

Integrating ESRI GIS Services with NetApp Storage Devices

Gangoor Sridhara | October 2003 | TR 3289

TECHNICAL REPORT

Network Appliance, a pioneer and industry leader in data storage technology, helps organizations understand and meet complex technical challenges with advanced storage solutions and global data management strategies.

Abstract

Network Appliance (NetApp) provides enterprise storage solutions that deliver simple, fast, and reliable features. NetApp developed SnapDrive™ to integrate Windows® software-based applications with NetApp filers. The iSCSI protocol is a transport protocol for SCSI-3 that uses a standard TCP/IP network to carry SCSI-3 requests and responses. iSCSI is a peer to other SCSI-3 transport protocols, including FCP-SCSI. iSCSI initiators and targets communicate using an iSCSI session. This tight integration makes it quick to implement database solutions and the ArcSDE solution on a NetApp filer. It enables businesses to easily utilize NetApp enterprise-class file storage solutions. This document describes the procedure to install iSCSI software initiator and SnapDrive software, and describes how to integrate ESRI GIS services with filers.

Table of Contents

Table of Figures

1	Purpose and Scope	5
2	Introduction	5
2.1	Configuration Design	6
2.2	System Design Using Filers	8
2.3	Using Backup Solution in a GIS Environment for SDE Data	8
2.4	Use of NetCache Solution	9
2.5	Assumptions	10
3	Supported GIS Products	10
3.1	Infrastructure	10
4	Software Required for Windows 2000	11
4.1	iSCSI Software	11
4.1.1	Overview of iSCSI	11
4.2	SnapDrive 3.0 Software	12
4.3	Virtual Disk Overview	12
4.4	ESRI GIS Software	13
4.5	Windows 2000 Server Host	13
4.6	Network Configuration	13
4.7	Configuration Overview—Windows 2000 Server	14
4.8	Network Appliance Filer	14
4.8.1	Volume Creation	14
4.8.2	Creating a qtree	14
4.8.3	CIFS Shares	15
4.9	Enable rsh	15
4.9.1	Checking the Gigabit Ethernet Connection	15
5	Software Installation Procedure	17
5.1	iSCSI Software	16

5.1.1	Pre Installation Requirements	16
5.1.2	iSCSI Software Installation	16
5.1.3	SnapDrive 3.0 Software Installation	17
5.2	Starting FilerView	17
5.3	SnapDrive Service Credentials	18
5.4	Creating a Virtual Disk	19
5.5	Computer Management Window	21
5.6	Using a LUN	22
5.7	Converting VLDs to LUNs	22
6	GIS Software Installation	22
6.1	Options Available - ArcGIS	23
6.2	Installation of DBMS Servers	23
6.2.1	Installing Oracle Database	23
6.2.2	Installing Microsoft SQL Server	25
6.2.3	Installing Informix Software	26
6.2.4	Installation of ArcSDE	26
6.2.5	Space Required for Database	27
6.2.6	Example of Creating a Database with One File Group Using SQL Server	27
7	Migrating Databases	27
8	Using Snapshot with Database Server	27
8.1	Creating Snapshot Copies with SnapDrive 3.0	28
8.2	DBMS Server Crash Recovery	28
8.3	Snapshot Copies of Databases in a Single Volume	28
8.4	Snapshot Batch Program of Online Database(s)	28
8.5	CRE_SNAPSHOT.BAT for SQL Server	28
9	SnapRestore	29
9.1	Using Writable Snapshot Backups	29
10	Disaster Recovery and High Availability	29

10.1	Asynchronous Mirroring for Disaster Recovery	29
10.2	Local Synchronous Mirroring for High Availability	29
10.3	Clustered Fail-Over (CFO)	30
11	References	30

Table of Figures

- Figure 1 – Cluster Configuration
- Figure 2 – Storage Configuration
- Figure 3 – ESRI Design Using Filers
- Figure 4 – ESRI Design Using NearStore for Backup and Recovery
- Figure 5 – Using NetCache for Streaming Solution
- Figure 6 – iSCSI Software Initiator Installation Status
- Figure 7 – SnapDrive InstallShield Wizard
- Figure 8 – SnapDrive Service Credentials
- Figure 9 – SnapDrive 3.0 Software Installation Status
- Figure 10 – iSCSI Initiator – Active Sessions
- Figure 11 – Create Disk Wizard
- Figure 12 – Computer Management – Disks Created
- Figure 13 – SnapDrive 3.0 Created LUNs
- Figure 14 – Choose GIS Installation Type
- Figure 15 – Database Properties for SALES Database

1) Purpose and Scope

The purpose of this document is to explain how to integrate ESRI software and to set up filer-specific software or utilities on Network Appliance™ filer(s). This paper explains the procedure to integrate iSCSI software and SnapDrive software to configure ESRI ArcGIS and ArcSDE products using virtual disks on the Windows server. This paper is intended for people who are interested in knowing how to integrate NetApp storage devices with ESRI products.

2) Introduction

Network Appliance solutions are based on Data ONTAP™, a highly optimized and scalable operating system (OS), to provide a total data storage solution. Data ONTAP supports multiple protocols, such as NFS, CIFS, HTTP, FCP, SCSI, etc., to provide cross-platform support. SecureShare® is a multiprotocol lock-management facility that is integrated into the Data ONTAP microkernel. Write Anywhere File Layout (WAFL®) is a highly optimized file system and provides better performance and other benefits. NetApp storage devices come with RAID protection.

In February 2003, Network Appliance, Inc. (NetApp), announced the support of iSCSI target mode across a range of NetApp storage configurations. This enables filers to provide block-mode access to storage over a standard Ethernet network infrastructure. By using the existing network infrastructure, storage area network configuration can be achieved by using iSCSI software. SnapDrive 3.0 software supports an iSCSI logical unit number (LUN) that is configured as a local SCSI disk to the Windows server.

The filer head has hardware redundancy such as multiple network interfaces, Fibre Channel multipath, and duplicate power supplies. In addition to the single filer head configuration, NetApp filers can be configured to provide Clustered Fail-over (CFO) capability for the filer heads. Each storage appliance can operate independently in active/active configuration. The partner filer assumes the additional functionality of the failed storage appliance.




Figure 1 Cluster Configuration

You can refer to Network Appliance documentation for additional information.

2.1) Configuration Design

NetApp filers can be configured either directly to a GIS application server, attached to a network, or in a storage area network (SAN). For database server environments, we recommend either a direct-attached or an IP-based SAN configuration. On the Windows 2000 platform, we recommend using our SnapDrive software to create virtual local disks using the iSCSI protocol.

General configuration design is shown below for both direct-attached and SAN environments.

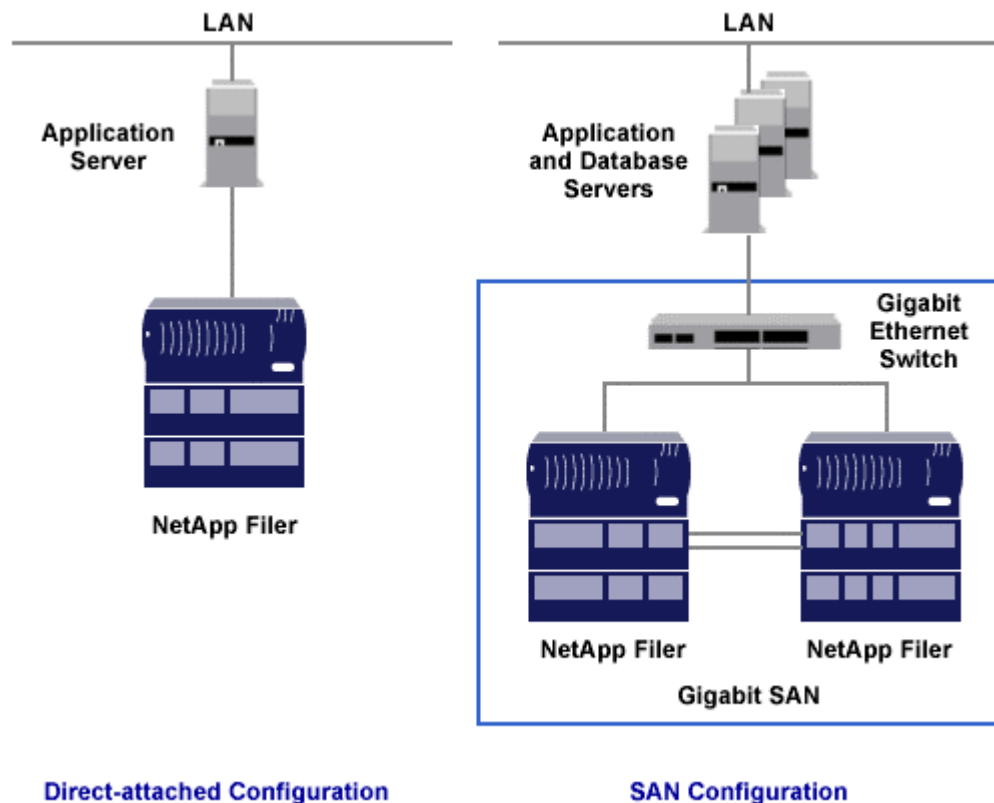


Figure 2 Storage Configuration

Before proceeding with the configuration, we suggest determining the disk space requirements to store GIS and other data. Once you have the information, increase the number by at least 50% to determine the amount of capacity. In this way, additional space requirements for data protection and any near future growth are allocated. If iSCSI LUNs are used with space reservation, it requires an equal amount of space to allow at least one Snapshot™ copy of the LUN. Initial Snapshot duplication of a LUN takes an equal amount of storage, and incremental backups do not require the same amount of space as the first Snapshot.

SnapDrive has evolved to utilize more I/O protocols and integrates with more NetApp filer features, but even so, it is extremely easy to integrate database servers with NetApp filer technology. SnapDrive 3.0 (use of qtree) has increased the flexibility when multiple devices utilize a single NetApp filer volume. Support for asynchronous mirroring has made it much easier to create a robust database server environment for disaster recovery. Support of the new LUN I/O protocol has drastically increased the I/O throughput SnapDrive can support, and it has made it easier to integrate into existing data center infrastructures, but it has also made it more complicated to determine which I/O protocol to integrate. The different protocols will be discussed when appropriate, but the focus of this document is not the I/O protocols.

This technical report will describe integration of ArcGIS and ArcSDE software with the filer. It explains in detail the NetApp solution-specific installation procedure rather than the installation procedure for

ESRI software. For the detailed ArcGIS and ArcSDE installation procedure, please refer to ESRI documents. For detailed information about qtrees-based SnapDrive features for the database environment, refer to the *SnapMirror® Installation Guide for SQL Server*.

If performance improvement is expected compared to existing iSCSI setup with a clean network and proper working of NetApp storage servers, we recommend configuring storage using the FCP protocol or iSCSI HBA hardware configuration.

2.2) System Design Using Filers

The ESRI system design using a storage area network (SAN) can be configured with single or multiple filers at various points for storing ArcGIS data where ArcGIS services are running, and for storing ArcSDE data on a different UNIX® server. Oracle® databases are known to perform well in a filer configuration. A sample design diagram for a SAN configuration is shown in figure 3.

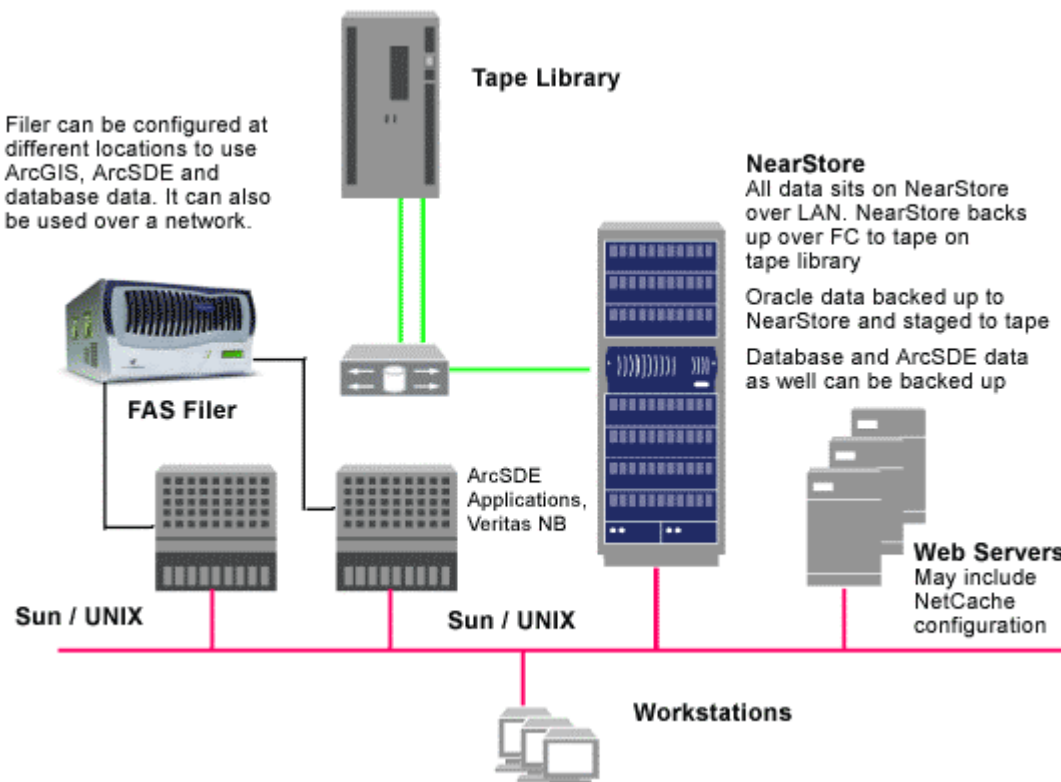


Figure 3 ESRI System Design Using Filer

2.3) Using a Backup Solution in a GIS Environment for SDE Data

The NetApp NearStore® R100 or R150 solution is efficient and cost-effective. This solution reduces the complexity of storage backup efforts involved with tapes. As NearStore provides tape compatibility, the data can be backed up to tapes from the nearline storage. Even if the data resides on a SAN or NAS environment, it can be backed up to R100/R150 systems. Since NearStore provides disk-based backup

and offers Snapshot and SnapRestore® features, backup and restoring the data are easy to manage. If the data is to be archived to tape, it can be dumped to tapes using NearStore configurations.

The following diagram shows the relationship between the primary storage, backup target, and secondary storage. You may use this information as a starting point in deciding the best configuration in your environment.

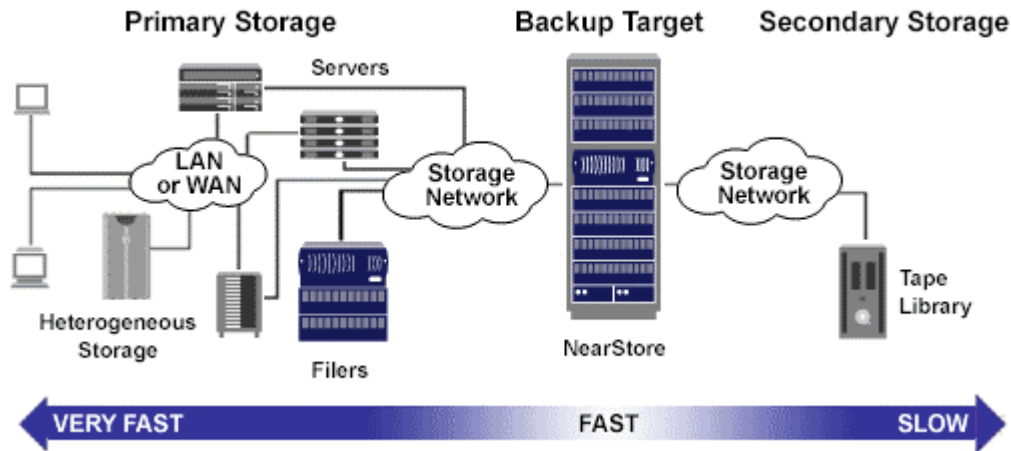


Figure 4 ESRI Design Using NearStore for Backup and Recovery

2.4) Use of the NetCache Solution

If an efficient streaming content solution requires Web servers in the above configuration, a NetCache solution may be configured. One such solution offered by NetCache is shown below.

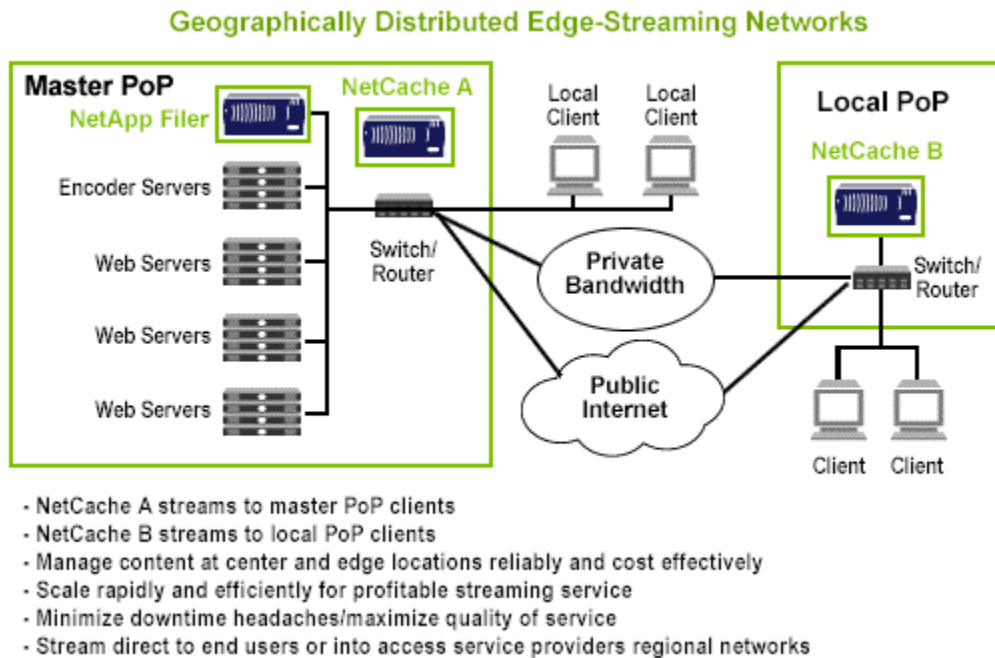


Figure 5 Using NetCache for Streaming Solution

2.5) Assumptions

This technical report assumes that the reader is knowledgeable about ESRI GIS products and Windows 2000 administration. For information on how to install a filer, please refer to the Network Appliance product manual [Data ONTAP System Administrator's Guide](http://now.netapp.com/Self-Service/Start/usrPage.asp) (<http://now.netapp.com/Self-Service/Start/usrPage.asp>), available on our Web site. For additional information about SnapDrive, please refer to *SnapDrive System Administrator's Guide*. Managing Network Appliance filers is easy and simple. For ESRI-GIS products, please refer to [ESRI](http://www.esri.com/) product reference guides. <http://www.esri.com/>

3) Supported GIS Products

SnapDrive v 3.0 is application-independent and all ESRI-supported database versions work well with the NetApp storage devices. On Windows platforms, virtual disks are suggested for installing ArcSDE and storing SDE data. On UNIX, either network-attached storage (NAS) or a SAN configuration is supported. All ESRI products are supported on the filer configuration. ArcSDE software has been installed and tested in a filer configuration.

3.1) Infrastructure

Filer and GIS products can be used in different configurations depending on the user requirements. The filer should be configured as a member of the Windows domain. If the user environment requires a local disk configuration, it can be accomplished on the filer by installing SnapDrive software on the Windows 2000 machine. SnapDrive 3.0 supports iSCSI LUNs as well as FCP LUNs.

The following infrastructure was used to configure and use SnapDrive in our test configuration:

The Windows 2000 SP4 Server

800 class or 900 class filer(s)

The name of the NetApp filer for GIS applications is guy

The name of the NetApp filer for database and SDE is sunny

The name of the server hosting GIS is cluster1

The location of the GIS binaries on the Windows 2000 system is K:\

The name of SDE and the Oracle database Windows machine: cluster2

The Oracle HOME directory on cluster2 machine is J:\ora

The target volume of the GIS on the NetApp filer is content

The qtree used for GIS software data is /vol/content/gis

The qtree used for GIS data is /vol/content/gisdata

The qtree used for ArcSDE is /vol/content/sde

CIFS share names used are gis, sde, gisdata, sdelog, oracle

The virtual disk (iSCSI LUN) data is on both cluster1 and cluster2

The name of the Windows domain is DBLAB

The name of the domain administrator account that has admin privileges is
DBLAB/Administrator

4) Software Required for Windows 2000

Installation of software on the Windows server requires installation of iSCSI initiator software and SnapDrive 3.0 software installation/upgrade from previous releases. This section discusses the procedure required to complete the storage configuration.

4.1) SCSI Software

4.1.1) Overview of iSCSI

The iSCSI protocol is a transport protocol for well-known SCSI-3 (SCSI) that uses a standard TCP/IP network to carry SCSI requests and responses. It's a peer-to-other SCSI-3 transport protocol such as FCP-SCSI. The TCP/IP protocol provides the guaranteed in-order delivery of data over a variety of underlying physical media.

In February 2003, NetApp announced the support for iSCSI target mode across a broad range of its storage appliances. This enables the NetApp storage appliances to provide block-oriented access to

storage over a standard Ethernet networking infrastructure. The SCSI-3 protocol defines a mechanism by which initiators access data provided by data servers known as targets.

NetApp filers support SCSI-3 block-level access using either an FCP-SCSI or iSCSI transport protocol. The data container used by SCSI-3 targets is called logical unit number (LUN). NetApp storage administrators can create and manage these LUNs using the appropriate transport mechanism such as FCP-SCSI or iSCSI. These LUNs can be included in advanced storage management provided by NetApp storage devices.

4.1.2) .iSCSI Software

Microsoft supports iSCSI software initiator on its Windows environment and has announced the support of iSCSI software initiator on its Windows platform. You can get the software from the Network Appliance [customer support site](http://now.netapp.com/Self-Service/Start/usrPage.asp) (<http://now.netapp.com/Self-Service/Start/usrPage.asp>) or from the [Microsoft web site](http://www.microsoft.com/downloads/search.aspx?displaylang=en) (<http://www.microsoft.com/downloads/search.aspx?displaylang=en>). As a prerequisite to the setup, download the latest version of software on the Windows server machine. This paper suggests using NetApp support iSCSI HBA hardware in order to use TCP Offload Engine (TOE) for improved performance. This will help the processor to do other jobs, thus improving performance.

Currently supported iSCSI HBA details may be found at our [web site](http://now.netapp.com/Self-Service/Start/usrPage.asp). <http://now.netapp.com/Self-Service/Start/usrPage.asp>

4.2) iSCSI Software

SnapDrive is software that integrates with Microsoft® Windows Volume Manager to enable the use of filers as virtual storage for application data in a Windows 2000, Windows XP, or Windows 2003 environment. To the Windows host, the storage on the filer is configured as a local SCSI disk.

SnapDrive software offers a rich set of capabilities to virtualize and enhance storage management for Microsoft Windows environments. It is tightly integrated with the Microsoft NTFS file system and provides a layer of abstraction between application data and physical storage associated with that data.

The SnapDrive 3.0 component requires Data ONTAP 6.4.2 or later software. SnapDrive resides in the Windows 2000 computer management storage tree, and it provides a user interface for configuring and managing the virtual disks.

If a SAN storage configuration is required, it can be accomplished by installing the FCP solution from Network Appliance. The FCP driver resembles a Windows 2000 storage-class disk driver, acts as a SCSI/RAID controller, comes with an HBA for Windows hosts, and gets installed like a hardware component. For additional information, please visit the [NetApp portal](http://now.netapp.com/Self-Service/Start/usrPage.asp). <http://now.netapp.com/Self-Service/Start/usrPage.asp>

SnapDrive manages virtual disks on a NetApp filer, making these virtual disks available as local disks on the Windows hosts. This allows the Windows hosts to interact with these disks as if they belonged to a directly attached RAID array. SnapDrive can be configured in a nonclustered configuration as well as Windows clusters (MSCS) and NetApp fail-over topologies.

4.3) Virtual Disk Overview

To the NetApp storage device, a virtual disk also known as a LUN is a logical representation of a physical unit of storage. In this context, the storage device handles each LUN as a single storage object. In a NetApp SAN environment, an initiator (Windows host) initiates a SCSI I/O operation to a target (NetApp storage device). The operation can be initiated using either FCP or the iSCSI protocol. Once a target receives a SCSI request, writing data to or fetching data from the LUN performs the necessary operation. SnapDrive 3.0 utility allows the management of these LUNs. Using the SnapDrive 3.0 management utility, LUN size can be increased but cannot be shrunk. The virtual disks (LUN) storage is configured as direct-attached storage with RAID protection to RDBMS and allows ArcSDE data to be stored on the filer seamlessly.

4.4) ESRI GIS Software

Before proceeding further, make sure you have all the necessary hardware, ESRI software such as ArcGIS and ArcSDE, and software license keys from ESRI. If you do not have license keys, please contact ESRI. In our test environment, we installed ArcSDE software on one Windows server machine and Oracle 8.1.7 and Oracle 9.0.1.2 databases on different Windows servers in the same network. Both SQL Server and Informix databases are supported in this configuration. This configuration is similar to that of most ESRI customers where ArcSDE is installed at a central location and clients access the GIS and SDE services. Using iSCSI software, we used one filer to configure local SCSI disks on two different Windows servers. In your environment, you can make appropriate changes to this configuration.

4.5) Windows 2000 Server Host

A Network Appliance configuration for ESRI GIS applications can be used in different environments. You can configure a filer using CIFS, NFS, iSCSI, or SAN. An Intel® processor-based server running Windows 2000 with Service Pack 3 with certain patches or ISP 4 and SnapDrive 3.0 or later are required for a SAN environment. Any member of the local or domain administrator's group can install GIS. The server should be a member of the same domain or active directory (AD) tree as the filer. Licensing iSCSI and CIFS on the filer enables the CIFS protocol to allow joining the filer into the domain or AD. This is necessary to enable SnapDrive 3.0 to perform commands on the filer via remote procedure call (RPC).

4.6) Network Configuration

Existing Ethernet networking infrastructure can be used to configure iSCSI LUNs. However, for improved performance reasons, use of a Gigabit Ethernet network connection between the ArcSDE server and the filer is strongly suggested. This can also be accomplished using a crossover cable in a simple networking topology or a dedicated switch with only GIS and NetApp filer workloads. The network used in this configuration consisted of 100Base-T full-duplex connectivity between the clients and ESRI GIS Server, and Gigabit Ethernet (1000Base-SX) full duplex between the GIS Server and filer. ArcSDE and ArcGIS clients can still connect to the services using their existing network infrastructure. A gigabit SAN or dedicated network connection between the GIS Server machine and the NetApp filer is suggested for the following reasons:

Any issues of contention or latency are eliminated if the GIS Server machines and NetApp filer are the only nodes in the network.

Security is no longer a major concern and is ensured by creating a private network connection. There is no issue of protecting the GIS databases from being exposed.

Adding NetApp filers to increase capacity or performance is assured the same security levels as installing the first NetApp filer.

It is strongly recommended that Gigabit Ethernet be used to support your GIS database environment on the filer. Refer to "[Architecting a Gigabit SAN for Database Application Environments.](http://www.netapp.com/tech_library/3112.html)"
http://www.netapp.com/tech_library/3112.html

4.7) Configuration Overview - Windows 2000 Server

For complete requirements, refer to the *SnapMirror System Administrator's Guide* from Network Appliance. With reference to SnapDrive 3.0:

1. Connect the Windows server to the filer via a Gigabit Ethernet connection.
2. Check the Intel Gigabit NIC setup for full-duplex mode.
3. Set the Intel Gigabit NIC flow control to FULL.

4.8) Network Appliance Filer

Installing and configuring NetApp storage devices is simple and easy to manage. It is also very easy to scale the configuration depending on growth of data. For more information, please refer to the filer installation guide. The above document is also available online from Network Appliance. Traditional installation of servers on the Windows 2000 platform suggests that the storage of all components reside on local disks.

4.8.1) Volume Creation

To use iSCSI LUNs, a volume must be created with a CIFS Share using the volume path. We suggest creating a qtree for each share on the same volume. To complete this task, either the FilerView® GUI utility or a command line mode can be used. To start FilerView, start a browser with http://filename/na_admin/ where the filer name is the name of the iSCSI target filer or its IP address.

FilerView → Volume → Add will guide you through the volume creation process. Once the volume is created, create a Qtree and CIFS Shares.

FilerView → Volume → Qtree → Add

FilerView → CIFS → Shares → Add. Before creating a CIFS Shares, CIFS must be running. To set up CIFS services, run the `cifs setup` command or use the FilerView GUI utility. To start the cifs services successfully, the filer must be added to the Windows domain.

The number of disks and RAID groups has to be configured appropriately to your environment. Default RAID group size is 14 disks, and it is strongly suggested that you don't increase the RAID group size beyond 14. Creating volume(s), qtree(s), and CIFS Shares can be done either at a command line or using the FilerView console.

```
filerA> vol create /vol/content 9 -r 14
```

Note that one disk for each volume will be used as a RAID disk to provide data integrity.

4.8.2) Creating a qtree

Qtrees can be created using the FilerView GUI utility. You can also create a qtree or folder on each volume using the following commands:

```
filerA> qtree create /vol/content/gis
filerA> qtree create /vol/content/gisdata

filerA> qtree create /vol/content/sde
filerA> qtree create /vol/content/oracle
```

4.8.3) CIFS Shares

Next, we need to create CIFS shares on the filer for the GIS data directory and logs using the following CIFS commands. On the NetApp filer console, enter:

```
filerA> cifs shares -add gisdata /vol/content/gisdata
filerA> cifs shares -add gis /vol/content/gis

filerA> cifs shares -add sde /vol/content/sde
filerA> cifs shares -add oracle /vol/content/oracle
```

4.9) Enable rsh

Using SnapDrive requires the filer to be enabled for remote shell (rsh) from the OS machine. To enable rsh, use the FilerView utility. Follow these steps:

Select Security ---> Manage Rsh

Enter the OS server IP address

Enter a user name similar to Administrator or the SnapDrive service account user name

Enter any comments followed by #

Click Apply

Click Close

After completing the above steps, test whether rsh has been successfully enabled by issuing a command on the OS server: `rsh filename ?` where `filer name` is name of the filer. This command should display all the commands available on that filer. If an error message is displayed, read it carefully and refer to the *Data ONTAP System Administrator's Guide*. Rsh has to be enabled for successful SnapDrive software installation.

4.9.1) Checking the Gigabit Ethernet Connection

Before installing any SnapDrive software, please check the Gigabit Ethernet connection between the OS machine and the filer. This can be achieved by using the ping command.

On the filer command line, when you ping the IP address of GigE on an OS machine, you should get a network connectivity response.

If you need further help, please check NetApp's [NOW™](http://now.netapp.com) site or contact the NetApp support organization. <http://now.netapp.com/Self-Service/Start/usrPage.asp>

5) Software Installation Procedure

SnapDrive 3.0 supports managing iSCSI LUNs and it is only supported under Windows 2000 Server, Windows XP, or Windows 2003 platforms. This section explains the steps involved with iSCSI software initiator and SnapDrive installation procedure.

5.1) iSCSI Software

NetApp supports the iSCSI protocol as an iSCSI target device and Microsoft supports it as an iSCSI software initiator device. Microsoft offers and supports iSCSI software for its customers and downloads the software by visiting their web site <http://www.microsoft.com/downloads/search.aspx?displaylang=en>. To enable the target device on the NetApp storage device, CIFS and iSCSI product licenses have to be enabled. If you have the CIFS license, you can obtain the iSCSI license by visiting [NetApp](http://now.netapp.com/Self-Service/Start/usrPage.asp) site. <http://now.netapp.com/Self-Service/Start/usrPage.asp> Download the iSCSI attach kit.

5.1.1) Pre Installation Requirements

In order to use the iSCSI protocol, the CIFS license has to be enabled. The iSCSI license can be downloaded from the NetApp site at no charge. However, the storage device should have a CIFS license activated. Even though SnapDrive 3.0 requires a Gigabit Ethernet network, iSCSI does not have such a restriction and normal 100Base Ethernet network infrastructure could be used to configure iSCSI LUNs.

5.1.2) iSCSI Software Installation

First we will install an iSCSI software initiator on the Windows server by selecting the iSCSI software initiator installation file and running it. The installation utility will guide you through the process.

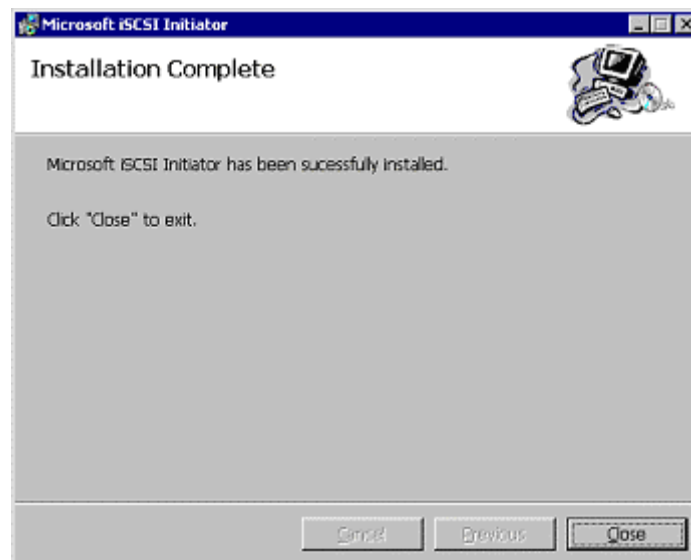


Figure 6 iSCSI Software Initiator Installation Status

5.1.3) SnapDrive 3.0 Software Installation

Login to <http://now.netapp.com> to download the SnapDrive 3.0 software. This site requires an authorized userid and password from NetApp. If you do not have one and wish to register one, please do so by going to [NOW](#) (NetApp on the Web). It's a good idea to check for network connectivity, such as a Gigabit Ethernet connection between the filer and the Windows server machine. The network can go through a switch or it can be connected on a point-to-point basis between the filer and the Windows machine. Make sure both Gigabit Ethernet connections are on the same subnet and use the same broadcast address. For details, please refer to the *SnapDrive Installation Guide*. Pinging the address from both the Windows server and the filer can check typical connectivity. With iSCSI, you can configure on your existing network infrastructure; still we strongly suggest that you have Gigabit Ethernet network connectivity between the NetApp storage device and Windows server machine. On our test setup, the file was saved to c:\temp. The following figure shows the progress of the software download and where it's saving the file.

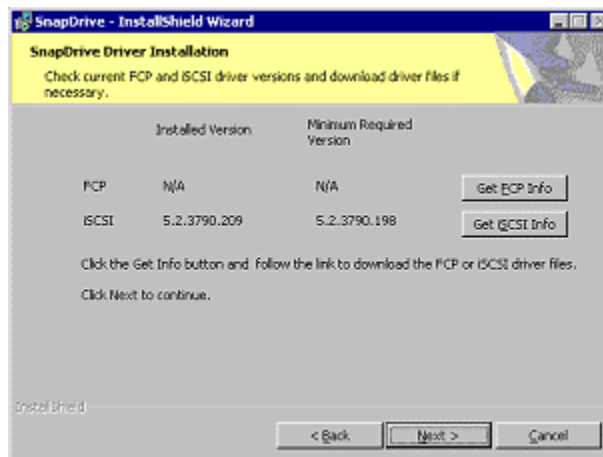


Figure 7 SnapDrive InstallShield Wizard

The SnapDrive installation wizard guides you through the process. Appropriate MPIO driver information will be displayed or installed if required on the Windows server that has a valid FCP configuration.

5.2) Starting FilerView

After saving the file, start FilerView to administer it. On our test environment, we started a browser with the command http://guy/na_admin/ where *guy* is the filer connected to the Windows servers *cluster1* and *cluster2*. Enter the filer root user and password.

FilerView allows you to manage the filer. Using the security, manage rsh, enable remote shell (rsh) to access the filer from the Windows system. To achieve this, enter the IP address and the user name of the Windows server on the filer's `/etc/hosts` file. In our case, we provided the IP address of the Gigabit Ethernet NIC address of the Windows server. To make sure rsh is enabled on the filer, open a command shell and execute the `rsh filename help` command where filename is to be replaced

with the actual name of the filer. This should display all the commands available on the filer. Before proceeding further, make sure you can execute rsh commands from the Windows server machine. If you still have issues, please refer to the SnapDrive installation guide.

5.3) SnapDrive Service Credentials

For virtual disks (logical unit number, or LUN) in iSCSI, you have to ensure that the specified account is a member of the filer's local administrator's group as well as the local administrator's group of the Windows OS system. For virtual disks, the user account must have root access to the filer. Please refer to the *SnapDrive Installation Guide* for more details.

Now select the correct SnapDrive Service Credentials to specify account information that will be used by the SnapDrive service. This can be administrator user on the Windows domain server.

Now enter the DOMAIN\administrator domain password (corresponding domain\user and password) that you used to log onto the machine. (Note that you are free to select any user name in place of the administrator account as shown in the figure, but the user must have administrator privileges on the local machine). After entering the domain account password, the following output is shown:

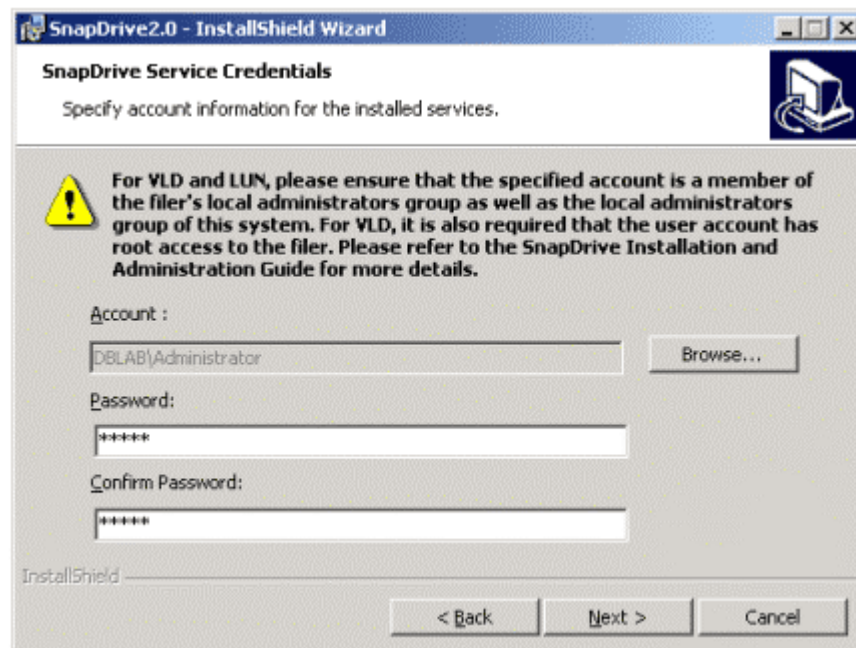


Figure 8 SnapDrive Service Credentials

Next, select the appropriate user ID and group ID. The InstallShield wizard will install the SnapDrive software. Select the Install button if you are ready for installation of the software. Continue with the InstallShield wizard to complete the SnapDrive software installation. In order to continue, make sure that iSCSI software initiator and SnapDrive software are installed successfully. On our system, the SnapDrive 3.0 installation wizard display appears as shown below:

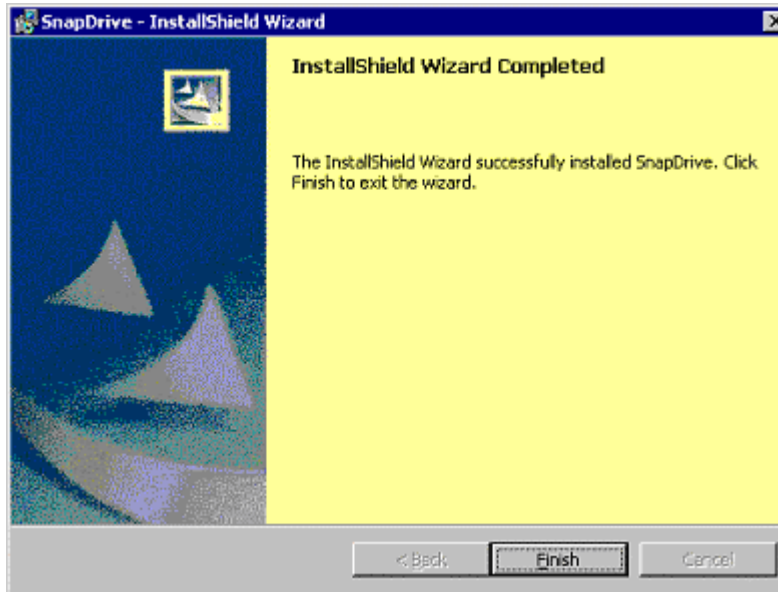


Figure 9 SnapDrive 3.0 Software Installation Status

5.4) Creating a Virtual Disk

Before proceeding further, verify that a volume, qtree, and CIFS shares are created. You can create multiple qtrees using the same volume or different volumes. Once the qtree is created, create CIFS shares using the qtree. This can be done from the FilerView menu. The procedures to create volumes, qtree, and CIFS shares have been discussed earlier in this report.

At this point, this paper assumes that you have been able to install iSCSI and the SnapDrive 3.0 software successfully. To create a LUN using SnapDrive 3.0, start the Computer Management and click SnapDrive under Storage, and then right-click on Disks and select Create a disk from the pop-up menu. The utility guides you through the process of disk creation. Note that during the disk creation process, you'll select the iSCSI initiator group (or create one and use it).

Before creating the virtual disk, check that the iSCSI initiator properties and active session are available as shown below:



Figure 10 iSCSI Initiator–Active Sessions

To create a virtual disk, start the Computer Management window and select SnapDrive under the Storage section. Right-click on Disks, select Create a virtual disk, and follow through the disk creation wizard instructions.

With the successful creation, the Create Disk wizard was displayed on our system.

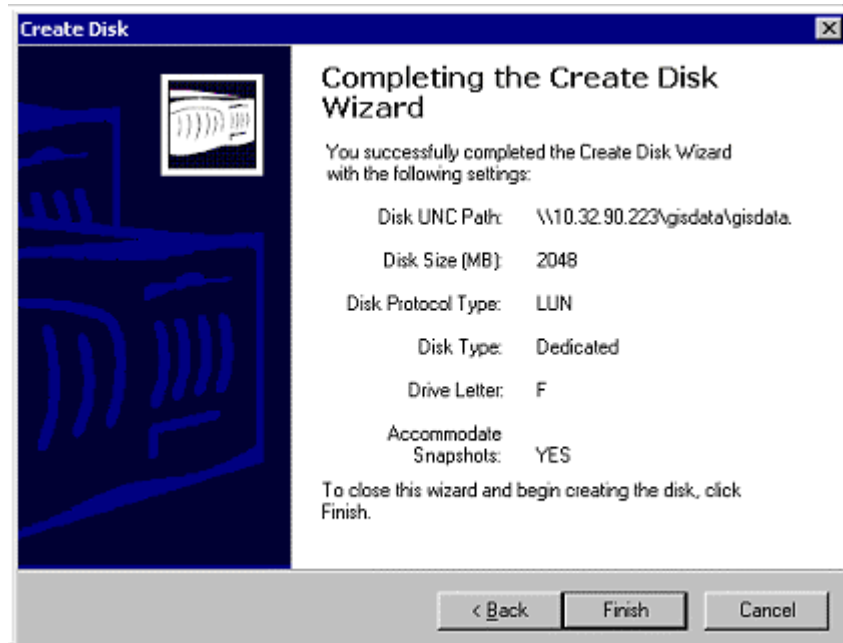


Figure 11 Create Disk Wizard

5.5) Computer Management Window

Disks created on the filer as LUNs can be viewed on the computer management window. On our test setup, this was the output:

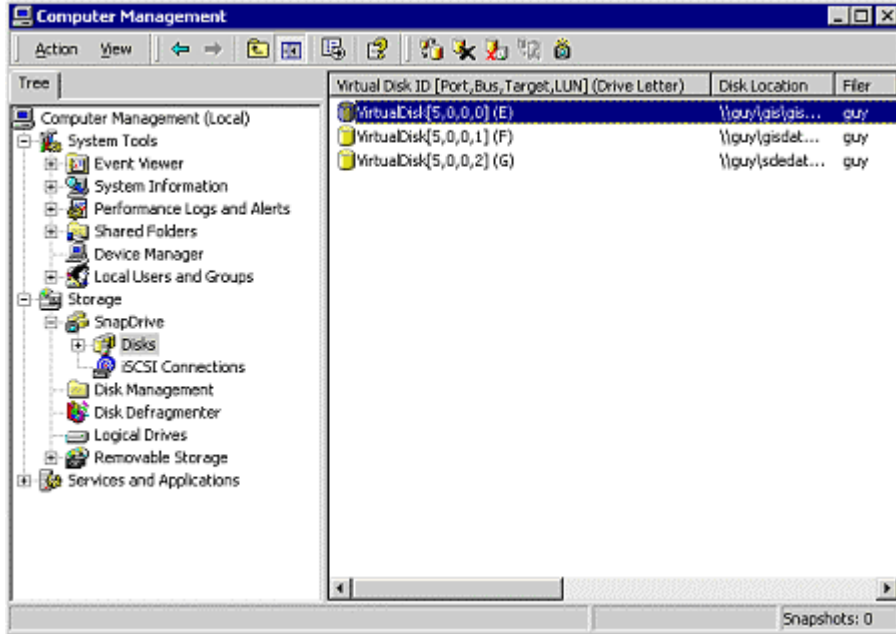


Figure 12 Computer Management - Disks Created on the Filer

At this time, the filer shows that the following iSCSI LUNs have been created and that similar virtual disks have been created on the ArcSDE server machine for installing the database server and ArcSDE services.

Manage LUNs ?		
LUNs → Manage		
Add New LUN	Hide Maps	
LUN	Description	Size
/vol/content/data/gisdata.lun	ArcGIS data disk	4.006 GB
/vol/content/gis/gis.lun	ArcGIS Installation	10.004 GB
/vol/content/oracle/oracle.lun	Oracle Home Directory	5.002 GB

Figure 13 SnapDrive 3.0 Created LUNs

5.6) Using a LUN

In system environments where I/O performance is a major requirement, an FCP LUN is an excellent choice. It delivers high-end I/O throughput and servicing times. A LUN is as stable and reliable as a local SCSI disk and will deliver greater maximum I/O throughput and shorter I/O latencies. A LUN together with Multi-Patted I/O (MPIO) potentially supports much higher I/O throughput compared with other storage configurations but with a higher initial cost.

5.7) Converting VLDs to LUNs

SnapDrive 3.0 supports iSCSI LUNs and NetApp suggests converting VLDs into LUNs. The procedure is given in the section above. Most customer sites already have some Gigabit Ethernet infrastructure configured and available for the new SQL Server environment. Even when a new Gigabit Ethernet infrastructure has to be created, it is very cost-efficient and reliable, and it will in most cases satisfy system requirements. If a Fibre Channel infrastructure is already in place and can be used with NetApp SAN technology, then it makes sense to use the LUN technology. In cases where the performance requirements are not well understood, start by testing the existing network infrastructure using iSCSI, and if that environment doesn't fulfill expectations, then use of a Gigabit Ethernet network connection is suggested.

6) GIS Software Installation

The ESRI ArcGIS has three key parts and they are:

- ArcGIS desktop software that is an integrated suite of advanced GIS applications such as ArcReader, ArcView, ArcEditor, and ArcInfo

- ArcSDE as an interface to manage spatial data in a GIS database system

- ArcIMS software, Internet-based GIS for distributing data and services

ArcGIS allows you to implement GIS for multiple users including single-user requirements. ArcGIS solutions can be deployed on a distributed network, on a network of servers, or on a single desktop. ArcGIS 8.3 supports a file-based data model of coverages, shapefiles, and attributed tables.

ArcMap, ArcCatalog, and ArcToolbox are designed to work together to perform all GIS tasks.

Installing ArcSDE on the filer adds benefits in addition to:

- Manage and maintain your geodatabase in the DBMS

- Serve multiple ArcGIS users anywhere on your network

- Open your ArcGIS database to other users

- Serve data openly and securely

Using storage solutions provided by NetApp filers can easily solve the above requirements/tasks.

To install ArcGIS, please refer to the ArcGIS product manuals from ESRI. The following figures explain the process of installing ArcGIS 8.2 on a Windows 2000 server in our test setup. Depending on the installation type, select the correct procedure. We do not attempt to explain the installation procedure for the ArcGIS product. If the license manager is not installed, it will prompt you to install it. It is a good

idea to keep the software license key code ready and to make sure that a hardware key is plugged into the Windows server.

ArcSDE can be configured in two different ways. One is to have direct connection between a client and the database management system (DBMS) on the server platform or to have an application server configuration. The most commonly visible configuration includes an ArcSDE application server.

Before proceeding further, refer to system configuration details provided by ESRI for ArcSDE implementation and the database vendor's configuration documents.

6.1) Options Available - ArcGIS

ArcGIS manuals provide detailed information on this subject. We suggest that you have access to all the relevant information before starting installation of the ESRI software. Based on your requirements, select the Typical, Complete, or Custom radio button. If you select Complete, it will guide you through the entire process of installing all the extensions. ArcGIS setup options are listed in the following figure.

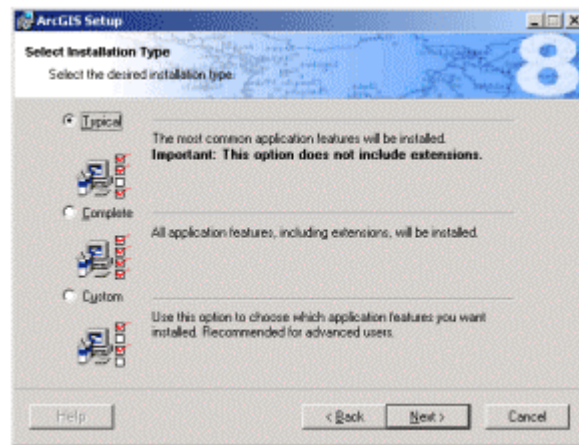


Figure 14 Choose the GIS Installation Type

Once all the extensions and ArcGIS are installed, prepare for installation of other ESRI software products, such as ArcSDE. ArcSDE 8.2 supports Oracle, Microsoft SQL Server, and Informix databases.

6.2) Installation of DBMS Servers

The filer supports all RDBMS software. Before installing database software, check with ESRI to find out the supported database-specific versions. Once you have that information, installation is a straightforward task.

6.2.1) Installing Oracle Database

Installing Oracle for GIS services involves the following the process:

Place Oracle Home on a filer

Create an Oracle Home server with a filer

Store multiple Oracle versions on a filer

Create a new instance with data files on a filer

Migrate an instance from local disk onto a filer

The following items are needed to run Oracle8[™] with a filer:

Oracle8 Server machine (this could be a UNIX or Windows 2000 Server)

NetApp filer

Network

Oracle user account

Filer NFS mount point (for UNIX environment)

Before the installation of Oracle software, choose the Oracle Home directory. Placing Oracle Home on the filer offers advantages. Among other things, the list includes:

With Oracle Home on a filer, you can easily configure a manual fail-over server. This will be the subject of a future technical paper.

With Oracle Home on a filer, you can run multiple instances of Oracle8 on separate database servers that use the same set of executables. This saves the space otherwise required to store multiple copies of the Oracle8 executables on separate database servers.

With Oracle Home on a filer, you can easily set up multiple versions of Oracle on the same server, and then switch between them. This is ideal in a lab environment, where multiple versions of Oracle must be tested.

It is also very important to know the disadvantages of installing Oracle Home directory as mentioned earlier. Placing Oracle Home on a filer creates the following issues:

You have an additional point of failure in your Oracle installation. If you plan to store Oracle database files on the filer anyway, this is probably not a serious problem. Loss of the filer would temporarily bring down the database anyway.

In the event of a filer outage, the error message files that Oracle normally keeps in the `${ORACLE_HOME}/rdbms/logs` directory will not be accessible. Thus, there is a chance that information required to diagnose the problem might not be available. For this reason, you should relocate these files onto the local hard drive of the Oracle8 server machine or onto a separate NetApp filer.

A NetApp filer can store a common set of Oracle executables that are used by multiple Oracle database servers. In the process, you can save the disk space required to support multiple separate installations of Oracle.

To avoid a situation where the Oracle diagnostic files are inaccessible due to a failure on the filer, it is advisable to store the Oracle alert and trace files in a local disk directory. To do this, for each Oracle instance that will use Oracle executables on a filer, follow these steps:

1. Shut down the Oracle instance (if it is running).
2. Create a local directory where the alert and trace files will reside (e.g., `/opt/oracle/log`).
3. Modify the `BACKGROUND_DUMP_DEST` parameter in the instance's `init.ora` file to point to the local disk directory. (This parameter controls the target directory for the alert and trace files.)
4. Back up any existing alert and trace files that might exist in the `$ORACLE_HOME/rdbms/log` directory on the filer.
5. Start the instance again.

At this point, the alert and trace files for this instance will reside on local disk, not on the filer. Thus, a failure on the filer will not make these important diagnostic files inaccessible.

Please note the installation details such as database name, SID, and password created. Before installation, we suggest that you create a VLD of sufficient size to be used as local disk for DBMS software/database installation. If you have already created VLDs, you can proceed with the installation.

6.2.2) Installing Microsoft SQL Server

Installing SQL Server on a filer is a simple procedure. It is worth noting the following information.

Login as the administrator local to that server.

Make sure that the administrator has the User Right Log on as a service from within the User Rights dialog in User Manager for Domains.

Set the MSSQLServer and SQLServerAgent services to log in as the administrator using the Services applet in Control Panel. Stop and start the MSSQLServer and SQLServerAgent services to force them to log on as the administrator.

Create a share on the filer with security set to Everyone Full Control.

SQL Server 7.0 supports the use of UNC paths to describe the location of the data files. UNC (Uniform Naming Convention) paths allow you to describe the location of any file or resource on a Windows network by using a combination of:

The machine name

The share name

The logical path to the resource within the share

The advantage of UNC paths lies in their comparison to mapped drives. While drive mapping occurs late in the boot process, UNC paths work very early in the process. This is because UNC paths are used as remote boot paths for Windows software-based diskless workstations. Thus, UNC paths must work almost from the beginning in the Windows boot process. For this reason, UNC paths are more reliable than mapped drives.

For example, the following is the UNC path to a SQL Server database primary data file on a filer whose host name is `FILER1`. The host name can also be replaced with its network IP address.

```
\\FILER\HOME\sql2000\data\testdata.mdf
```

For further details, we suggest to reading our technical report TR3082.

http://www.netapp.com/tech_library/3082.html

An example of databases and their properties is shown in the figure below.

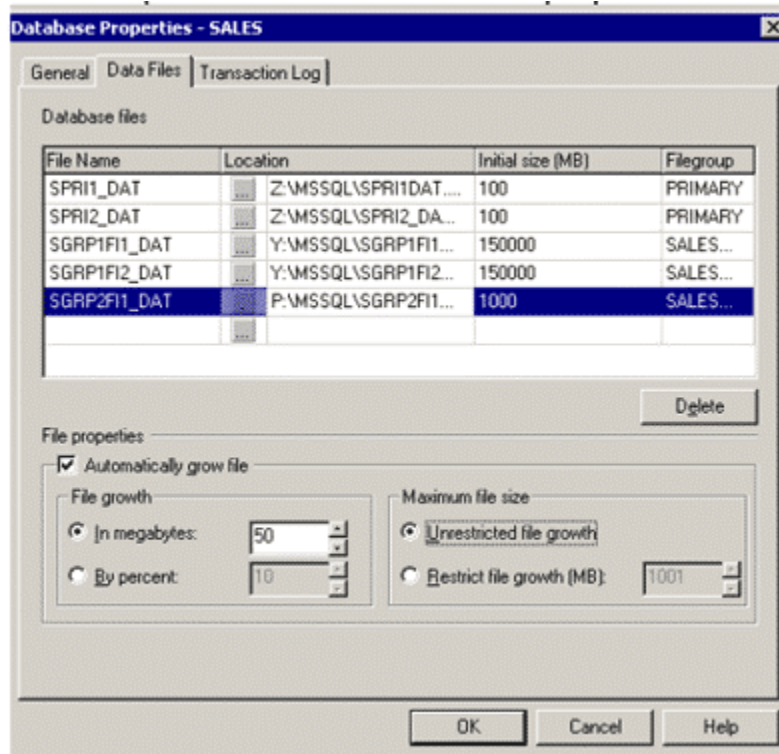


Figure 15 Database Properties for SALES Database

6.2.3) Installing Informix Software

As mentioned earlier, ESRI ArcSDE supports different database products. On our test machine, we installed Oracle 8.1.7 and Oracle 9.0.1.2 database instances. In this paper, we do not intend to explain the detailed steps for database product installation. For details, please refer to ESRI and Oracle product installation guides. For integrating any database software with filers, please refer to our technical library. http://www.netapp.com/tech_library/

6.2.4) Installation of ArcSDE

After creating virtual disks (treated as local SCSI disks), we installed ArcGIS software and database servers. For ArcSDE installation, we use iSCSI LUNs with SnapDrive 3.0 software to manage the LUNs. As a Pre Installation check, make sure that virtual disks are available to install the ArcSDE service. Select the right ArcSDE software media and start the installation. Before installing ArcSDE, select the installation directory. For this purpose, use a virtual disk created during the Create virtual

disk process using the SnapDrive program. This starts with SDE for the database setup program and the installation wizard, and guides you through the installation.

After the installation of ArcSDE, you need to complete the post-installation task. For example, in an Oracle configuration this helps you to create an SDE tablespace and user, set up the SDE repository, and create the ArcSDE service. Note the SDE user name and password, and make sure installation is successful.

6.2.5) Space Required for Database

How much space is needed for the database files?

The initial space and future growth requirements have to be known at the time of volume(s) creation. Of more concern at this point is the size of the database files. It is necessary to first clarify the difference between extending a database tablespace and expanding a database file. A fixed growth size is easier to control and can always be changed through Enterprise Manager.

6.2.6) Example of Creating a Database with One File Group Using SQL Server

The following CREATE DATABASE statement creates a SALE database consisting of one data file in the PRIMARY file group. Both the data and log files can grow unlimited until the device is full. Both the data file and the log file are located on the 'M' device.

```
CREATE DATABASE SALE
ON PRIMARY
( NAME = TEST_DATA,
  FILENAME = 'M:\SALE_DB\SALE_DATA.MDF',
  SIZE = 30GB,
  FILEGROWTH = 1 GB)
LOG ON
( NAME = SALE_LOG,
  FILENAME = 'M:\SALE_DB\SALE_LOG.LDF',
  SIZE = 30MB,
  FILEGROWTH = 1 MB)
```

7) Migrating Databases

When a NetApp filer is introduced into an existing database server environment, databases have to be migrated from their current storage locations to the new filer storage. It is not complicated to migrate system databases to the new storage, but different methodologies have to be used for each system database. For detailed information refer to TR3247: *Guide to Integrate SQL Server with SnapDrive v 2.0.1* written by Gerson Finlev of Network Appliance. http://www.netapp.com/tech_library/3247.html

8) Using Snapshot with Database Server

Snapshot copies have to be created and restored through the SnapDrive Manager. The methodology used when creating Snapshot backups depends on several factors: number of databases, whether the database is stored on one or several devices over how many volumes, or whether the database engine can be shut down for a short duration. Snapshot creation depends on SnapDrive functionality, which has to be well understood. Note that using iSCSI LUNs, the first Snapshot copy takes an equal amount of the space on the LUN, and additional copies take only the changes from the previous Snapshot

copy.

A Snapshot backup using a LUN can only be used as point-in-time backup and recovery.

8.1) Creating Snapshot Copies with SnapDrive 3.0

Creating Snapshot copies depends on SnapDrive Manager functionality. Using SnapDrive, you can take a Snapshot duplicate of a LUN, whereas SnapDrive restores a *LUN device* and not the volume. It is essential to understand the effect this has when a database utilizes several devices in one single volume. When a Snapshot copy is created on a device in a volume, all other devices will be duplicated at the same time. This is not a big issue but it is important to understand that all devices used by a single database have to be restored at the same time; if one device is not restored, the database will probably not be able to go online because of inconsistencies.

8.2) DBMS Server Crash Recovery

For details, refer to our **TR3247** Technical Report: *Guide to Integrate SQL Server with SnapDrive v 2.0.1 and NetApp Filer*. http://www.netapp.com/tech_library/3247.html

8.3) Snapshot Copies of Databases in a Single Volume

SnapDrive 3.0 allows the user to take a Snapshot of an iSCSI or FCP LUN. SnapDrive 3.0 will flush all the metadata before a Snapshot copy is taken. Using SnapDrive 3.0, a single Snapshot duplicate of a LUN is taken at a time. If a Snapshot is taken without SnapDrive 3.0, make sure to flush the metadata into the Windows server.

8.4) Snapshot Batch Program of Online Database(s)

You can create a batch program to create Snapshot backups of online databases.

8.5) CRE_SNAPSHOT.BAT for SQL Server

It is easiest to control and to understand if one batch job is created per database.

When several devices utilize one volume, then only one single Snapshot copy should be created per volume.

```
isql -Usa -P -S<SQL Server> -i cr_checkpoint.sql

( SDCLI snap delete -d <drive letter 1> -s <oldest_data_1>
SDCLI snap rename -d <mount point 1> -o <oldest-1_data_1> -n <oldest_data_1 >
SDCLI snap rename -d <mount point 1> -o <oldest-2_data_1> -n <oldest-1_data_1 >
SDCLI snap rename -d <mount point 1> -o <oldest-3_data_1> -n <oldest-2_data_1>

// Continue all current Snapshots on all mount points

SDCLI snap delete -d <mount point 2> -s <oldest_data_2>
SDCLI snap rename -d <mount point 2> -o <oldest-1_data_2> -n <oldest_data_2>
SDCLI snap rename -d <mount point 2> -o <oldest-2_data_2> -n <oldest-1_data_2>
SDCLI snap rename -d <mount point 2> -o <oldest-3_data_2> -n <oldest-2_data_2>
```

```

SDCLI snap delete -d <mount point 3> -s <oldest_log>
SDCLI snap rename -d <mount point 3> -o <oldest-1_log> -n <oldest_log>

SDCLI snap rename -d <mount point 3> -o <oldest-2_log> -n <oldest-1_log>
SDCLI snap rename -d <mount point 3> -o <oldest-3_log> -n <oldest-2_log>

// Continue creating new Snapshots of all data files and than

SDCLI snap create -s <newest_data1> -D <mount point 1>
SDCLI snap create -s <newest_data2> -D <mount point 2>

SDCLI snap create -s <newest_data2> -D <mount point 3>

```

9) SnapRestore

Before a database can be restored, it has to be taken offline. SnapRestore has to be executed from the SnapDrive Manager or SDCLI.

SnapDrive restores single files, and in this case, the single file is the device that contains the database files. The single file that is restored can be a LUN.

All devices utilized by the database to be restored have to be restored before the database is taken online. If one of the devices utilized by the database is not restored, then the database will probably be prevented from coming online.

9.1) Using Writable Snapshot Backups

Snapshot backups can be mounted in writable mode. Backups mounted in writable mode are excellent for testing or other short-term use. However, they should never be used as the base for a new database image. Never create copies of a writable Snapshot copy. It is easy to run out of free space if Snapshot backups are created using writable copies.

10) Disaster Recovery and High Availability

SnapDrive 3 supports asynchronous SnapMirror and local synchronous mirroring. SnapMirror is supported by SnapDrive functionality, whereas local synchronous mirroring is filer-level mirroring transparent to SnapDrive.

10.1) Asynchronous Mirroring for Disaster Recovery

Asynchronous SnapMirror is an excellent solution for disaster recovery of DBMS Server databases. The recommendation is to use rolling Snapshot copies for disaster recovery. Each time a Snapshot copy has been created, SnapDrive will send it to the mirror if SnapMirror is turned on. Snapshot mirrors can be used in read-only mode.

10.2) Local Synchronous Mirroring for High Availability

Local synchronous mirror (LSM) is single-filer-level mirroring and is transparent for SnapDrive and DBMS Server. LSM creates a volume consisting of two duplexes, where each duplex is an exact copy of the other at any point in time. LSM is one component of a high-availability system environment, which makes it very unlikely that a volume will fail because of multiple disk failure in a duplex RAID group.

10.3) Clustered Fail-over (CFO)

Filer Clustered Fail-over (CFO) and local synchronous mirroring (LSM) are important high-availability components. CFO safeguards against single filer-head failure, and LSM safeguards against double disk failure in a single RAID group. Both CFO and LSM are transparent for both SQL Server and SnapDrive.

11) References

Guide to Integrating SQL Server with SnapDrive™ V.2.0.1

http://www.netapp.com/tech_library/3247.html

File System Design for an NFS File Server Appliance

http://www.netapp.com/tech_library/3002.html#File System

NetApp Products and Technology

<http://www.netapp.com/products/software.html>

Clustered Fail-over High-Availability Solution

<http://www.netapp.com/products/filer/clustered.html>



Network Appliance, Inc.

495 East Java Drive
Sunnyvale, CA 94089
www.netapp.com

Network Appliance, Inc. Company Confidential and Proprietary.

© 2002 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, FAServer, FilerView, NetCache, SecureShare, SnapManager, SnapMirror, SnapRestore, and WAFL are registered trademarks and Network Appliance, ApplianceWatch, BareMetal, Camera-to-Viewer, Center-to-Edge, ContentDirector, ContentFabric, ContentReporter, DataFabric, Data ONTAP, EdgeFiler, HyperSAN, InfoFabric, MultiStore, NearStore, NetApp Availability Assurance, NetApp ProTech Expert, NOW, NOW (NetApp on the Web), RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, Smart SAN, SnapCache, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapMigrator, Snapshot, SnapSuite, SnapVault, SohoCache, SohoFiler. The evolution of storage, Vfiler, and Web Filer are trademarks of Network Appliance, Inc. in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.