



# iSCSI Software Boot for Windows with NetApp Storage

Deepak S N, NetApp, Inc.

January, 2008 | TR-3644

## Abstract

This technical document demonstrates to the storage administrator audience how to prepare, configure, and boot servers using NetApp storage systems. This document enumerates how to configure servers with iSCSI enabled network adapters to achieve remote boot and provides detailed steps for configuring iSCSI software boot using NetApp storage systems.

# Table of Contents

|           |   |           |
|-----------|---|-----------|
| <b>1</b>  | <b>Introduction</b>   | <b>3</b>  |
| <b>2</b>  | <b>Intended Audience and Assumptions</b>                      | <b>3</b>  |
| <b>3</b>  | <b>Working Methodologies</b>                                  | <b>3</b>  |
| <b>4</b>  | <b>Intel Network Adapter Configuration</b>                    | <b>4</b>  |
| <b>5</b>  | <b>iSCSI Adapter Configuration</b>                            | <b>5</b>  |
| 5.1       | Software Sources  | 5         |
| 5.2       | Configuring Host-Side iSCSI Flash Boot                        | 5         |
| 5.3       | Performing iSCSI Flash Update Using ISCSIUTL Tool in DOS Mode | 11        |
| 5.4       | Performing Boot Configuration on iSCSI Card on Server         | 11        |
| 5.4.1     | Configuring iSCSI Initiator with Static Setting               | 13        |
| 5.4.2     | Configuring iSCSI Initiator DHCP Setting                      | 14        |
| 5.5       | Network Configuration   | 18        |
| <b>6</b>  | <b>Storage Configuration</b>                                  | <b>18</b> |
| <b>7</b>  | <b>Microsoft Windows Setup for iSCSI Boot</b>                 | <b>23</b> |
| 7.1       | Booting Windows Server 2003 without Local Disk                | 23        |
| 7.1.1     | Installing the Windows Operating System                       | 23        |
| 7.1.2     | Preparing iSCSI Remote Boot                                   | 24        |
| 7.2       | Booting Windows Server 2003 with Local Disk                   | 24        |
| 7.2.1     | Configuring sysprep on Windows Server                         | 25        |
| <b>8</b>  | <b>Cloning</b>  | <b>25</b> |
| 8.1       | Purpose   | 25        |
| 8.2       | Preparing the Image   | 25        |
| 8.3       | NetApp FlexClone Configuration                                | 26        |
| 8.3.1     | Performing Cloning on Volume and LUN Level on the Storage     | 26        |
| <b>9</b>  | <b>Best Practices</b>   | <b>26</b> |
| 9.1       | Volume Layout   | 26        |
| 9.2       | Operating System Page File Placement                          | 27        |
| 9.3       | Page File Configuration                                       | 27        |
| 9.3.1     | Page File Best Practice                                       | 27        |
| 9.4       | High Availability   | 28        |
| <b>10</b> | <b>References</b>   | <b>29</b> |
| <b>11</b> | <b>Summary</b>  | <b>29</b> |

# 1 Introduction

This document provides detailed instructions on configuring systems for booting from a Logical Unit Number (LUN) in an Internet Small Computer System Interface (iSCSI) storage solution that is configured on Gigabit Ethernet adapter supplied by Intel and Broadcom. This paper mainly describes how to configure Intel Gigabit network adapters for iSCSI boot; however, the same can be followed for configuring Broadcom NetXtreme adapters.

The following sections describe how to prepare host and storage to perform remote iSCSI boot with and without local disk attached to the server.

The following topics are discussed in detail:

- Steps involved in configuration for iSCSI firmware on network adapter family
- Sources to look for software required to perform iSCSI remote boot on NetApp storage
- Detailed configuration for flash iSCSI firmware on iSCSI-capable network cards
- How to prepare storage configuration for performing remote iSCSI boot
- Configuring boot from iSCSI using Intel® PRO network adapters on a system without local disk
- Configuring boot from iSCSI using Intel Pro network adapter on a system with local disk

## 2 Intended Audience and Assumptions

This paper is intended for system and storage architects for designing NetApp storage appliances into IT environments. For the concepts, methods, and procedures in this document to be useful for readers, the following assumption is made.

- The reader has a minimum general knowledge of NetApp hardware and software solutions, particularly in the area of Fibre Channel and iSCSI.

## 3 Working Methodologies

iSCSI boot is a process whereby the Operating System is initialized from a storage disk array across a storage area network (SAN) rather than from the locally attached hard disk drive. The NetApp approach is more cost-effective and delivers the same benefits of any SAN, including better manageability, higher availability, and lower cost of ownership, while removing the high acquisition costs usually associated with deploying SANs by combining these networking and storage functions in a single controller device. Servers equipped with standard Gigabit network adapters, which can be configured to perform iSCSI offloading chip technology, are now able to connect to SANs with complete iSCSI functionality, including boot capabilities under Windows®. This technology eliminates the high up-front acquisition costs of adding storage networking to a server by allowing IT professionals to avoid having to purchase a server with a separate HBA controller preinstalled. In the past, IT professionals had to purchase separate controllers to perform simultaneous data and storage networking functions versus purchasing a server equipped with a standard network adapter capable of iSCSI software boot, designed to provide both functions in a single network device.

The working of iSCSI software boot can explained in the following steps:

1. The system starts up with the flashed iSCSI firmware on the network adapter for performing remote boot.
2. The iSCSI firmware has the following stacks, which are initiated during the boot phase to establish session with iSCSI target (storage) to obtain the LUN information over the TCP/IP network:

- The iSCSI stack, built into the firmware to carry SCSI commands over the TCP/IP network. In the initial boot phase it uses information supplied in the configuration.
  - For example, initiator/target IP address, LUN ID, and so on to establish an iSCSI session, which is required for booting from the storage.
  - The TCP/IP stack, the medium that carries iSCSI packets between Initiator and target and vice versa. It is a connection-oriented protocol.
3. System BIOS comes up and contacts DHCP for IP address, gateway, iSCSI target name, LUN ID, and so on.
  4. BIOS logs into iSCSI target and transfers off to BIOS boot loader code.
  5. BIOS boot loader code reads the first sector of the disk (partition table) and determines the active partition and starts loading NT loader to the system memory.
  6. The NT loader reads the kernel and all boot drivers from the iSCSI LUN and transfers control to the Windows OS.
  7. Windows starts up and initializes the boot drivers, including the NIC, TCP, and iSCSI drivers.
  8. The iSCSI software initiator takes over boot operation and continues to load from the iSCSI disk.

## 4 Intel Network Adapter Configuration

This paper provides information about Intel iSCSI boot that is designed for the Gigabit family of PCI-Express Server adapters. Intel iSCSI boot provides the capability to boot from a remote iSCSI disk volume located on an iSCSI-based storage area network (SAN). This paper covers hardware and software installation and setup procedures for installing and booting the system using these adapters. Broadcom has a specific iSCSI solution for some of its OEM partners. The iSCSI firmware and drivers for the network cards and Operating Systems are supported by the OEM vendors. IBM, Dell and HP ship their servers with onboard NetXtreme Gigabit Ethernet adapters, which can be flashed with the iSCSI firmware. However, they may already have been programmed by the server vendor in the factory. The iSCSI software installation and configuration for the Broadcom and NetExtreme are similar to the Intel iSCSI software configuration.

The Intel SCSI boot solution is able to use two different mechanisms to obtain appropriate configuration information to perform an iSCSI boot:

- **Static configuration:** In a Static configuration the client system is configured with a Static IP address (including Net mask and Gateway), an assigned initiator IQN, and targets. These values are programmed into the NVRAM associated with either the onboard NIC or the Gigabit Ethernet adapter and can only be changed through MBA.
- **Dynamic configuration:** In a Dynamic configuration, either the onboard NIC or the Gigabit Ethernet adapter acquires its IP address, initiator IQN, and target IQN, using a DHCP service.

## 5 iSCSI Adapter Configuration

### 5.1 Software Sources

The standard Intel Gigabit Network Adapters capable of performing iSCSI software boot have to be flashed with iSCSI firmware to provide support for software boot functionality.

The Intel iSCSI Network Adapter Flash Utility (iSCSIUTL), a DOS utility for programming Intel iSCSI boot firmware on iSCSI boot-capable adapters, can be extracted by installing ISboot.exe which is available at <http://www.intel.com/support/network/adapter/pro100/sb/cs-023748.htm>.

**Note:** ISBoot.exe should be downloaded for a specific Intel Gigabit card family.

To install, do as follows:

1. Click ISBoot.exe to extract the iSCSIUTL.exe tool for upgrading iSCSI firmware on Intel adapters.
2. Accept the license agreement.  
The files will be extracted and copied to the default location, that is, the c:\Intel12.2 folder.
3. To change the location, click the Change button and provide an alternate path.
4. The extracted folder contains necessary iSCSI drivers for different Windows Operating Systems (32 or 64 bit), which can be found in the Win32 folder.
5. The iSCSIUTL tool used to flash iSCSI firmware on all Intel boot-capable Gigabit Network Adapters can be found in the iSCSIUTL folder.

### 5.2 Configuring Host-Side iSCSI Flash Boot

1. Create a DOS bootable floppy from the Windows XP system and copy the iSCSIUt1.exe file located in the C:\Intel12.2\iSCSIUt1\DOS extracted folder to floppy disk.
2. Boot the server with the DOS bootable floppy created in Step 1.

After booting with the DOS bootable floppy, the iSCSIUTL tool can be invoked with the following different options:

[General Options](#)

[Flash Programming Options](#)

[Configuration Options](#)

[Adapter Selection Options](#)

[Adapter Flash Enable/Disable Options](#)

#### A. General Options

These options provide information on help, exit codes, and versions when invoked with iSCSIUTL.

| General Options |   |
|-----------------|---|
| Parameter       | Description   |
| -HELP or -?     | Displays command line help and exits. When -HELP is used, all other parameters are ignored. |

|                      |   |
|----------------------|---|
| -EXITCODES           | Displays exit codes returned to the OS by iSCSIUtl and exits. When -EXITCODES are set, all other parameters except -HELP are ignored.                               |
| -IMAGEVERSION or -IV | Displays the version number of the iSCSI Boot image stored within iSCSIUtl. When -IMAGEVERSION is set, all other parameters except -EXITCODES and -HELP are ignored |

**Figure 1) General Options**

```
A:\ISCSIUTL -IV

Intel® iSCSI Boot Utility

iSCSIUTIL version 2.0.0.7.00

Copyright © 2002-2006 Intel Corporation

Intel ® iSCSI Boot image version
```

```
Gigabit iSCSI : 2.0.07
```

**B. Flash Programming Options**

These options are used to flash iSCSI boot image to all iSCSI-capable network adapters identified by the ISCSIUTL tool by -UP option. If it is invoked with -QUIET options it does not warn the user and silently updates the firmware.

| <b>Flash Programming Options</b> |   |
|----------------------------------|---|
| <b>Parameter</b>                 | <b>Description</b>  |
| -UP                              | Programs the Intel iSCSI boot ROM image into the flash ROM of all iSCSI boot-capable network adapters in the system. If programming with the -ALL command, the preexisting PXE/RPL flash ROM in the network adapters will be overwritten. If -QUIET is not specified, the user is prompted to write over the existing image before upgrading. |
| -QUIET                           | Runs iSCSIUtl without confirmation prompts. This option is used with the -UP or -RESTORE option.  |

Figure 2) Flash Programming Options (1)

```
A:\ISCSIUTL -up -quiet -all

Intel® iSCSI Boot Utility

iSCSIUTIL version 2.0.0.7.00

Copyright © 2002-2006 Intel Corporation

Programming flash on Port 1 with iSCSI Boot image Flash updated
successful.
```

| Port | Network Address | Series  | WOL | Boot ROM Type   | Version |
|------|-----------------|---------|-----|-----------------|---------|
| 1    | 001517382270    | Gigabit | YES | iSCSI (Enabled) | 2.0.07  |

Figure 3) Flash Programming Options (2)

```
A:\ISCSIUTL -Up -All

Intel® iSCSI Boot Utility

iSCSIUTIL version 2.0.0.7.00

Copyright © 2002-2006 Intel Corporation

Programming flash on Port 1 with iSCSI Boot image

An Intel® Boot Agent image has been detected on this device.
Programming this device with Intel® iSCSI Boot will erase this image.
Would like to continue? (Y)es or (N)o:
```

### C. Configuration Options

These options help to back up iSCSI configuration on flashed network adapters, for example, initiator and target setting, by using `-SAVE` option with ISCSIUTL. To restore the setting to the replaced card, use the `-RESTORE` option with ISCSIUTL. To back up the information, provide a file with `-FILE` option.

**Note:** When saving or restoring from the file, ensure that the CHAP setting is disabled in iSCSI configuration due to security reasons since the password and username will be in plain text when saving to a file.

| Configuration Options |   |
|-----------------------|---|
| Parameter             | Description   |
| -SAVE                 | Saves iSCSI boot settings in the EEPROM to a disk file on the selected iSCSI-capable port. Need to specify a filename using |

|          |  |
|----------|--|
|          | the -FILE parameter.   |
| -RESTORE | Restores previously saved iSCSI boot settings on the selected iSCSI boot-capable port. Need -FILE parameter to specify the saved data. |
| -FILE    | Specifies a filename for -SAVE or -RESTORE parameters by -FILE = [filename].   |

Figure 4) Configuration Options (1)

```
A:\ISCSIUTL -save -nic=1 -file=test
Intel® iSCSI Boot Utility
iSCSIUTIL version 2.0.0.7.00
Copyright © 2002-2006 Intel Corporation
Warning: For security reasons it is recommended that CHAP setting are
disabled when using save/restore.
Saving iSCSI setting on NIC 1 to file test...done

Port   Network Address  Series    WOL    Boot ROM Type    Version
===   =====  =====  =====  =====
1      001517382270  Gigabit   YES     iSCSI (Enabled)  2.0.07
```

Figure 5) Configuration Options (2)

```
A:\ISCSIUTL -restore -nic=1 -file=test
Intel® iSCSI Boot Utility
iSCSIUTIL version 2.0.0.7.00
Copyright © 2002-2006 Intel Corporation

Warning: The current configuration on port 1 is about to be over-
written.
Would like to continue? (Y)es or (N)o:
```

#### D. Adapter Selection Options

These options are used to select a single network port by using -NIC=XX[port number] option or all available ports on the system by providing -ALL with ISCSIUTL tool. The network port on which this operation is performed, can be identified by using BLINK option with ISCSIUTL tool.

| Adapter Selection Options |  |
|---------------------------|--|
| Parameter                 | Description  |
| -ALL                      | Selects all Intel iSCSI boot-capable adapters found in system.                                       |
| -NIC=XX                   | Selects a specific iSCSI boot-capable adapter in the system. Not valid with use of -UP parameter.    |
| -BLINK                    | Provides a method for identifying an adapter. Blinks the LED on the selected adapter for 10 seconds. |

Figure 6) Adapter Selection Options (1)

```
A:\ISCSIUTL -blink -nic=all
Intel® iSCSI Boot Utility
iSCSIUTIL version 2.0.0.7.00
Copyright © 2002-2006 Intel Corporation

The LED on port 1 is now blinking.
It will blink for about 10 seconds.

The LED on port 2 is now blinking.
It will blink for about 10 seconds.

Port      Network Address  Series    WOL    Boot ROM Type    Version
===      =====  =====  =====  =====
1         001517382270    Gigabit   YES     iSCSI (Enabled)   2.0.07
2         001517382271    Gigabit   YES     iSCSI(Enabled)   2.0.07
```

Figure 7) Adapter Selection Options (2)

```
A:\ISCSIUTL -blink -nic=1
Intel® iSCSI Boot Utility
iSCSIUTIL version 2.0.0.7.00
Copyright © 2002-2006 Intel Corporation

The LED on port 1 is now blinking.
It will blink for about 10 seconds.
```

| Port | Network Address | Series  | WOL   | Boot ROM Type   | Version |
|------|-----------------|---------|-------|-----------------|---------|
| ===  | =====           | =====   | ===== | =====           | =====   |
| 1    | 001517382270    | Gigabit | YES   | iSCSI (Enabled) | 2.0.07  |

### E. Adapter Flash Enable/Disable Options

These options are used to enable the iSCSI card for flash update by using `FLASHENABLE`. To disable the iSCSI card, use the `FLASHDISABLE` option with `ISCSIUTL` tool.

| Adapter Flash Enable/Disable Options           |   |
|--|---|
| Parameter                                      | Description   |
| <code>-FLASHENABLE</code> or <code>-FE</code>  | Enables iSCSI boot flash ROM on the selected network adapter port. At least one iSCSI boot-capable port in the system must be enabled in order to use iSCSI boot. |
| <code>-FLASHDISABLE</code> or <code>-FD</code> | Disables the iSCSI boot ROM on the selected network adapter port. The FLASH must be disabled on all iSCSI boot-capable ports.                                     |

Figure 8) Adapter Flash Enable/Disable Options (1)

```
A:\ISCSIUTL -FD -ALL
Intel® iSCSI Boot Utility
iSCSIUTIL version 2.0.0.7.00
Copyright © 2002-2006 Intel Corporation

Enabling boot ROM on port 1 .... Done
Reboot the system to enable the boot ROM on port 1.

Port Network Address Series WOL Boot ROM Type Version
=== =====
1 001517382270 Gigabit YES FLASH(Disabled) 2.0.07
```

Figure 9) Adapter Flash Enable/Disable Options (2)

```
A:\ISCSIUTL -FE -ALL
Intel® iSCSI Boot Utility
iSCSIUTIL version 2.0.0.7.00
Copyright © 2002-2006 Intel Corporation
```

Enabling boot ROM on port 1 .... Done

Reboot the system to enable the boot ROM on port 1.

| Port | Network Address | Series  | WOL | Boot ROM Type  | Version |
|------|-----------------|---------|-----|----------------|---------|
| 1    | 001517382270    | Gigabit | YES | FLASH(Enabled) | 2.0.07  |

### 5.3 Performing iSCSI Flash Update Using ISCSIUTL Tool in DOS Mode

1. Check for iSCSI-capable cards on your system by running:

```
ISCSIUTL -ALL
```

If the O/P is none, then check for network adapters qualified for iSCSI remote boot or not.

2. Before performing flash update on the network adapter, check if all the network ports are flash enabled or not by running one of the following commands:

```
ISCSIUTL -FLASHENABLE -ALL
```

```
ISCSIUTL -FE -ALL
```

```
ISCSIUTL -FE -NIC=XX
```

You can also flash enable on selected network adapters by typing:

```
-NIC=xx[port number].
```

3. To check for which network adapter to perform flash update, run:

```
ISCSIUTL -BLINK -NIC=XX[ network port number]
```

The selected network adapter blinks on the system.

4. To program iSCSI-enabled network adapter, run:

```
ISCSIUTL -UP -ALL -QUIET
```

If `-QUIET` option is not supplied, then network adapters will be programmed silently without warning the user.

5. After iSCSI firmware is updated, the system needs to be rebooted for EEPROM to be enabled with new firmware.
6. Once the system is rebooted, you can configure iSCSI setting by pressing `CTRL-D` when system is booting.

### 5.4 Performing Boot Configuration on iSCSI Card on Server

1. Press `CTRL - D` keys at the system startup to configure iSCSI boot configuration on the card ([Figure 10](#)).

**Figure 10) Performing Boot Configuration- Step 1**

Intel® iSCSI Boot version 2.0.07

Copyright © 2002-2006 Intel Corporation. All rights

Press ESC Key to skip iSCSI boot initialization.

```
Press <Ctrl-D> to run setup...
```

**Figure 11) Performing Boot Configuration (1)**

```
=====iSCSI Port Selection=====
MAC: 00151729FF18 Dev : 105E Loc:6:0:0 BOOT: DISABLED
MAC: 00151729FF19 Dev : 105E Loc:6:0:0 BOOT: DISABLED
=====
```

2. Enable the iSCSI NIC port for configuring initiator and target information by pressing **P** for primary boot or **S** for secondary boot if primary fails ([Figure 12](#)).

**Figure 12) Performing Boot Configuration- Step 2**

```
=====iSCSI Port Selection=====
MAC: 00151729FF18 Dev : 105E Loc:6:0:0 BOOT: PRIMARY
MAC: 00151729FF19 Dev : 105E Loc:6:0:0 BOOT: SECONDARY
=====
```

3. Press the **Enter** key by selecting primary card to invoke port configuration for specifying information.

**Figure 13) Performing Boot Configuration- Step 3**

```
=====iSCSI Port Selection=====
          iSCSI    Boot    Configuration
          iSCSI    CHAP    Configuration
          Discard changes and Exit
          Save changes and Exit
=====
```

4. Select iSCSI boot configuration to provide initiator and target information via DHCP or by static entry. If you choose to provide manually, then below are the parameters that you need to define (([Figure 13](#))).

## 5.4.1 Configuring iSCSI Initiator with Static Setting

### 1. Initiator Name: `iqn.1991-05.com.microsoft:hostname`

Here initiator name, which is IQN (iSCSI qualified name), is a standard defined by T11 body.

**Note:** The same initiator name must be defined in the storage too when creating an igroup. The storage configuration is defined in the storage configuration section below.

**Note:** The same IQN name must be used in the setting you have defined in the storage. If the IQN name does not match, you will get an error message when initiator is connecting to target and initial phase in remote boot.

**Note:** Acceptable characters for hostname are "a-zA-Z0-9" and special characters ". " Or "-".

### 2. Initiator IP Address: `XXX.XXX.XXX.XXX`

The IP address of the initiator used for participating in the iSCSI remote boot must be defined.

### 3. Initiator Subnet Mask: `XXX.XXX.XXX.XXX`

Based on the IP address, the allocation subnet mask must be defined.

### 4. Initiator Gateway: `XXX.XXX.XXX.XXX`

If the target IP address belongs to the same network subnet, then gateway address is not required. If the target (storage) is in different subnet, then router IP address must be defined in the gateway.

### 5. Target Name: `iqn.1992-08.com.netapp: storagename`

Here storage name is tagged with target IQN name defined by the T11 body. The target IQN name can be obtained in the storage, which is defined in the storage configuration section 6 below. The IQN name contains a date string, the domain of the naming authority, a unique string to identify the node.

For example, `iqn.1992-08.com.netapp:sn.5099945438`.

### 6. Target IP: `XXX.XXX.XXX.XXX`

Target IP address must be defined for establishing iSCSI session. If you have configured different IP addresses for different network ports on the storage you can provide the IP address which is enabled for iSCSI.

### 7. Target Port : `XXXX`

By default for iSCSI protocol

By default port 3260 to listen

### 8. Boot LUN: `XX`

The LUN is defined in the storage when mapping is made to an igroup. The LUN number and mapping configuration on the storage can be referenced in storage section below.

**Note:** The LUN number is defined in decimal value in storage end.

Figure 14) Configuring iSCSI Initiator with Static Setting

```
=====iSCSI Boot Configuration =====
      [ ] Use dynamic IP Configuration ( DHCP )
Initiator Name : iqn.1991-05.com.remote
Initiator IP :      10.10.10.11
Subnet Mask:      255.255.255.0
Gateway:          10.10.10.1
Vlan ID:

[ ] Use DHCP for iSCSI Target information
Target Name :      iqn.1992-08.com.netapp:sn.101173382
Target IP:         10.10.10.12
Target Port :      3260
Boot LUN:          1

[OK]                                     [ Cancel ]
=====
```

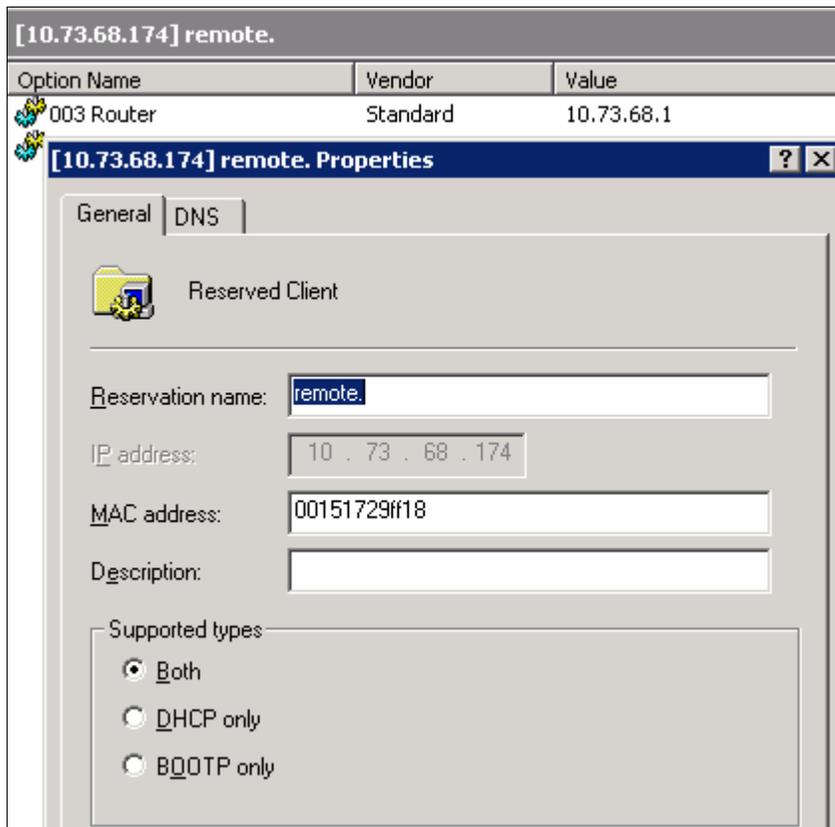
#### 5.4.2 Configuring iSCSI Initiator DHCP Setting

Dynamic Host Configuration Protocol (DHCP) is a protocol used by network computers (clients) to obtain IP address and other parameters such as subnet/gateway/DNS server dynamically instead of manually configuring on each device. The DHCP server ensures that all IP addresses are unique (IP address conflicts) releases only if lease is expired. There are different modes for allocating IP address. In this paper we are using dynamic mode where we can reserve IP address permanently assigned to client for a given MAC hardware address which is burned to each network adapter and is unique.

To define these configurations, do as follows:

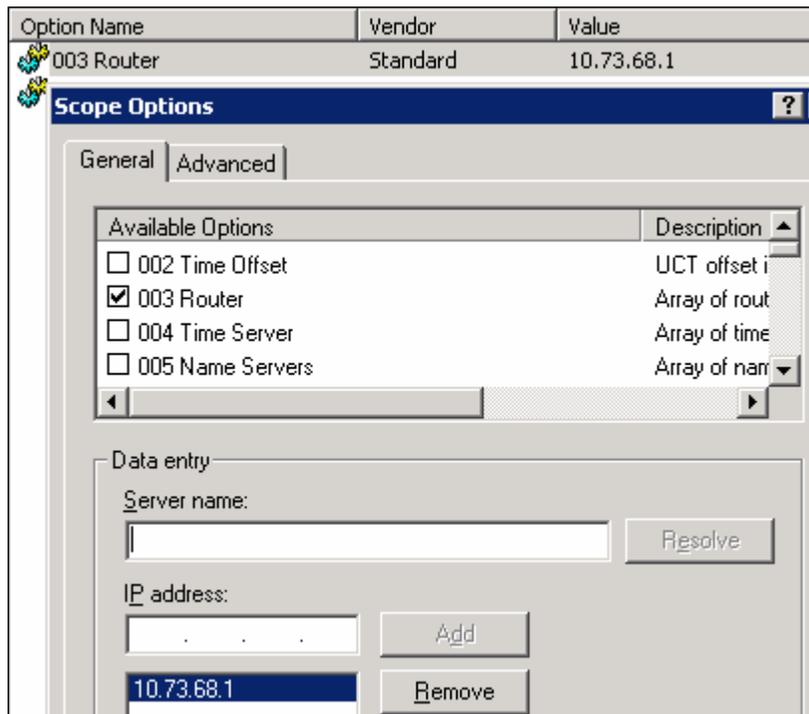
1. Windows 2003 DHCP server is used for providing initiator IP address, subnet mask, and gateway.
2. Configure initiator IP address reservation by providing the MAC address of the iSCSI network port, which can be obtained in iSCSI Port Configuration Menu ([Figure 15](#)).

Figure 15) iSCSI Port Configuration Menu



3. Define the Router IP address in the DHCP router list option 3 under Scope Options if your target and initiator belong to different network subnets ([Figure 16](#)).

**Figure 16) Scope Options**



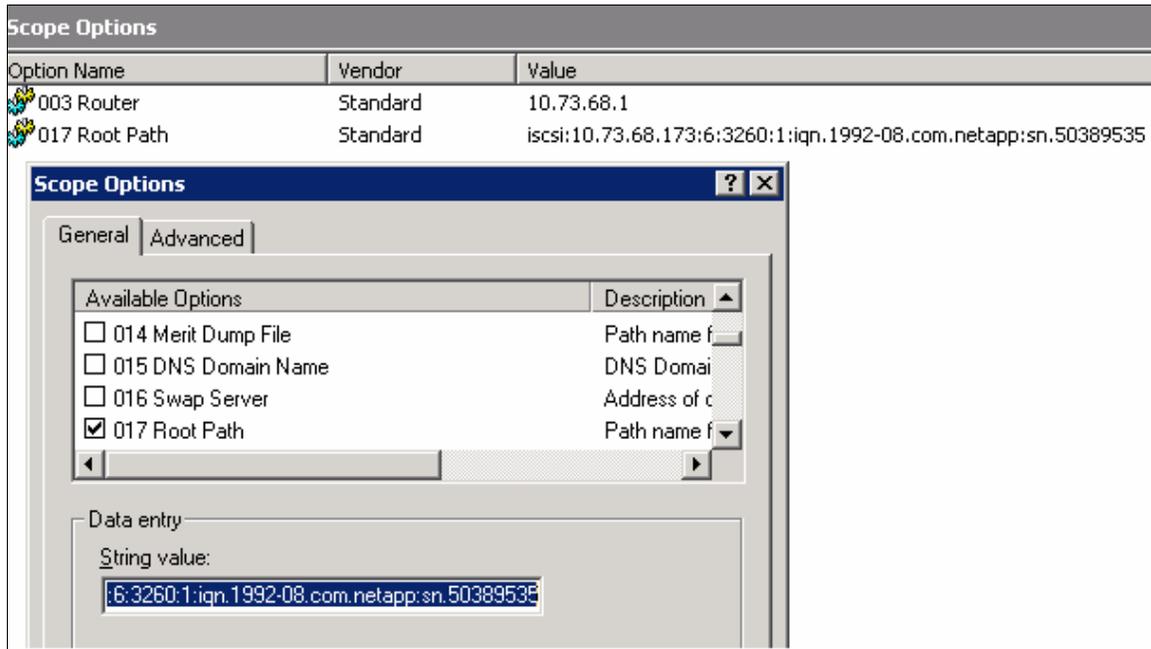
4. If you are using Dynamic Host Configuration Protocol (DHCP), the DHCP server must be configured to provide the iSCSI boot configurations to the iSCSI initiator. It is necessary to set up the DHCP server to specify root path option 17 to send iSCSI target information back to the iSCSI initiator ([Figure 17](#)).

```
iscsi:<storage name or IP  
address>:<protocol>:<port>:<LUN>:<targetname>
```

Where:

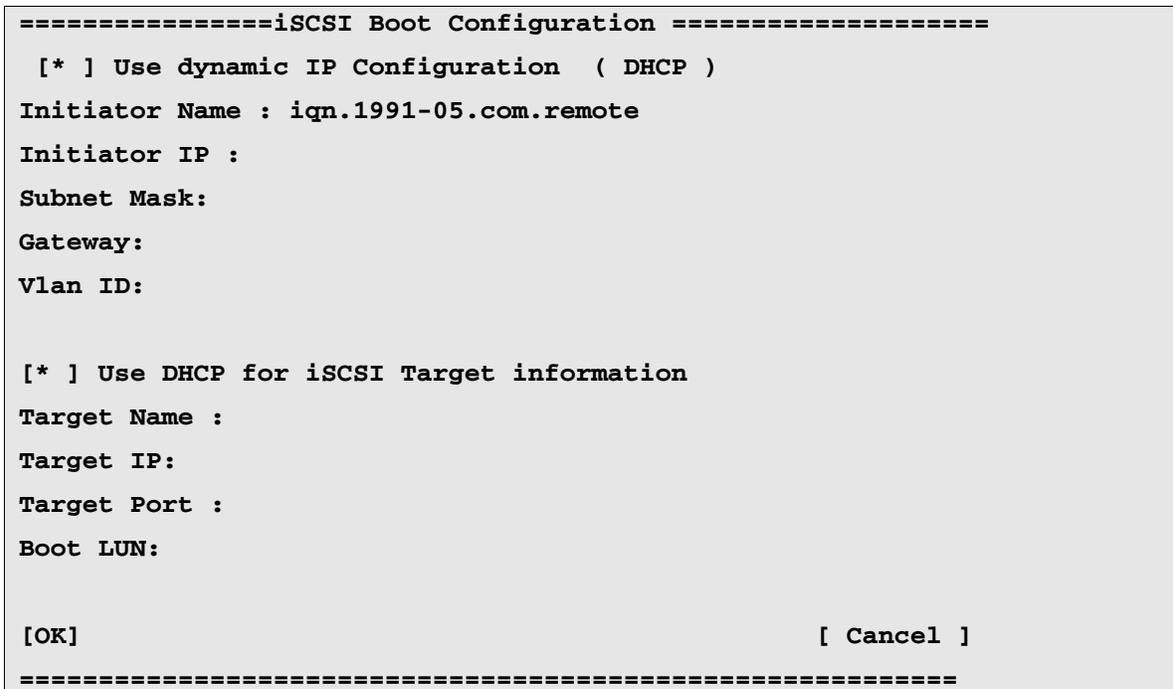
- <Storage name or IP address>: IQN or IP address of the storage.
- <Protocol>: Transportation protocol used by iSCSI is TCP, which is, 6. No other protocol is currently defined.
- <Port>: 3260 is defined as the default port for iSCSI; if it is blank then 3260 is used.
- <LUN>: The LUN number defined in the storage for access control.
- <Target Name>: The node name of the storage, which can be recorded by running a few commands defined in the storage section 6 below.

**Figure 17) Scope Options String Value**



- Configuring iSCSI network adapter for DHCP setting can be performed under boot configuration, which is invoked by pressing CTRL + D at boot prompt.

**Figure 18) iSCSI Boot Configuration**



## 5.5 NETWORK CONFIGURATION

iSCSI remote boot on Intel network adapter with boot version 2.0.07 has to be operated on network switch ports where spanning tree portfast has to be enabled. If not iSCSI enabled network adapters cannot login to the target and access the LUN exposed to the initiator.

## 6 Storage Configuration

This section covers the configuration components necessary to implement the iSCSI SAN boot environment. It is assumed storage is set up and accessible to host. Access the storage subsystem active controller through console via serial cable, or Telnet to the controller or use the FilerView® GUI

In this document both Telnet and FilerView sessions have been used for configuration.

To obtain the storage node name, type:

```
Storage> iscsi nodename
```

```
iscsi target nodename: iqn.1992-08.com.netapp:sn.101173498
```

Where:

<nodename> is the iSCSI target name

In FilerView, navigate to LUN section, click the iSCSI option. The Target node name appears under Manage Names.

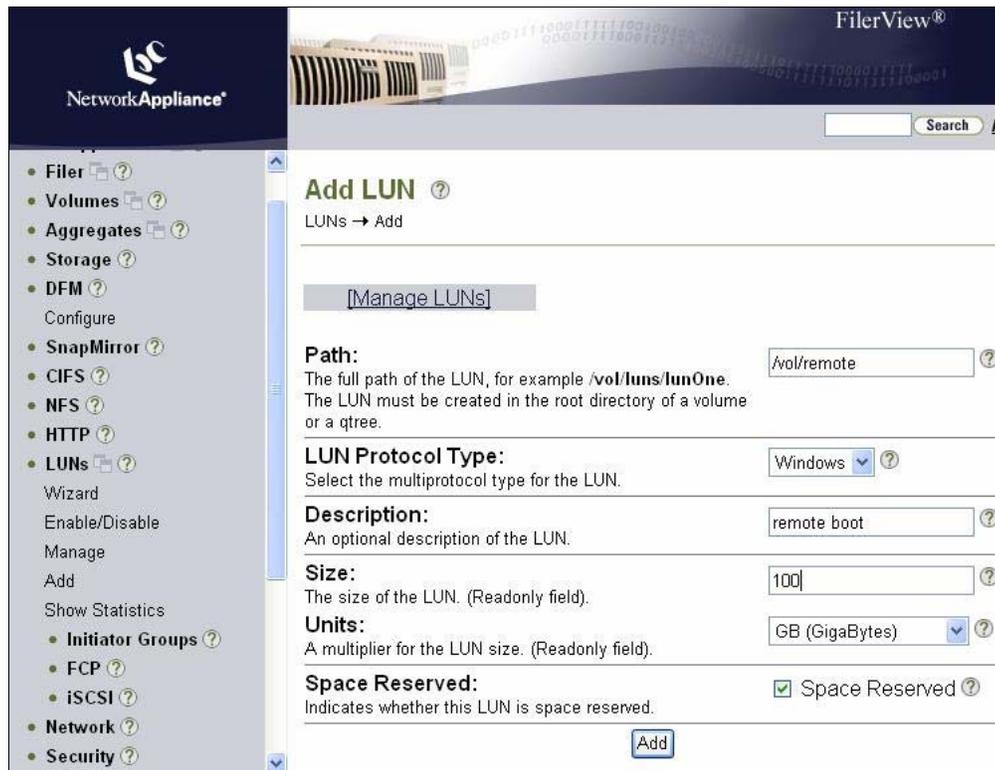
**Figure 19) Manage iSCSI Names**



- Create a boot LUN, size 100GB, type Windows, using the command below, or using FilerView you can navigate to LUN section, then click Add command and supply the information to create the LUN, which is used for remote boot purpose.

```
Storage> lun create -s 100g -t windows /vol/remotelun
```

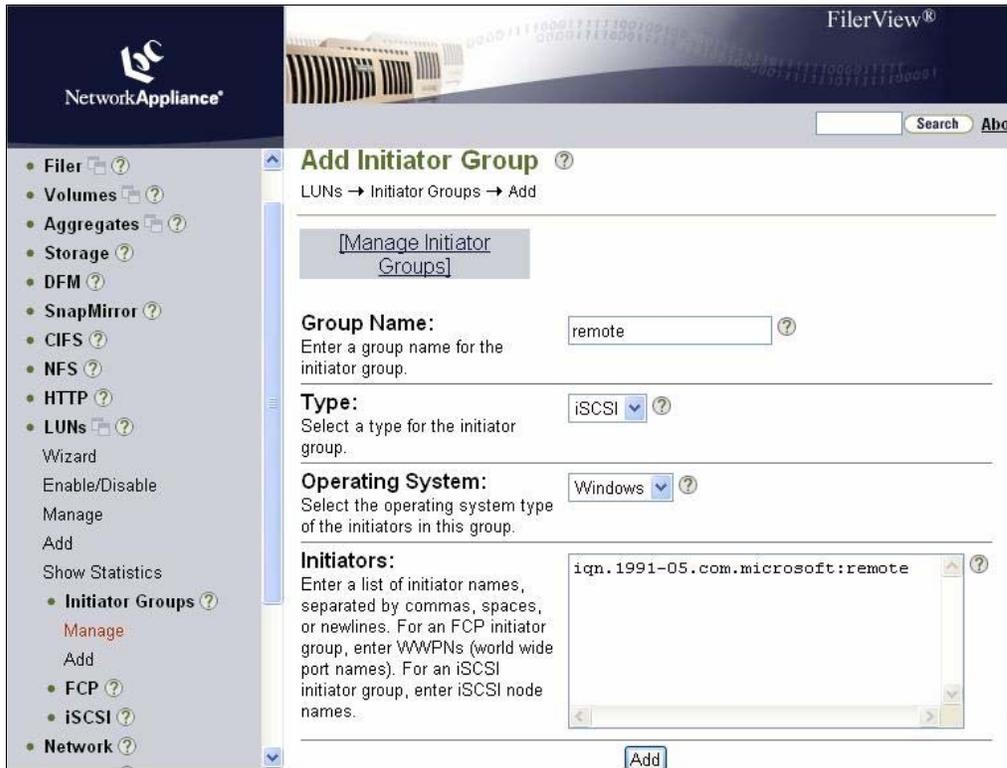
Figure 20) Add LUN



- Create igroup to define the initiator's IQN name for providing access to the LUN. This initiator IQN name is defined in iSCSI configuration, which is defined in iSCSI configuration sections 5.3.1 or 5.3.2 above. The command below is used; also, using FilerView you can navigate to LUN section, and under that you click Initiators Groups and supply the information.

```
Storage> igroup create -i -t windows remote iqn.1991-05.com.microsoft:remote
```

Figure 21) Add Initiator Group



- To access the LUN you need to map to the igroup where the initiator is defined with LUN ID. This LUN ID is to be defined in the iSCSI target configuration section. The below command is used; also, using FilerView you can navigate to LUN section, under that you click Manage and see which LUN has to be mapped. Click No Maps, which bring No Maps view. There you click Add Groups to Map, which brings up LUN Map Add Groups view. There you select the igroup name and provide the LUN ID.

```
Storage>lun map /vol/remotelun remote 1
```

Figure 22) Manage LUNs

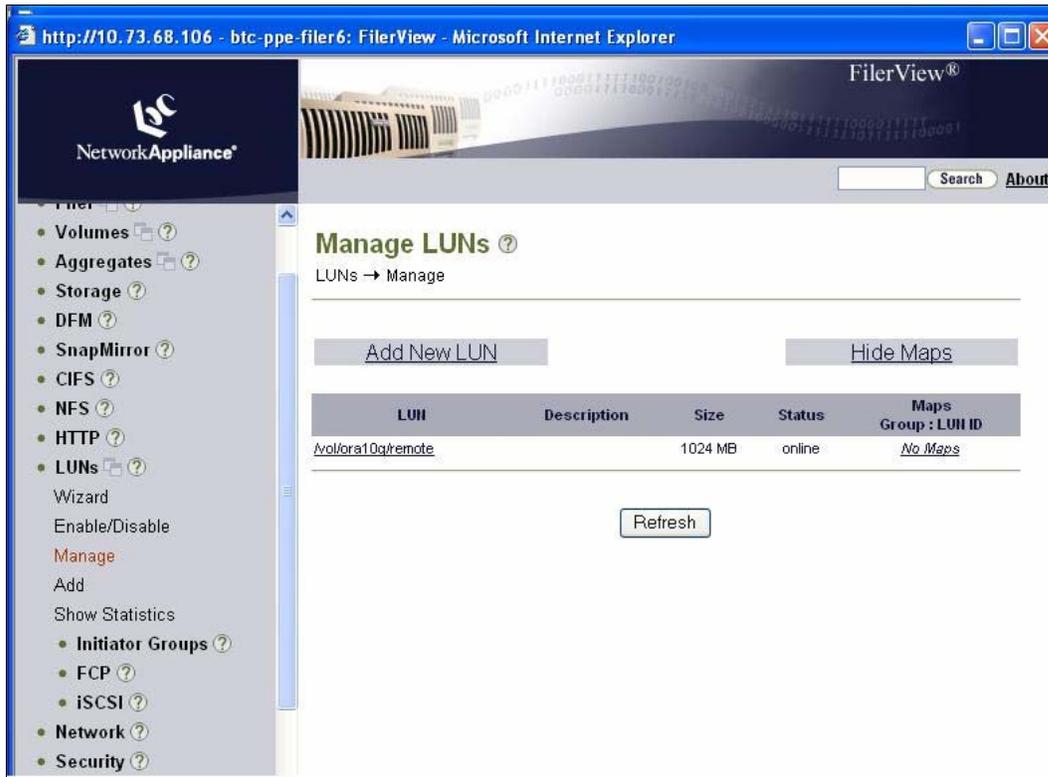


Figure 23) Map LUNs



Figure 24) Add Groups



Figure 25) Map LUNs



## 7 Microsoft Windows Setup for iSCSI Boot

The following preinstallation requirements are needed before installing the Windows Operating System:

- Follow installation steps to install the Intel iSCSI boot-capable adapters with Intel iSCSI boot firmware support.
- Create a disk target (LUN) on an available iSCSI target system. Configure this LUN to be accessible to host by defining igroup with host IQN and map to the LUN that is used for software boot.
- Make sure the iSCSI initiator system starts the iSCSI boot firmware. The firmware should be configured properly, be able to connect to iSCSI target, and detect the boot disk.
- Intel PRO/1000 PCIe network interface card (NIC).
- NetApp FilerView with FlexClone® (optional) for image replication.
- Any disk imaging software such as Symantec® Ghost.
- You will need Microsoft® iSCSI software initiator version 2.04 with integrated software boot support. This boot version of the initiator is distributed via <http://connect.microsoft.com>, invitation code: ms-8RR8-6K43.
- To set up Windows Server\* 2003 boot system when a local disk is available, follow the steps for installing Windows 2003 with local disk.
- To set up Windows Server 2003 boot system when a local disk is not available, follow the steps for installing Windows 2003 without local disk.
- To perform a diskless install directly to an iSCSI target using RIS, follow the steps for RIS installation of Windows 2003.
- To upgrade or install network drivers, OS service packs, or new hardware, follow the steps for upgrading/installing network drivers, OS service packs, or new hardware.

Note: If you find any Intel iSCSI remote boot issues you can refer the below link. There are known problems documented.

<http://www.intel.com/support/network/adapter/pro100/sb/CS-025776.htm#anchor3>

There are two setup methods which is defined below that can be followed to install the Windows Operating System using iSCSI-capable network adapter.

### 7.1 Booting Windows Server 2003 without Local Disk

#### 7.1.1 Installing the Windows Operating System

1. Copy `iSSetup.sys`, `iSSetup.inf`, and `txtsetup.oem` files from the Intel iSCSI boot media to a floppy disk.
2. Start the system and boot from the Windows Server 2003 Installation CD.
3. When prompted "Press F6 if you need to install a third-party SCSI or RAID Driver..." press F6.  
**Note:** The option to press F6 is available in the very start of the OS installation process and is easy to miss. You are only given three seconds to press F6. If you fail to press F6, restart the system and begin the OS installation again.
4. When prompted to install third-party storage device, press <s> and select "Intel (R) PRO/1000 iSCSI Setup Driver."
5. Press Enter and continue to complete the OS installation.

## 7.1.2 Preparing iSCSI Remote Boot

1. Install network driver for Intel PRO/1000 adapters. Install the driver via Device Manager and not via SetupBD or the `Install Drivers and Software` button in Autorun, which will use SetupBD.
2. You will not be able to install the network driver for the network port that is being used by the iSSetup driver, so install the network driver for other network port.
3. Once the Operating System is installed, ensure that Windows 2003 SP1 or SP1 R2 service pack is installed and then install the Microsoft Hotfix for Windows Server, KB902113 NDIS QFE. Windows 2003 SP2 contains the QFE.
4. Microsoft iSCSI initiator version 2.04 and later with integrated software boot can be downloaded from the Microsoft Web site at [www.connect.microsoft.com](http://www.connect.microsoft.com); log in using invitation code ms-8RR8-6K43.
5. Reboot the system.
6. Install Microsoft iSCSI software initiator version 2.04 with integrated software boot support. Select MPIO multipath support when more than one network ports are available for iSCSI boot. During installation when prompted select Intel iSCSI boot-capable Gigabit network adapter as a boot device. Intel Gigabit adapter is identified as "e1express" during the installation.
7. Reboot the system.
8. Confirm the remote disk is accessible and configurable in Windows disk management computer management. Note that you do not need to configure Microsoft iSCSI software initiator for Windows to be able to detect the iSCSI disk. The Microsoft iSCSI software initiator automatically retrieves the iSCSI configurations from the Gigabit network adapter iSCSI boot firmware.
9. Open the Microsoft iSCSI software initiator, go to the `Target` tab, and ensure it is loaded and connected to the iSCSI target.
10. Run the iSCSIApp to uninstall the Intel PRO/1000 iSCSI setup driver.
11. Reboot the system.
12. Install network driver for the port that was being used by Intel PRO/1000 iSCSI setup driver.
13. Run `iscsibcg.exe` with `/verify` and `/fix` command line parameters.
14. System is now ready for Windows iSCSI. You can use any third-party image tools to clone the local disk OS image to an iSCSI LUN presented by storage.

## 7.2 Booting Windows Server 2003 with Local Disk

1. If the local disk is not already connected, connect the hard drive temporarily to the system to be set up for diskless boot.
2. Create a disk target (LUN) on an available iSCSI target. Configure this LUN to be accessible to the MAC address of the NIC card of the host being booted.
3. Install Microsoft Windows 2003 SP1, SP1 R2, or SP2 Operating System on a temporarily connected local hard drive of the host system.
4. Install network drivers for Intel Gigabit network adapters that are to be used for iSCSI boot.
5. If you installed Windows 2003 SP1 or SP1 R2, then install the Microsoft Hotfix for Windows Server, KB902113 NDIS QFE. Windows 2003 SP2 contains the QFE.
6. Reboot the system.
7. Install Microsoft iSCSI software initiator version 2.04 with integrated software boot support.
  - a. Select MPIO Multipath support when more than one network port is available for iSCSI boot.
  - b. During installation when prompted select an Intel iSCSI boot-capable Gigabit network adapter as a boot device.
  - c. Select the checkbox to configure iSCSI boot.  
The Intel Gigabit network adapter is identified as `e1express` during the installation.
8. Microsoft iSCSI initiator version 2.04 and later with integrated software boot can be downloaded from the Microsoft Web site at [www.connect.microsoft.com](http://www.connect.microsoft.com); log in using invitation code ms-8RR8-6K43.

9. Reboot and confirm the remote disk is accessible and configurable in Windows disk management computer management.  
**Note:** You don't need to configure Microsoft iSCSI software initiator for Windows to be able to detect the iSCSI disk. Microsoft iSCSI software initiator automatically retrieves the iSCSI configurations from the Gigabit network adapter iSCSI boot firmware.

### 7.2.1 Configuring sysprep on Windows Server

1. Use sysprep to prepare the OS image for deployment.
2. Create a `sysprep` directory in the root of your system drive.
3. Locate the `SUPPORT\TOOLS\DEPLOY.CAB` archive on your Windows 2003 CD, extract files, and copy all the files within the archive to the `sysprep` directory.
4. Create the `Sysprep.inf` file using `setupmgr.exe` wizard and answer all the questions.
5. Edit the resulting `sysprep.inf` file by adding the `LegacyNic=1` option in the [Unattended] section, like so:  
[Unattended]  
OemSkipEula=Yes  
LegacyNic=1
6. Run `Sysprep`. Select `Reseal` and `Shutdown` options for preparing an OS image.
7. Shut down the machine and capture OS image to a local disk partition.
8. In order to image your disk to iSCSI LUN, you can use various third-party software: for example, Symantec Ghost, Winboot/i, and so on.
9. Transfer image from local hard drive to iSCSI target.
10. Remove local hard drive or keep local hard drive for page file use. Configure the system BIOS to boot from iSCSI disk and boot.

## 8 Cloning

### 8.1 Purpose

It would be cumbersome to install Windows and some assorted other software (Office suites, etc.) on dozens of machines. To circumvent this, It would be helpful to just install everything on one machine, configure it to your liking, and then clone it to the other machines. A direct cloning is almost never possible, because there's always something that needs to be different between machines (for example, machine names), but you can make an image of the original installation, dump it on a storage, and be presented with a mini-setup wizard that asks for those machine-specific things such as the hostname, workgroup/domain, administrator password, and so on.

### 8.2 Preparing the Image

To create initial image for performing cloning operation to other disks, do as follows:

1. For the image preparation we use a small command-line utility from Microsoft called `Sysprep`. This utility relies on the information declared in a text file called `sysprep.inf` to automate some of mini-setup's steps (for example, skip the EULA and/or the product serial number screens). You can obtain `sysprep` by downloading it from Microsoft's Web site.
2. You can use third-party software, for example, Symantec Ghost, WinBoot, and so on, which is designed to clone disk-disk or partition-partition on the system.
3. Once your primary image is created after cloning, it can be dumped in a central shared area on storage and later can be cloned using NetApp FlexClone seamlessly without using any third-party tools.
4. LUNs which are being cloned have to be backed up by Snapshot™.

## 8.3 NetApp FlexClone Configuration

You can leverage NetApp FlexClone to instantaneously create a volume-level clone. The cloning can be performed at storage end while the primary disk is accessed by the application for I/O operation by the end user. With LUN cloning capabilities present in the NetApp Data ONTAP® Operating System, it is possible to rapidly make clones of existing LUNs for provisioning to new servers. For example, with one golden boot image you can clone hundreds of LUNs to boot on servers, which reduces time spent to install the Operating System.

- Clone volumes are writable, point-in-time copies of a parent FlexVol® volume. They reside in the same aggregate as their parent volume. Changes made to the parent volume after the clone is created are not inherited by the clone.
- Clone volumes are fully functional volumes; you manage them using the `vol` command, just as you do the parent volume. Clone volumes themselves can be cloned.
- When you create clone volumes, LUNs in the parent volume are present in the clone, but they are not mapped, and they are offline. To bring the LUNs in the clone online, you must map them to igroups.
- When the LUNs in the parent volume are backed by Snapshot copies, the clone also inherits the Snapshot copies.

### 8.3.1 Performing Cloning on Volume and LUN Level on the Storage

- Volume-level cloning:

When volume-level cloning is created, all LUNs created under the volume will be cloned; however, all cloned LUNs are offline and unmapped and need online and create mapping to igroup in order to access. Parent volume and volume, which has to be cloned, have to be in the same aggregate.

To create/split volume-level cloning, type:

```
Storage*> vol clone create clonew1 -b /vol/new1
Storage*>vol clone split status clonew1
Storage*>vol clone split start clonew1
Storage*>vol clone split stop clonew1
```

- LUN-level cloning :

LUN-level cloning can be performed on individual LUNs instead of cloning the entire volume. When split, the cloned LUN reservation and fractional space for that individual LUN will occupy their own disk space in the volume. Cloned LUNs will be unmapped, and to provide access you need to define igroup and create mapping.

```
Storage*> lun clone create /vol/new1/firstclone -b /vol/new1/first
hourly.0
Storage*> lun clone split start /vol/new1/firstclone
Storage*> lun clone split status /vol/new1/firstclone
Storage*> lun clone split stop /vol/new1/firstclone
```

## 9 Best Practices

### 9.1 Volume Layout

Volumes containing boot LUNs should be separated from application data to preserve Snapshot data integrity and prevent Snapshot locking when using LUN clones. Even though volumes containing boot LUNs may not require much physical disk space, give the volume

enough spindles such that performance is not bound by disk activity. With Data ONTAP 7G and later, volumes with boot LUNs can be created on the same aggregate in which the data volumes reside to maximize storage utilization without sacrificing performance.

## 9.2 Operating System Page File Placement

To free physical memory, Windows writes seldom-used blocks from memory to disk to the Operating System page file. This operation is called paging. By default, Windows puts the paging file on the boot partition where the Operating System is installed and creates a default size of the paging file that is 1.5 times the physical RAM, up to a maximum of 4095MB.

## 9.3 Page File Configuration

The optimal solution, other than the solution of adding more physical memory, is to do the following:

1. Create one paging file on the boot partition by using the default settings.
2. Create another paging file on a less frequently used partition on a separate physical disk or RAID volume.

To configure additional page files on a different LUN, do as follows:

1. To change the size of the virtual memory paging file1, in Control Panel, open System.
2. On the Advanced tab, under Performance, click Settings.
3. In the Performance Options dialog box, click the Advanced tab.
4. Under Virtual Memory, click Change.
5. Under Drive [Volume Label], click a different drive to create paging file.
6. Under Paging File Size for selected drive, click Custom Size, and type a new paging file size in megabytes in the Initial size (MB) or Maximum size (MB) box, and then click Set.

### 9.3.1 Page File Best Practice

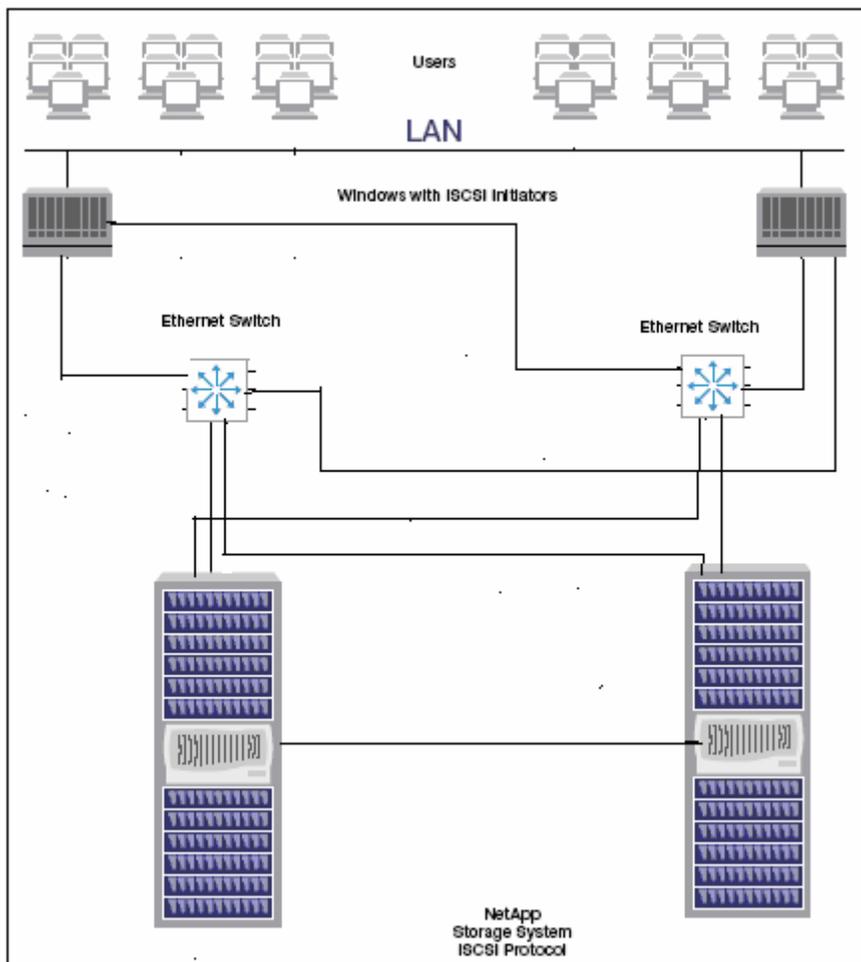
- Completely removing the paging file from the boot partition does not allow Windows to create a crash dump file (`Memory.dmp`) should a kernel mode STOP error occurs. Not having this crash dump file could lead to extended server downtime should the STOP require a debug to be performed, so always create any additional page file on a different partition that is configured on another iSCSI LUN.
- Placing the paging file on the boot partition does not optimize performance because Windows has to perform disk I/O on both the system directory and the paging file. Therefore, it is recommended that you create additional paging files on a different partition that is configured on a different iSCSI LUN presented by the storage controller.
- Windows will use the paging file on the less frequently used partition on iSCSI LUN over the paging file on the heavily used boot partition. Windows uses an internal algorithm to determine which page file to use for virtual memory management.
- It is recommended you create a small paging file on boot partition and then a separate page file on different volume on the system. This is because when you create a Snapshot copy of the boot volume on the storage, then space will be consumed, and if you have a large page file on it where data is not so important in case of recovering. Create a large page file on a separate volume that is an iSCSI LUN on a separate volume with Snapshots turned off.

## 9.4 High Availability

An Active-Active configuration is two storage systems (nodes) whose controllers are connected to each other either directly or through switches. You can configure the two nodes so that they share access to a common set of disks, subnets, and tape drives, or you can configure them to have two distinct sets of storage, each owned by one of the controllers. The nodes are connected to each other through a cluster adapter or an NVRAM adapter, which allows one node to serve data to the disks of its failed partner node. Each node continually monitors its partner, mirroring the data for each other's nonvolatile RAM (NVRAM).

- Ensure that the two power supplies for both the controllers and disk shelves are on different power circuits, so that a single power outage does not affect both components.
- Use VIFs (Virtual Interfaces) to provide redundancy and improve availability of network communication.
- Ensure that each node has sufficient resources to adequately support the workload of both nodes during takeover mode.
- Configure separate network switches with dual path availability to server and storage for iSCSI connectivity.
- Configure separate VLANs on network switches for security and to isolate network traffic.

Figure 26) iSCSI Protocol



## 10 References

- Trouble shooting on Intel Gigabit Network adapters.  
<http://www.intel.com/support/network/adapter/pro100/sb/CS-025776.htm#anchor3>
- White Paper: Boot from SAN in Windows Server 2003 and Windows 2000 Server  
<http://www.microsoft.com/windowsserversystem/storage/solutions/sanintegration/bootfromsaninwindows.mspx>
- STORport in Windows Server 2003: Improving Manageability and Performance in Hardware RAID and Storage Area Networks  
<http://www.microsoft.com/windowsserversystem/storage/technologies/storport/storportwp.mspx>
- Server Clusters: Storage Area Networks - For Windows 2000 and Windows Server 2003  
<http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/clustering/starenet.mspx>
- Microsoft Storage Technologies – Boot From SAN  
<http://www.microsoft.com/windowsserversystem/storage/technologies/bootfromsan/default.mspx>
- Support for Booting from a Storage Area Network (SAN)  
<http://support.microsoft.com/default.aspx?scid=kb;en-us;305547>
- iSCSI and Fiber Channel Configuration Guide  
[http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/QuickRef/fc\\_iscsi\\_config\\_guide.pdf](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/fc_iscsi_config_guide.pdf)

## 11 Summary

As discussed, using NetApp storage with iSCSI capability can be used to lower operating costs through centralized data consolidation and enable rapid deployment of an Operating System.