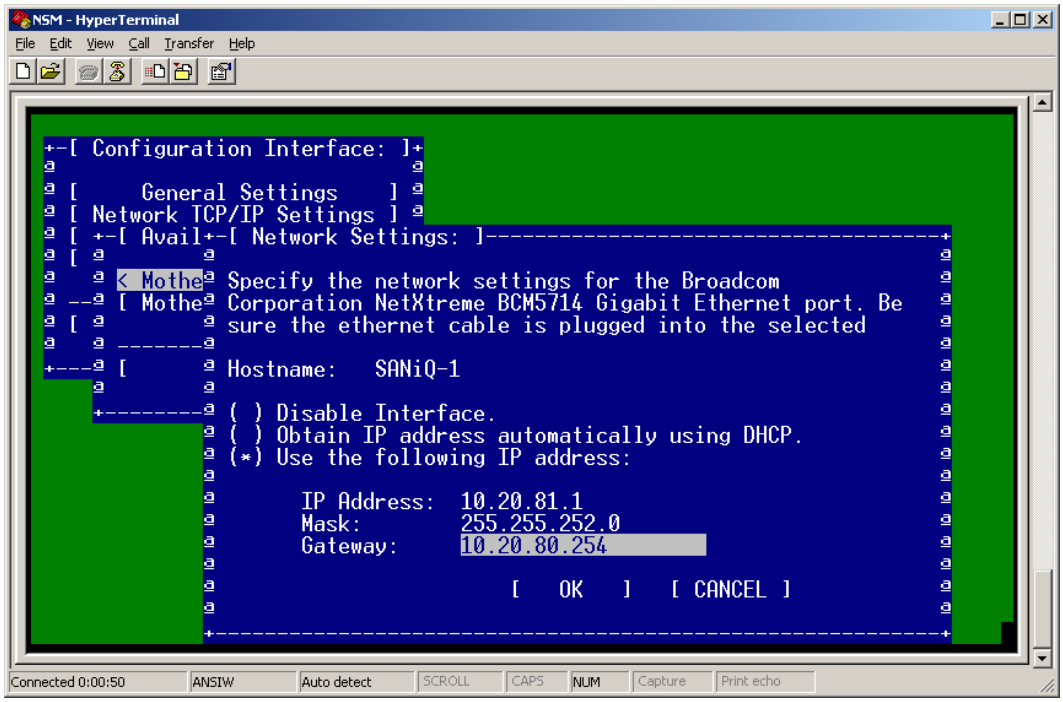

2. Launch Management Console

Launching the Management Console begins the discovery process and displays the NSMs. The three basic steps to configure the system including creation and assignment of all the targets are listed on the main screen, each driven by a wizard.

The LeftHand Networks SAN/iQ solution uses the concept of management groups to organize its storage clusters. The “Management Group, Clusters and Volumes” wizard steps the administrator through the initial management configuration and creation of the first volume.

The wizard asks few questions to complete the initial management configuration, including the name of the management group, virtual IP address to be used for the cluster and the first volume information.

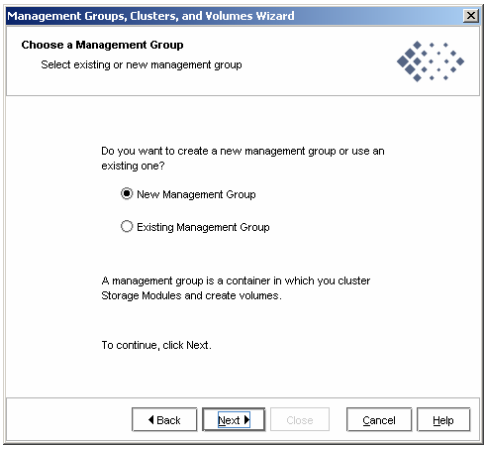

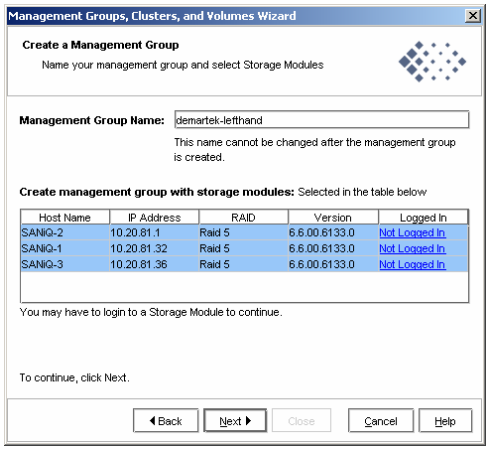

The storage cluster is given a virtual IP address, that will be used for all access to the volumes assigned to the cluster. The LeftHand Networks solution presents the virtual IP address to the clients and manages all fail-over and load balancing functions behind this virtual address.

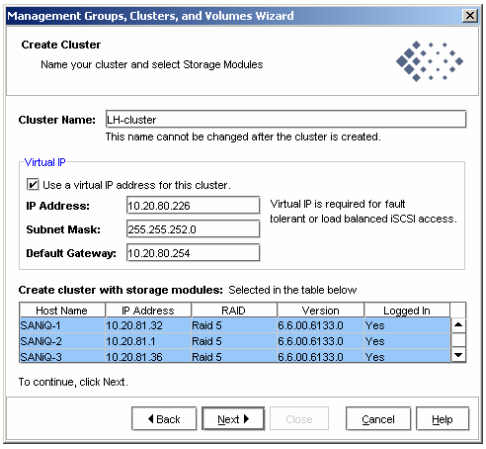

After the wizard has the virtual IP address, it prompts for the information to create the first volume, including volume name, replication features and capacity. In this case, 2-way replication is selected, which tells SAN/iQ to provide an additional layer of data protection for this volume.

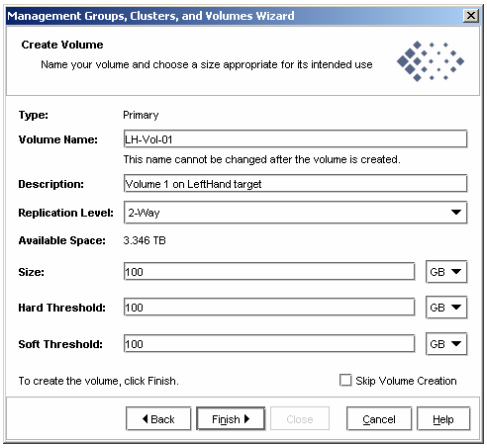

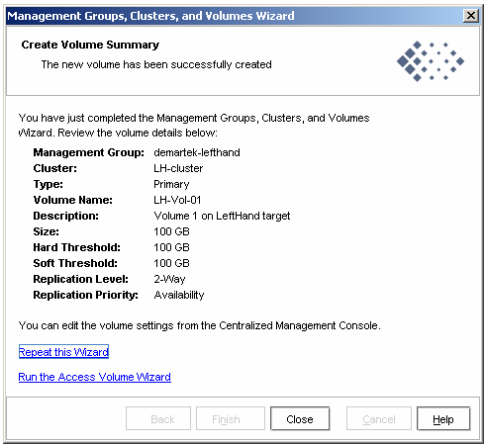

This wizard can be repeated to created additional volumes. Note that the “Access Volume Wizard” can be run to complete all the remaining target management steps.

3. Create LUNs on Disk Array

The LUNs are created using the “Management Group, Clusters and Volumes” wizard as described above in step 2.

4. Make LUNs Ready for Use

To make the volumes ready to use, they must be assigned to a host and appropriate security applied. The “Access Volume Wizard” is run to complete this process. This step can be run directly at the conclusion of the previous wizard from step 2 above.

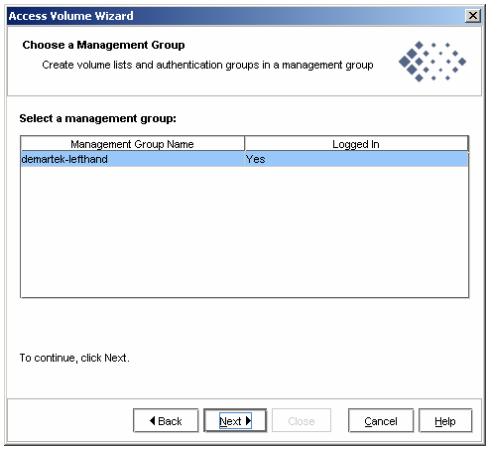

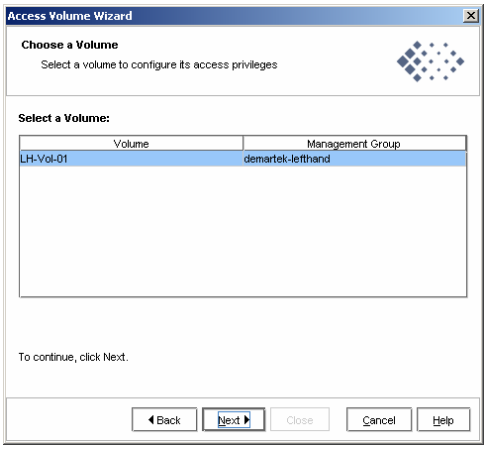

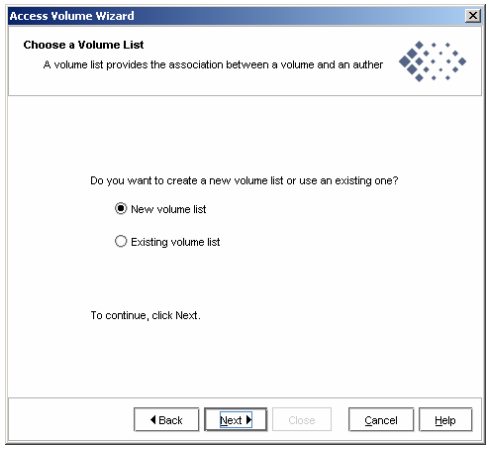

A volume list provides the association between the volume and a host. After the first volume has been associated with a host, other volumes can be added to the volume list, and these volumes are automatically associated with the same host.

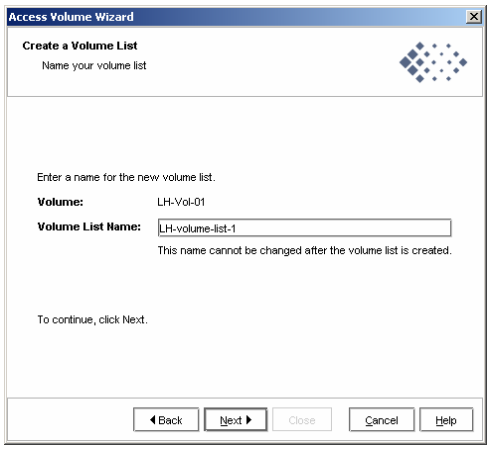

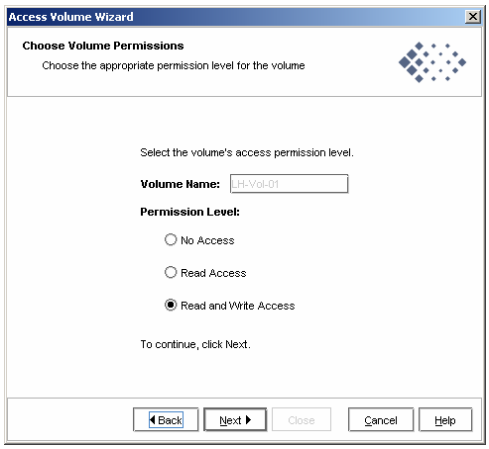

The authentication group provides the information about the hosts that will access the volumes in this volume list. In this case, the name of the host is used as the name of the authentication group.

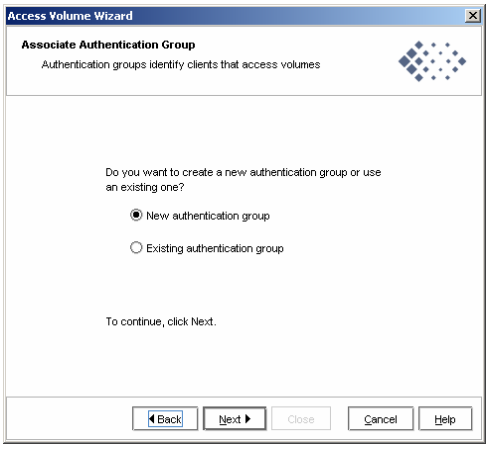

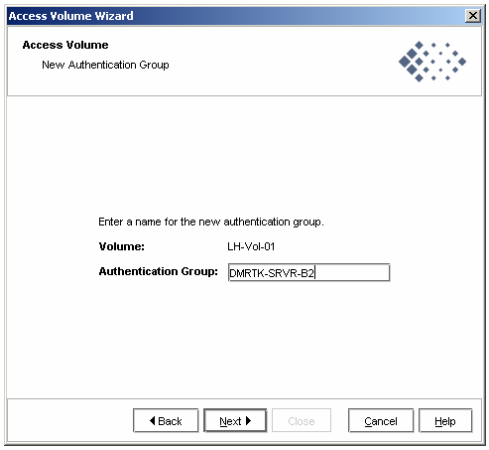

The iSCSI initiator name of the host is provided, along with any desired load balancing.

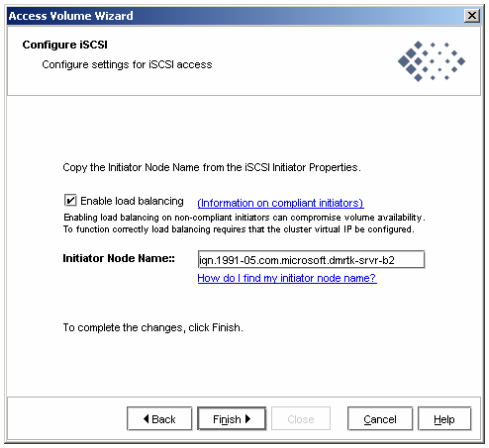

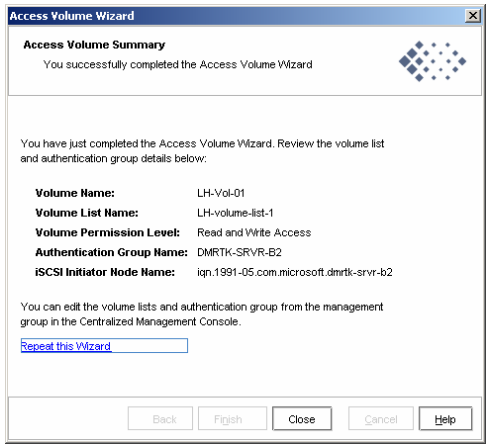

This wizard can be repeated for additional volume lists or host information.

5. Create iSCSI Targets

The targets are created in step 2 above.

6. Create Multi-path I/O for iSCSI Targets (optional)

Multi-path I/O is automatically performed by the storage cluster on the target side. Initiator side MPIO is configured below in the Initiator Configuration Steps.

7. Configure Security for iSCSI Targets (optional)

Additional security, such as CHAP, can be configured by editing the authentication group for the volume.

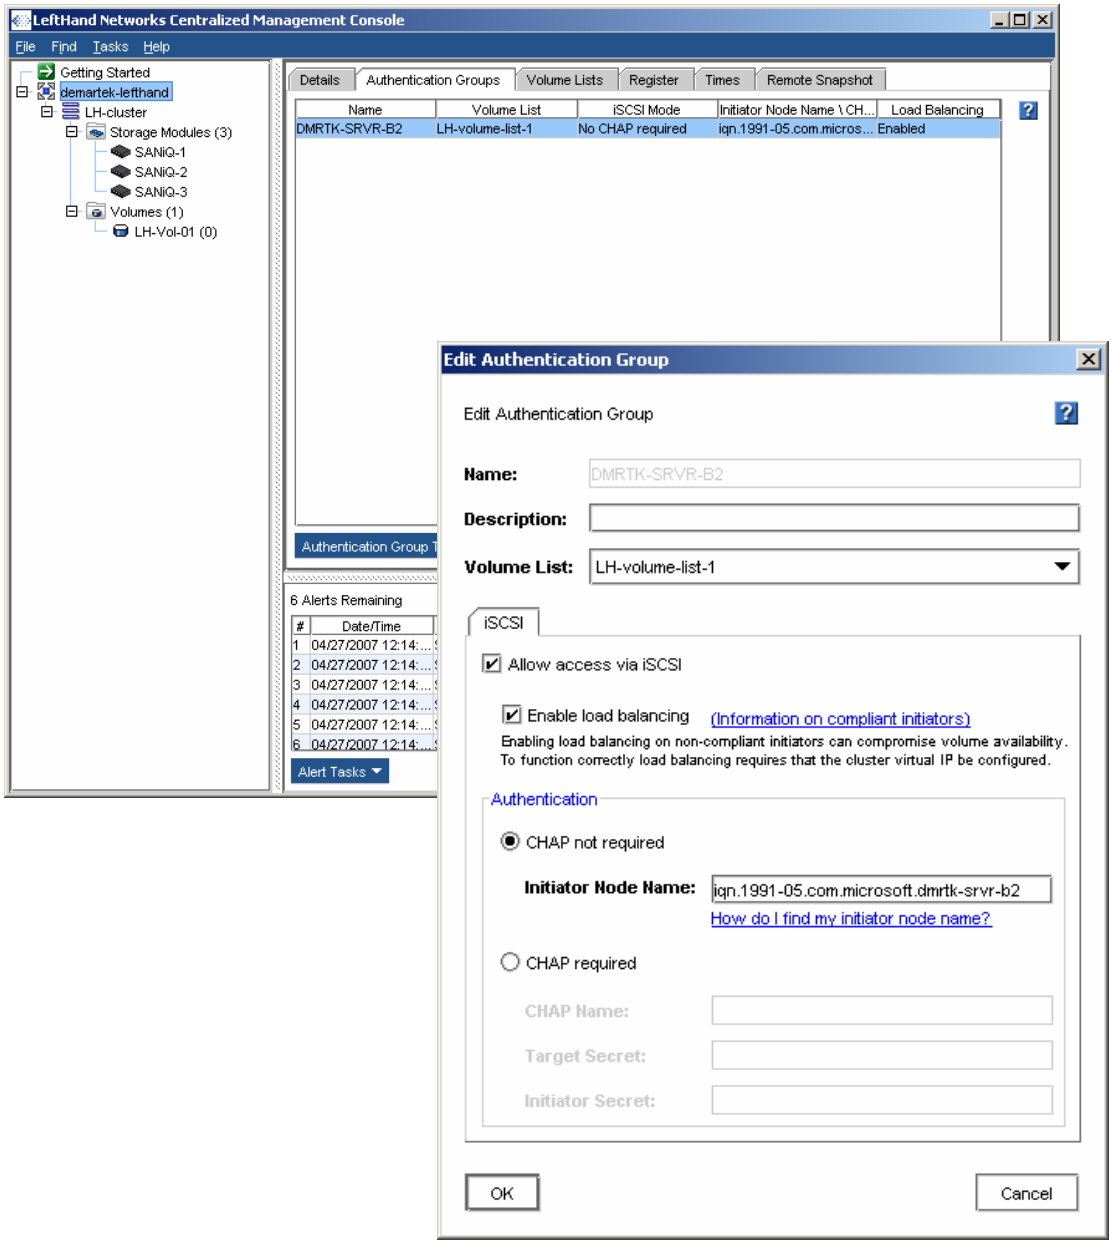

8. Make iSCSI Targets Ready for Use for iSCSI Initiators

No additional steps are needed to make the targets ready for use.

Initiator Configuration Steps

1. Configure Multi-path I/O from Application Host

MPIO for the initiator can be enabled by running the SAN/iQ Solution Pack from the application host.

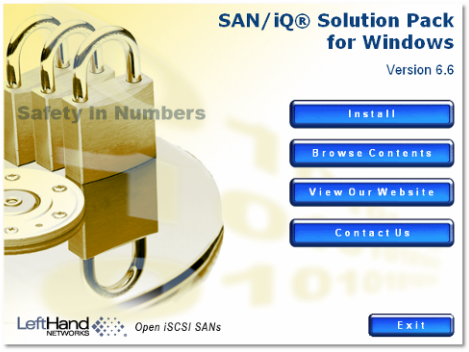

The SAN/iQ DSM for MPIO is selected, which begins the installation for the MPIO DSM.

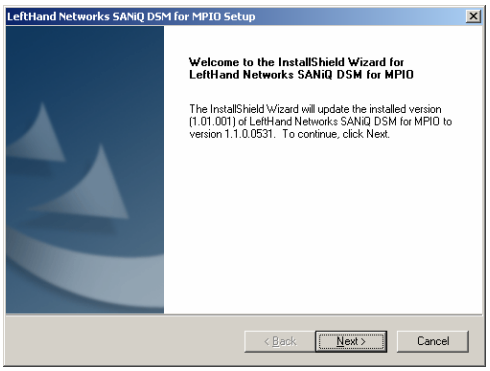

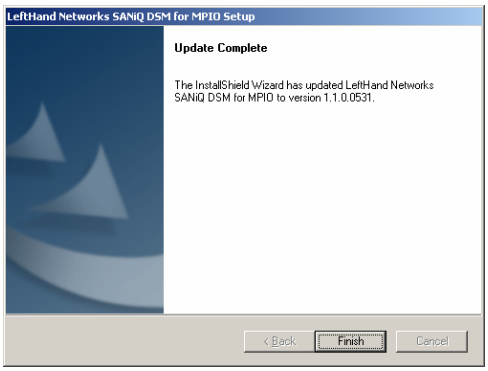

When the iSCSI initiator is launched to logon to the LeftHand Networks targets, the default addresses are selected, and MPIO is enabled for all the paths and targets. The "Advanced" tab in the initiator logon process is not needed.

Basic Performance Results

The following performance data is not intended to be viewed as a comprehensive performance benchmark, but to provide a general performance overview for the LeftHand Networks SAN/iQ solution.

Selected performance results are shown below, using a standard server-class network adapter, without receive-side scaling on the host. This configuration used two paths from two hosts, two I/O workers (one from each host), simultaneously accessing two target volumes and a queue depth of 20. Each host accessed its target volume with a pair of NICs configured as one “teaming” NIC. The target was configured as “non-mirrored.”

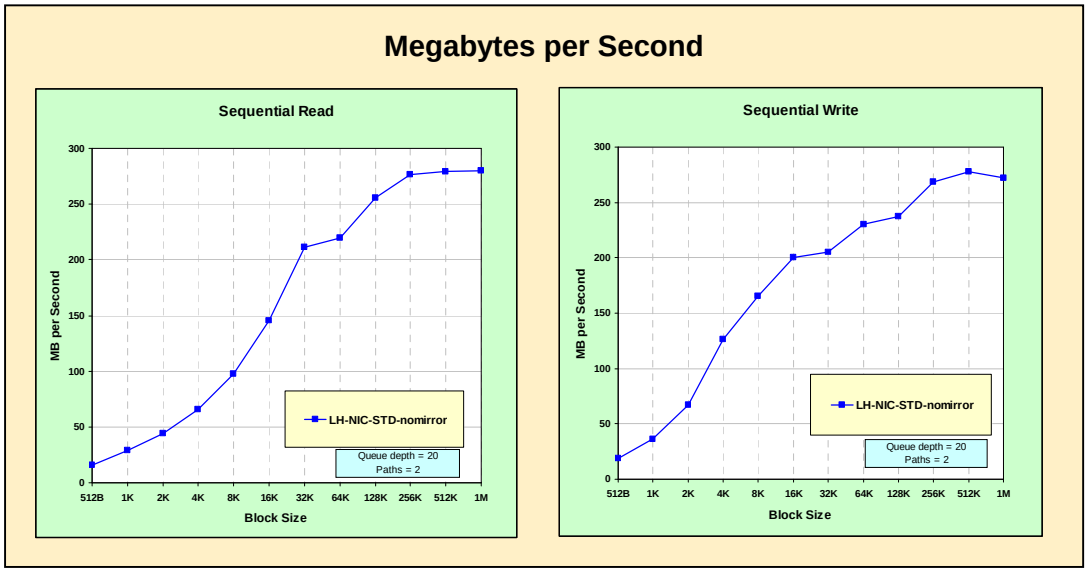

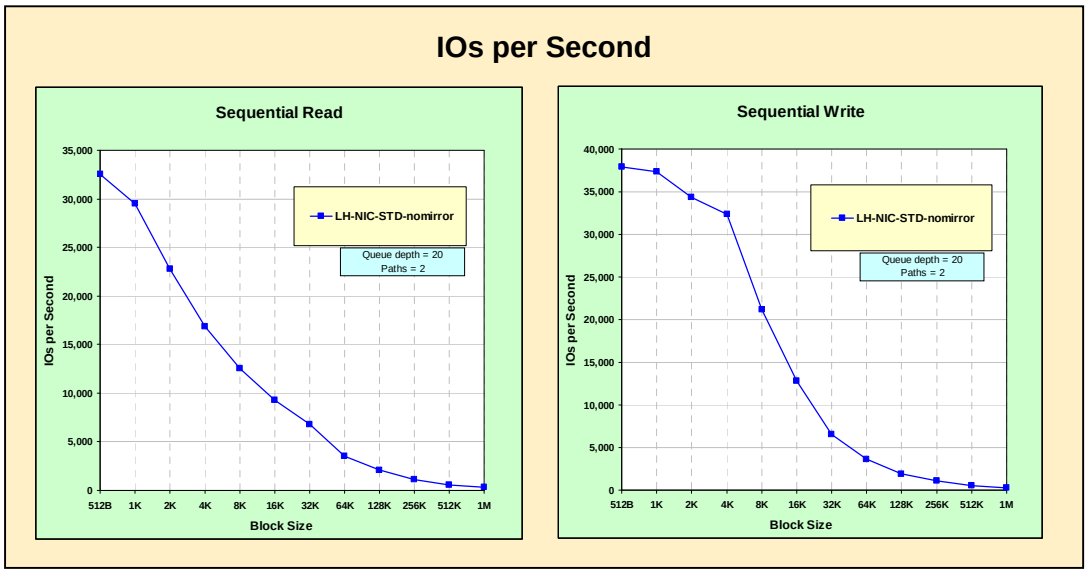

Storage Management Notes

Efficient Storage Management

Storage Manager for SANs

Storage Manager for SANs (SMfS) is a Microsoft Management Console snap-in that administrators can use to create and manage the logical units (LUNs) that are used to allocate space on storage arrays in both Fibre Channel and iSCSI environments. Administered through a conventional snap-in, Storage Manager for SANs can be used on storage area network (SAN) based storage arrays that support Virtual Disk Server (VDS) using a hardware VDS provider. Because of hardware, protocol, transport layer and security differences, configuration and LUN management differ for the two types (iSCSI and Fibre Channel) of supported environments. This feature will work with any type of Host Bus Adapter (HBA) or switches on the SAN. A list of VDS providers that have passed the Hardware Compatibility Tests (HCT) is available on http://www.microsoft.com/storage.

LUN management for Fibre Channel subsystems

On a Fibre Channel storage subsystem, LUNs are assigned directly to a server, which accesses the LUN through one or more Host Bus Adapter (HBA) ports. The administrator needs only to identify the server that will access the LUN, and enable one or more HBA ports on the server to be used for LUN I/O traffic. When the server is assigned to a LUN, the server can immediately access the LUN to create, augment, delete, and mask (or unmask) the LUN.

Support for multiple I/O paths. If a server supports Microsoft Multi-path I/O (MPIO), Storage Manager for SANs can provide path failover by enabling multiple ports on the server for LUN I/O traffic. To prevent data loss in a Fibre Channel environment, make sure that the server supports MPIO before enabling multiple ports. (On an iSCSI subsystem, this is not needed: the Microsoft iSCSI initiator (version 2.0) that is installed on the server supports MPIO.)

LUN management for iSCSI subsystems

Unlike on a Fibre Channel storage subsystem, LUNs on an iSCSI subsystem are not directly assigned to a server. For iSCSI, a LUN is assigned to a *target* – a logical entity that contains one or more LUNs. A server accesses the LUN by logging on to the target using the server's iSCSI initiator. To log on to a target, the initiator connects to *portals* on the target; a subsystem has one or more portals, which are associated with targets. If a server's initiator is logged on to a target, and a new LUN is assigned to the target, the server can immediately access the LUN.

Securing data on an iSCSI SAN. To help secure data transfers between the server and the subsystem, configure security for the login sessions between initiators and targets. Using Storage Manager for SANs, you can configure one-way or mutual Challenge Handshake Authentication Protocol (CHAP) authentication between the initiator and targets, and you can also configure Internet Protocol security (IPSec) data encryption.

Summary

Internet SCSI (iSCSI) can be a useful and relatively inexpensive way to provide storage for new applications or to provide a networked pool of storage for existing applications. Microsoft and its storage partners provide a variety of storage solutions that can be implemented relatively easily. This report allows administrators and IT managers to explore iSCSI technology and see actual deployment examples.

There is no question that iSCSI storage solutions and technology have a place in many IT environments. The performance of iSCSI storage solutions is adequate for many applications and iSCSI technology provides the benefits of storage area network technology for a lower cost than Fibre Channel storage solutions.

Related Links

For more information on storage for Windows Server Storage and iSCSI in particular, see the following:

- Microsoft Storage at http://www.microsoft.com/storage/
- Microsoft iSCSI Storage at http://www.microsoft.com/WindowsServer2003/technologies/storage/iscsi/default.mspx
- Microsoft Windows Storage Server at http://www.microsoft.com/windowsserversystem/wss2003/default.mspx
- Microsoft Windows Unified Data Storage Server 2003 at http://www.microsoft.com/windowsserversystem/storage/wudss.mspx
- Microsoft Storage Technical Articles and White Papers at http://www.microsoft.com/windowsserversystem/storage/indextecharticle.mspx
- Microsoft Scalable Networking Pack at http://www.microsoft.com/technet/network/snp/default.mspx
- Microsoft Exchange Solution Reviewed Program – Storage at http://technet.microsoft.com/en-us/exchange/bb412164.aspx
- Microsoft Cluster Server at http://www.microsoft.com/windowsserver2003/technologies/clustering/default.mspx

For more information on the Microsoft storage partner products mentioned in this report, see the following:

- Dell PowerVault NX1950 Networked Storage Solution at http://www.dell.com/content/products/productdetails.aspx/pvaul_nx1950?c=us&cs=555&l=en&s=biz
- EqualLogic PS3800XV at http://www.equallogic.com/products/view.aspx?id=1989
- HDS TagmaStore AMS1000 at http://www.hds.com/products_services/adaptable_modular_storage/
- HP StorageWorks 1200 All-in-One Storage System at http://www.hp.com/go/AiOStorage
- LeftHand Networks SAN/iQ at http://www.lefthandnetworks.com/products/nsm.php

For more information on RFC documents, see the following:

- RFC1334: CHAP and PAP at http://rfc.net/rfc1334.html
- RFC3720: iSCSI specification at http://rfc.net/rfc3720.html
- RFC4301: IPSec at http://rfc.net/rfc4301.html

For more information on IOMeter, the open-source I/O load generator and performance analysis tool:

- http://sourceforge.net/projects/iometer/