

April 2004

**DB2** Information Management Software



# Linux iSCSI for IBM<sup>®</sup> DB2<sup>®</sup> Universal Database<sup>™</sup> on IBM eServer<sup>™</sup> BladeCenter

By  
Baris Naciterhan  
IBM Software Group

<b>Contents</b>	
<b>1. INTRODUCTION</b>	<b>1</b>
1.1 Objective	1
1.2 Overview	1
1.3 Basic Requirements	2
<b>2. TEST ENVIRONMENT</b>	<b>3</b>
2.1 General Characteristics of the Linux iSCSI Cluster	3
2.2 IP Setup of the Test Environment	4
<b>3. LINUX iSCSI SETUP</b>	<b>5</b>
3.1 Hardware Setup	5
3.2 Software Setup	6
3.2.1 The software setup of the IBM BladeCenter	6
3.2.1.1 The initial configuration	6
3.2.1.2 The firmware and the driver update	13
3.2.1.3 The initiator software installation	17
3.2.2 The software setup of the Netapp FAS250	18
<b>4. LINUX iSCSI CONFIGURATION</b>	<b>20</b>
4.1 Linux iSCSI configuration of the initiator server	20
4.2 Linux iSCSI configuration of the Target storage	23
4.3 Creating a file system on the iSCSI initiator	31
<b>5. CONFIGURING DB2 UDB FOR LINUX iSCSI</b>	<b>36</b>
5.1 Installing DB2 UDB	36
5.2 Testing DB2 UDB	36

## 1. INTRODUCTION

### 1.1 Objective

This white paper describes the setup and implementation of Linux iSCSI for IBM® DB2® Universal Database™ (UDB) on IBM® eServer™ BladeCenter™. In this step-by-step guide, you will find all the information required to set up and run Linux iSCSI with DB2 UDB Database.

### 1.2 Overview

#### *What is Linux iSCSI?*

Linux iSCSI is the iSCSI platform supported by Linux-based networks. iSCSI is a protocol that provides the ability to access a storage unit or storage-area networks through an IP network. The iSCSI protocol makes the data communication possible between the host and the storage device over IP networks. The main advantage of iSCSI is its structure, which is composed of two main widely adapted protocols, SCSI and TCP/IP. The iSCSI protocol offers the ease of scalability and management of the storage area networks (SAN) without the high costs of SCSI or Fiber Channel. As iSCSI uses an IP gateway, it eliminates the distance limitations and complexity problems of other storage protocols.

#### *How iSCSI works*

When a user application on the iSCSI initiator sends a data request, Linux generates the necessary SCSI commands for the request and then encapsulates them with the data request. The next step is the encryption of the packet, depending on the needs and the security level of the network environment. Before sending over the network, the last step is adding an IP packet header to the packet.

When a packet is received by an iSCSI target, the first step is to check if the packet has been encrypted. If it has, then the packet is decrypted in order to access its contents; otherwise, the packet is passed on for further processing. The next step is sending the SCSI commands to the SCSI controller. The SCSI controller transmits the SCSI commands to the storage unit to process the requested data. Before the data is sent to the initiator server, it passes through the same steps described above to become an IP packet. Figure 1 illustrates the Linux iSCSI process.

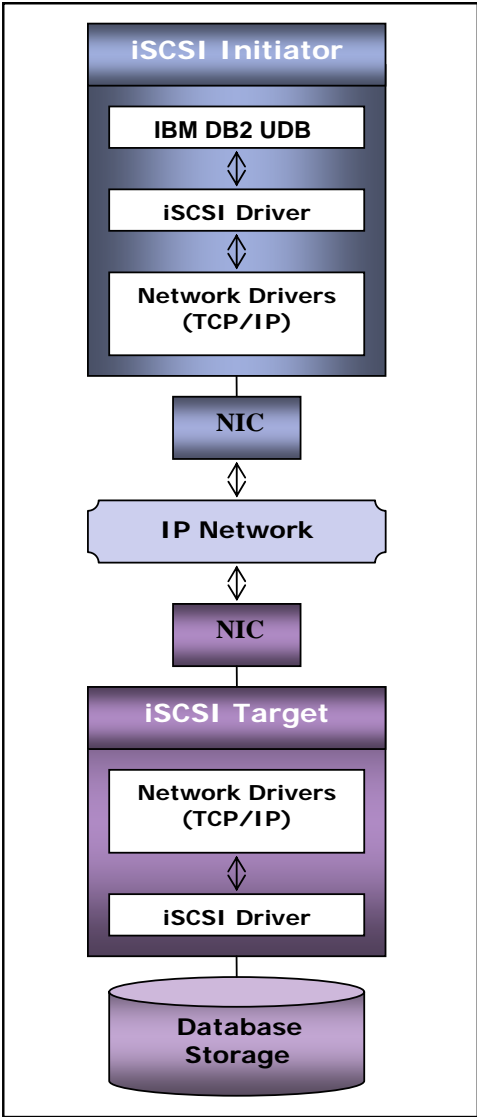


Figure 1: The process schema of the Linux iSCSI for IBM DB2 UDB.

1.3 Basic Requirements

- ✓ One or more Linux servers to be configured as iSCSI initiators
- ✓ An external storage unit with iSCSI support

- ✓ A cross-over cable for private network option
- ✓ DB2 UDB Enterprise Version 8.1 (or higher)

Note that in order to configure and run Linux iSCSI you must log in as root.

## 2. TEST ENVIRONMENT

### · BladeCenter

As all blade servers in the BladeCenter are identical, the following configuration represents all of them. To test Linux iSCSI on a partitioned database environment, four blade servers on the BladeCenter are configured.

- BladeCenter : HS20 8678
- CPU (dual) : Intel® Xeon 2.0 GHz
- Memory : 2.5 GB
- NICs : Broadcom Gigabit
- Hard Disk : SCSI 37.2

### · Netapp FAS250

- CPU (dual) : 1.25 GHz
- Memory : 510 MB
- NICs : Gigabit Dual Ethernet Controller
- Storage : 14 Disks ≈ 1 TB

### 2.1 General Characteristics of the Linux iSCSI Cluster

The first NIC (Eth0) of BladeCenter is connected to the local network. The second NIC (Eth1) of the BladeCenter is connected to the second NIC (Eth1) of the storage unit. The first NIC (Eth0) of the storage unit is connected to the local network.

## 2.2 IP Setup of the Test Environment

### *Turtle01 (iSCSI Initiator 1)*

Eth0 : 9.26.97.201  
Eth1 : 172.16.0.1

### *Turtle02 (iSCSI Initiator 2)*

Eth0 : 9.26.97.202  
Eth1 : 172.16.0.2

### *Turtle03 (iSCSI Initiator 3)*

Eth0 : 9.26.97.206  
Eth1 : 172.16.0.3

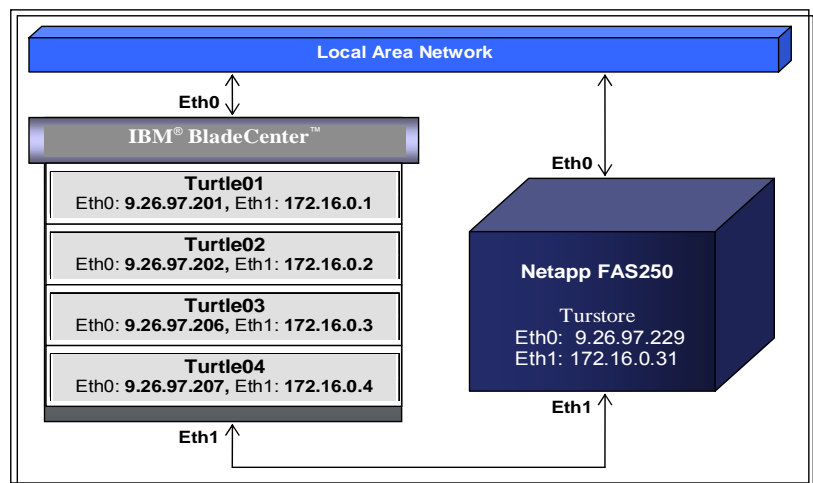
### *Turtle04 (iSCSI Initiator 4)*

Eth0 : 9.26.97.207  
Eth1 : 172.16.0.4

### *Turstore (iSCSI Target)*

Eth0 : 9.26.97.229  
Eth1 : 172.16.0.31

Figure 2 demonstrates the IP schema of the Linux iSCSI test cluster:



**Figure 2:** IP Schema of the Linux iSCSI Test Cluster

### 3. LINUX iSCSI SETUP

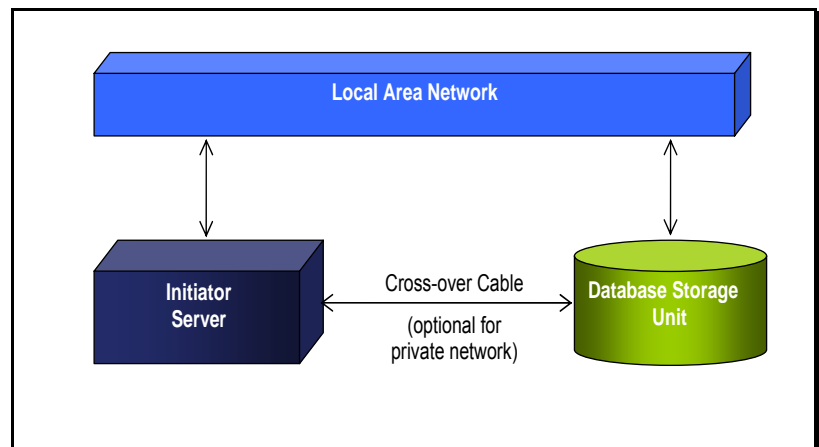
The Linux iSCSI for DB2 UDB setup process has two parts:

- Hardware setup
- Software setup

The performance of the system depends on both the hardware and the software properties of the whole environment.

#### 3.1 Hardware Setup

Two units must be configured in order to run Linux iSCSI for DB2. The first one is the BladeCenter, which is named as the initiator server in the iSCSI setup. And the second one is the Netapp FAS250, which is named as the iSCSI target storage in the iSCSI setup. For the connection between the units, there are two options. You can either use your local network connection or you can set up a private network between the units. Although using the local network connection is the more common method, you may choose to use a private network setup for performance reasons or to avoid having to rely on a local connection. Figure 3-1 demonstrates the Linux iSCSI hardware setup.



**Figure 3-1:** *The hardware schema of the Linux iSCSI for DB2 UDB.*

Two network interface cards (NIC) are configured on both units. One network card on each unit is connected to the external network, and the other network card is connected to the network card on the other unit via a cross-over cable.

## 3.2 Software Setup

The iSCSI Software Setup is divided into two main sections. The first one is the software setup of the initiator server, and the second one is the software configuration of the target unit.

### 3.2.1 The software setup of the BladeCenter

This part includes the initiator software installation on the BladeCenter. But before this step, the BladeCenter should be prepared for the Linux iSCSI environment. The preparation process consists of:

- The initial configuration of the BladeCenter
- Firmware update

#### 3.2.1.1 The initial configuration of the BladeCenter

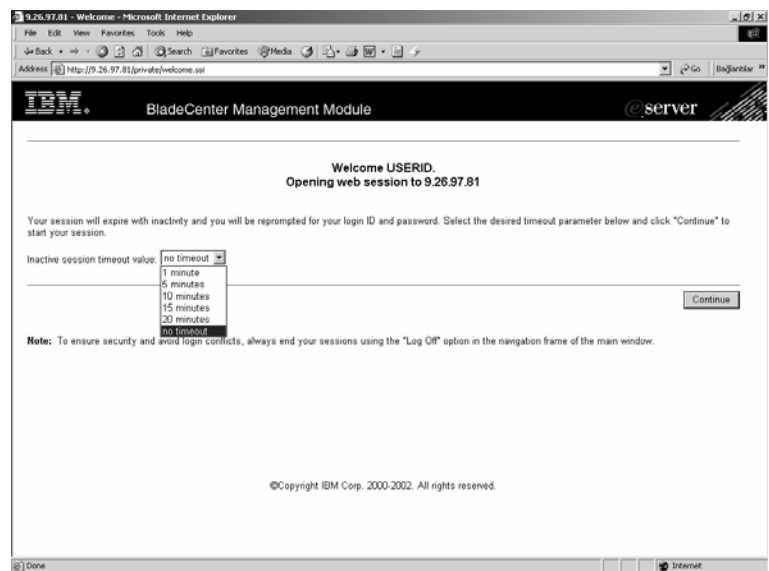
The Web site of the management module is used for initial configuration of the BladeCenter. There are two ways to log on to the Web site of the management module:

- I. If a DHCP server is available on your network, simply connect the Ethernet port of the management module to the local network using a standard Ethernet cable, open the Web browser and type the host name assigned by the DHCP server.
- II. If a DHCP server is not available on your network or if you want to connect directly to the management Web site, connect the Ethernet port of the management module to the target computer using a cross-over Ethernet cable, open the Web browser and type the default IP of the management module. The default static IP of the management module is 192.168.70.125.



**Figure 3-2:** *The Network Password window.*

Depending on your connection type, after typing the necessary host name/default IP into the address field of the Web browser, the network password window will come up as seen in Figure 3-2. Type **USERID** (Uppercase) into the 'User Name' field and **PASSWORD** (Uppercase and the character after 'W' is the number zero.) into the 'Password' field. These are the default values of the user ID and password fields for the Web site of the management module. After clicking **OK**, the resulting screen should look similar to Figure 3-3.



**Figure 3-3:** Welcome screen of the management module.

At this window, choose the session timeout length and then click **Start New Session**. Available session time options are:

- 1 minute
- 5 minutes
- 10 minutes
- 15 minutes
- 20 minutes
- no timeout

Select your session choice and then click **Continue**. You will see a window similar to Figure 3-4.



**System Status Summary**

System is operating normally. All monitored parameters are OK.

The following links can be used to view the status of different components.

[Blade Servers](#)  
[I/O Modules](#)  
[Management Modules](#)  
[Power Modules](#)  
[Blowers](#)  
[Front Panel](#)

**Blade Servers**

Click the icon in the Status column to view detailed information about each blade server.

Bay	Status	Name	Pwr	Owner <sup>1</sup>		Network		WOL <sup>1</sup>	Local Control			BSE <sup>1</sup>
				KVM	MT <sup>1</sup>	Onboard	Card		Pwr	KVM	MT <sup>1</sup>	
1	●	SN#K10UJ33M1HD	On	X	X	Eth	Fib   ...	On	X	X	X	
2	●	SN#K10UJ33N17B	On			Eth	Fib   ...	On	X	X	X	
3	●	SN#K10UJ33N162	On			Eth	Fib   ...	On	X	X	X	
4	●	SN#K10UJ33F124	On			Eth	Fib   ...	On	X	X	X	
5	●	SN#K10UJ33N114	On			Eth	Fib   ...	On	X	X	X	
6	●	SN#K10UJ33M159	On			Eth	Fib   ...	On	X	X	X	
7	●	SN#K10UJ33H159	On			Eth	Fib   ...	On	X	X	X	
8	●	SN#K10UJ33N166	On			Eth	Fib   ...	On	X	X	X	
9	●	SN#K10UJ33N146	Off			Eth	Fib   ...	On	X	X	X	

**Figure 3-4:** Main screen of the management module.

The menu on the left side of the window includes all functions you need to control the BladeCenter, and is divided into the four main sections listed below.

- Monitors
  - System Status
  - Event Log
  - LEDs
  - Hardware VPD
  - Firmware VPD
- Blade Tasks
  - Power/Restart
  - Remote Control
  - Firmware Update
  - Configuration
- I/O Module Tasks
  - Power/Restart
  - Management
  - Firmware Update

- MM Control
  - General Settings
  - Login Profiles
  - Alerts
  - Network Interfaces
  - Network Protocols
  - Security
  - Configuration File
  - Firmware Update
  - Restore Defaults
  - Restart MM
- Log Off

The **Monitors** section includes the functions to view the status and the settings of the whole system:

*System Status:* Using this feature, you can view the overall status of the blades.

*Event Log:* Using this feature, you can view all entries that are currently stored in the event and error log files of the management module.

*LEDs:* Using this feature, you can check the current status of the LEDs of the blades and the whole system.

*Hardware VPD:* Using this feature, you can check the hardware description of the BladeCenter and all modules present on the BladeCenter.

*Firmware VPD:* Using this feature, you can check the software description of the BladeCenter and all modules present on the BladeCenter.

The **Blade Tasks** section includes the functions to control and change the settings of the blade servers.

*Power/Restart:* Using this feature, you can power the blade servers on/off, restart the service processor, enable/disable the local power control and enable/disable the 'Wake on LAN' option for each blade server.

*Remote Control:* Using this feature, you can remotely control a blade server from any computer connected to the Web site of the management module.

*Firmware Update:* Using this feature, you can update the firmware of each blade server.

*Configuration:* Using this feature, you can display the blade information and change the boot sequence of each blade server.

The **I/O Module Tasks** section includes the functions to manage and change the settings of the modules present in the BladeCenter:

*Power/Restart:* Using this feature, you can power on/off and restart the modules of the BladeCenter.

*Management:* Using this feature, you can change the IP configuration of the modules.

*Firmware Update:* Using this feature, you can update the firmware of each module present in the BladeCenter to its most recent version.

The **MM Control** section includes the functions to configure the settings of the management module:

*General Settings:* Using this feature, you can define a name for the management module, set the date and time, and also write the contact/location information.

*Login Profiles:* Using this feature, you can create login profiles with specified read/write permissions.

*Alerts:* Using this feature, you can set and manage the alerts of the BladeCenter.

*Network Interfaces:* Using this feature, you can configure the network connections of the BladeCenter.

*Network Protocols:* Using this feature, you can configure the settings for the SNMP, SMTP and DNS network protocols.

*Security:* Using this feature, you can configure the settings of the secure socket layer (SSL).

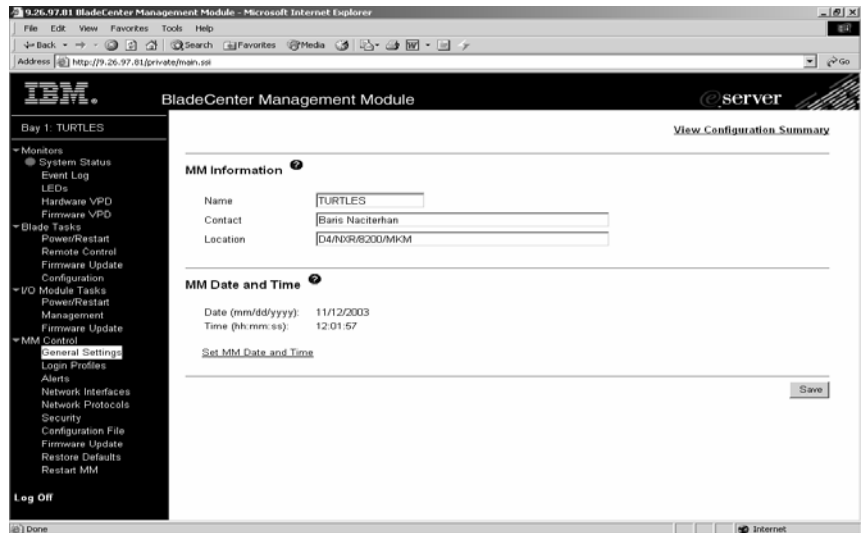
*Configuration File:* Using this feature, you can backup/restore the configuration of the management module.

*Firmware Update:* Using this feature, you can update the firmware of the management module.

*Restore Defaults:* Using this feature, you can restore the default configuration.

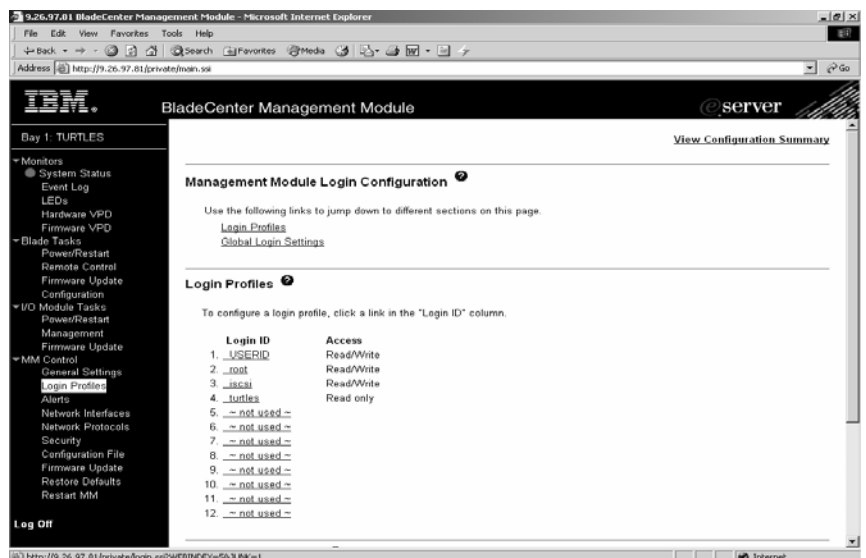
*Restart MM:* Using this feature, you can restart the management module. To complete the initial configuration of the BladeCenter, follow the steps written below:

I. Click the *General Settings* link in the MM Control section. Fill the Name, Contact and Location fields and set the date&time by clicking the 'Set MM Date and Time' link. Figure 3-5 shows the General Settings window of the management module.



**Figure 3-5:** General Settings window of the BladeCenter management module.

II. Click the *Login Profiles* link in the MM Control section. For security reasons, define a new user name and password by clicking any of the '~not used~' links. Figure 3-6 shows the Login Profiles window of the management module.



**Figure 3-6:** Login Profiles window of the management module.

III. Click the *Network Interfaces* link in the MM Control section. Enter your network information into the appropriate fields. Figure 3-7 shows the *Network Interfaces* window of the management module.

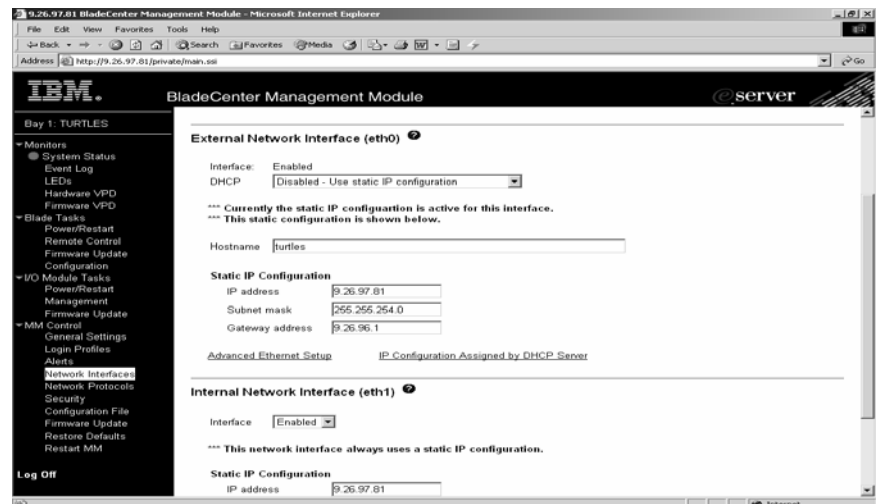


Figure 3-7: *Network Interfaces* window of the management module.

IV. Choose your network protocol type by clicking the *Network Protocols* link in the MM Control section, then enter your network protocol information into the appropriate fields. Figure 3-8 shows the *Network Protocols* window of the management module.

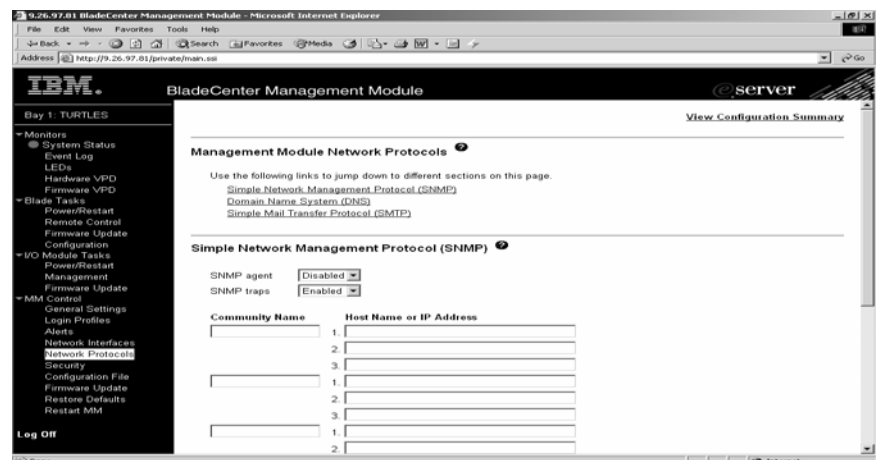
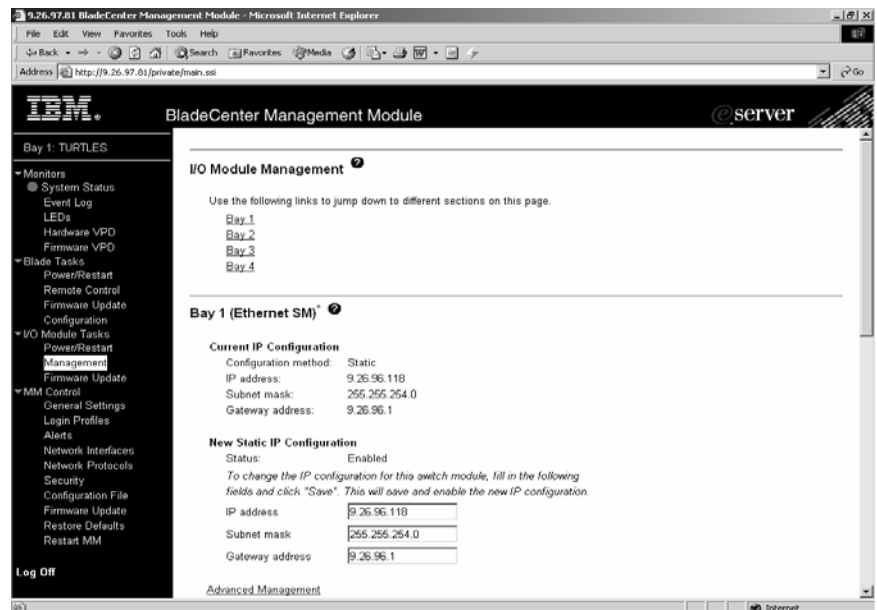


Figure 3-8: *Network Protocols* window of the management module.

V. Back up the configuration of the management module by clicking the *Configuration File* link in the MM control section. If a failure occurs in the management module, you can restore the configuration settings from the backup.

VI. Configure the network settings of the network switch modules present in your system by clicking the *Management* link in the I/O Module Tasks section. Enter the network information of the modules into the appropriate fields. Figure 3-9 shows the *I/O Management* window of the management module.



**Figure 3-9:** I/O Module Management window of the management module.

### 3.2.1.2 The firmware and the driver update of the BladeCenter

For a more stable environment and to get the best performance from the BladeCenter, it is very important to update the firmware and driver versions of your system. You should update the firmware of each module and the firmware of each blade server in the BladeCenter.

Use the following Web site to get the most recent firmware and driver versions of the BladeCenter.

<http://www.ibm.com/pc/support/site.wss/>

Follow the firmware and driver installation instructions there. You can view and update the firmware versions of the components of the BladeCenter by using the management module Web site.

While searching for the driver and firmware for the BladeCenter, you may need to check the hardware information (including module name, manufacturer ID, Machine Type/Model, Serial No, and so on) of the whole system. It is possible to view and edit the current hardware information of the BladeCenter by using the management module's Web site. Simply click *Hardware VPD* under the Monitors section in order to view/edit the hardware information of the BladeCenter.

Figure 3-10 shows the *Hardware VPD* window of the management module.

The screenshot shows the BladeCenter Management Module web interface. The main content area is titled "BladeCenter Hardware VPD" and contains a table of hardware components. The table has the following columns: Bay(s), Module Name, Manuf. ID, Machine Type/Model, Machine Serial No., Hardware Revision, Manuf. Date, Part Number, FRU Number, and FRU Serial. The table is divided into sections for "Chassis and Media Tray" and "Blade Servers".

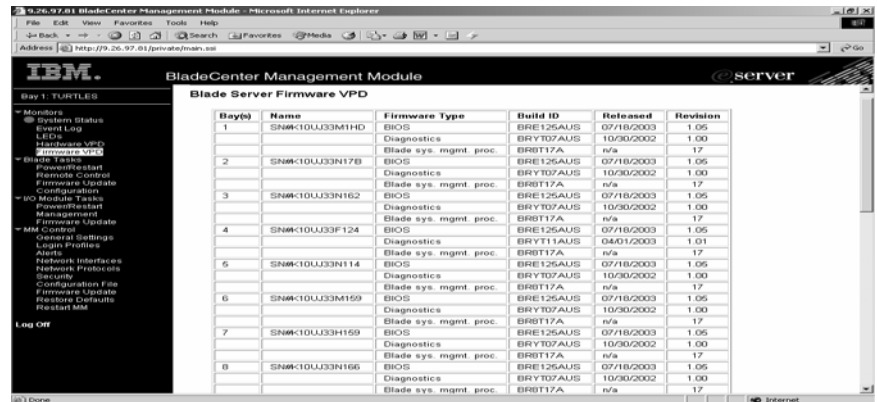
Bay(s)	Module Name	Manuf. ID	Machine Type/Model	Machine Serial No.	Hardware Revision	Manuf. Date	Part Number	FRU Number	FRU Serial
<b>Chassis and Media Tray</b>									
	Chassis	SLRM	86771XX	78C6261	5	1602	71P8807	59P6609	J1PH3C
1	Media Tray	---	n/a	n/a	0	3800	32P1936	59P6629	--
<b>Blade Servers</b>									
1	SNMK10LU33M1HD	SLRM	867821X	KPV8956	8	1203	71P8790	59P6610	K10LLE
	Daughter Card	SLRM	n/a	n/a	5	0103	59P6623	59P6624	J1R2M
2	SNMK10LU33N17B	SLRM	867821X	KPV8956	8	1203	71P8790	59P6610	K10ULL
	Daughter Card	SLRM	n/a	n/a	5	1303	59P6623	59P6624	J1R3H3
3	SNMK10LU33N162	SLRM	867821X	KPV8841	8	1203	71P8790	59P6610	K10ULL
	Daughter Card	SLRM	n/a	n/a	5	0103	59P6623	59P6624	J1R2M
4	SNMK10LU33F124	SLRM	867821X	KPV8763	8	1203	71P8790	59P6610	K10ULL
	Daughter Card	SLRM	n/a	n/a	5	1303	59P6623	59P6624	J1R3H3
5	SNMK10LU33N114	SLRM	867821X	KPV8902	8	1203	71P8790	59P6610	K10ULL

Figure 3-10: *Hardware VPD* window of the management module.

» You can view and edit the firmware information of the BladeCenter.

It is possible to view the current firmware information of the BladeCenter by using the management module's Web site. You can view the firmware information of the blade servers and the modules installed in the BladeCenter.

Click *Firmware VPD* under the Monitors section in order to view the firmware information of your BladeCenter. Figure 3-11 shows the *Firmware VPD* window of the management module.



Bay(s)	Name	Firmware Type	Build ID	Released	Revision
1	SNM-C10U33M1HD	BIOS	BRE126AUS	07/18/2003	1.05
		Diagnostics	BRYT07AUS	10/30/2002	1.00
		Blade sys. mgmt. proc.	BRBT17A	n/a	17
2	SNM-C10U33N17B	BIOS	BRE126AUS	07/18/2003	1.05
		Diagnostics	BRYT07AUS	10/30/2002	1.00
		Blade sys. mgmt. proc.	BRBT17A	n/a	17
3	SNM-C10U33N162	BIOS	BRE126AUS	07/18/2003	1.05
		Diagnostics	BRYT07AUS	10/30/2002	1.00
		Blade sys. mgmt. proc.	BRBT17A	n/a	17
4	SNM-C10U33F124	BIOS	BRE126AUS	07/18/2003	1.05
		Diagnostics	BRYT11AUS	04/01/2003	1.01
		Blade sys. mgmt. proc.	BRBT17A	n/a	17
5	SNM-C10U33N114	BIOS	BRE126AUS	07/18/2003	1.05
		Diagnostics	BRYT07AUS	10/30/2002	1.00
		Blade sys. mgmt. proc.	BRBT17A	n/a	17
6	SNM-C10U33M159	BIOS	BRE126AUS	07/18/2003	1.05
		Diagnostics	BRYT07AUS	10/30/2002	1.00
		Blade sys. mgmt. proc.	BRBT17A	n/a	17
7	SNM-C10U33H159	BIOS	BRE126AUS	07/18/2003	1.05
		Diagnostics	BRYT07AUS	10/30/2002	1.00
		Blade sys. mgmt. proc.	BRBT17A	n/a	17
8	SNM-C10U33N166	BIOS	BRE126AUS	07/18/2003	1.05
		Diagnostics	BRYT07AUS	10/30/2002	1.00
		Blade sys. mgmt. proc.	BRBT17A	n/a	17

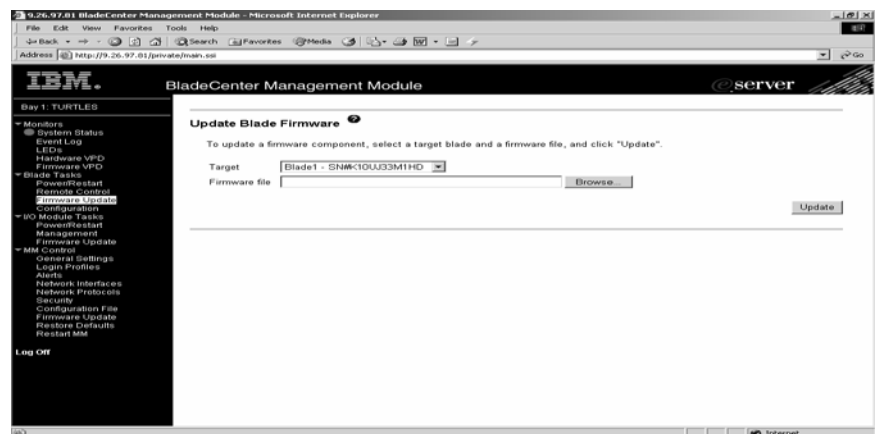
Figure 3-11: *Firmware VPD* window of the management module.

» You can update the firmware component of any selected blade server.

Simply click *Firmware Update* under Blade Tasks section in order to update the firmware component of your BladeCenter.

Figure 3-12 shows the *Firmware Update of the Blades* window of the management module.

To update a firmware component, select a target blade and a firmware file from the pull-down menus and then click "Update".



**Update Blade Firmware**

To update a firmware component, select a target blade and a firmware file, and click "Update".

Target:

Firmware file:

Figure 3-12: *Firmware Update of the Blades* window of the management module.

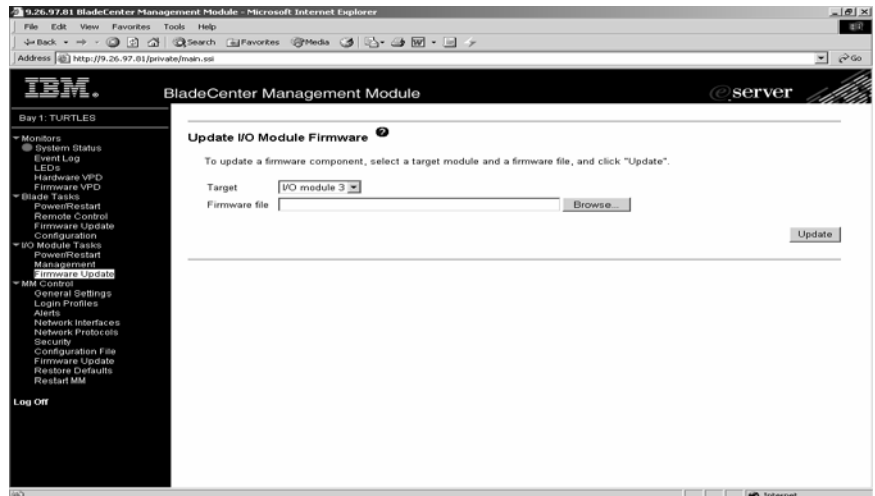


» You can update the firmware of any selected I/O module.

Simply click *Firmware Update* under the I/O Module Tasks section in order to update the firmware component of any selected module.

Figure 3-13 shows the *Firmware Update of the I/O Modules* window.

To update a firmware component, select a target module and a firmware file from the pull-down menu and then click "Update".



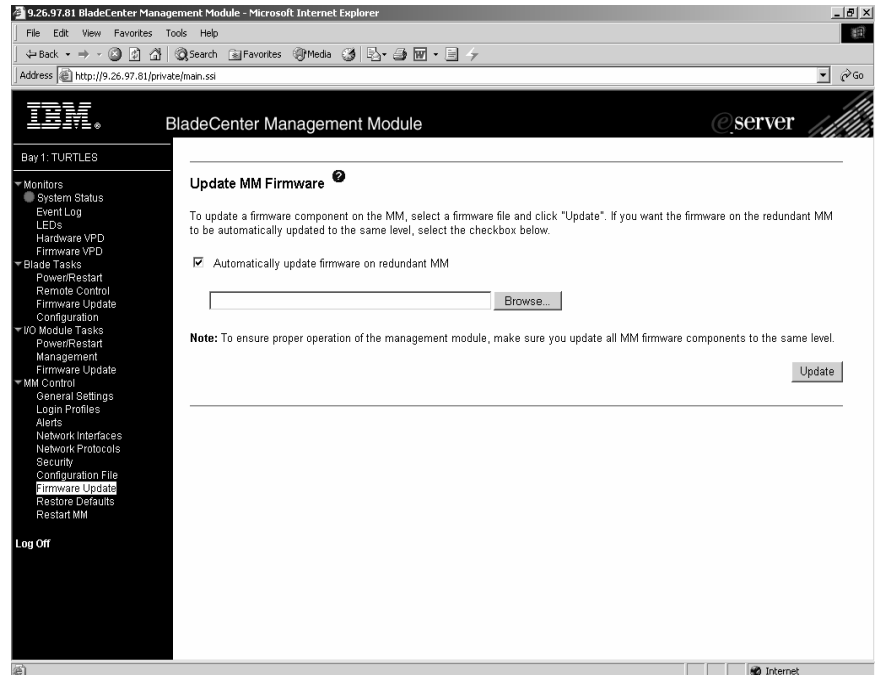
**Figure 3-13:** *Firmware Update of the Modules* window of the management module.

» You can update the firmware of the management module.

To update the firmware of the management module, which is different from updating the I/O modules, simply click *Firmware Update* in the MM Control section.

Figure 3-14 shows the Management Module update window.

To update a firmware component in the MM, select a firmware file using the 'Browse' button and then click "Update". Also, depending on your system, if you want the firmware on the redundant MM to be automatically updated to the same level, select the check box on the same window.



**Figure 3-14:** Management Module firmware update window.

### 3.2.1.3 The initiator software installation on the BladeCenter

For a complete installation of the initiator software on the BladeCenter, you must complete the following steps:

- ✓ Install the most current version of your Linux distribution that is supported by iSCSI on the blades. Ensure that the Linux distribution has the most current packages by updating to the most recent fix pack. You can check the operating system requirements on the SourceForge Web site.
- ✓ Ensure that the kernel source matches the kernel running on each blade.
- ✓ Download a released version of the Linux iSCSI initiator driver from the SourceForge Web site and install the driver following the installation instructions provided on the Web site at:

<http://sourceforge.net/projects/linux-iscsi/>

### 3.2.2 The software setup of the Netapp FAS250

After completing the software setup on BladeCenter, the next step is to deal with the target storage unit. Since our target device supports iSCSI and the iSCSI driver is already installed, we only need to prepare the operating system environment to be able to work with BladeCenter. To connect to the unit for the first time, you can use the serial connection method and connect to the storage unit via a serial cable using a connection program such as 'minicom'. To configure the storage unit Netapp FAS250, you will need to enter some information detailing your environment. Enter '**setup**'. The resulting screen will be similar to the one below:

```
NetApp Release 6.4.1: Fri May 2 14:54:02 PDT 2003
System ID: 0084165423 (turstore)
System Serial Number: 1039689 (turstore)
slot 0: System Board
    Processors:      2
    Processor revision: B2
    Processor type:  1250
    Memory Size:    510 MB
slot 0: FC Host Adapter 0b
    14 Disks:      952.0GB
    1 shelf with EFH
slot 0: FC Host Adapter 0c
slot 0: SB1250 Gigabit Dual Ethernet Controller
    e0a MAC Address: 00:a0:98:01:28:ec (auto-100tx-fd-up)
    e0b MAC Address: 00:a0:98:01:28:ed (auto-1000t-fd-up)
slot 0: NetApp ATA/IDE Adapter 0a (0x000000000000001f0)
    0a.0          122MB
Please enter the new hostname:
```

Enter a hostname for your storage unit. After this point, you will continue entering some values (such as IP address, netmask, host name, etc.), depending on your network environment. Check the example screen on the next page.

```
Please enter the new hostname : turstore
Do you want to configure virtual network interfaces? [n]: n
Please enter the IP address for Network Interface e0a : 9.26.97.229
Please enter the netmask for Network Interface e0a : 255.255.254.0
Please enter media type for e0a {100tx-fd, 100tx, auto (100/1000)}
[auto]:
Please enter flow control for e0a {none, receive, send, full} [full]:
Do you want e0a to support jumbo frames? [n]:
Please enter the IP address for Network Interface e0b : 172.16.0.31
Please enter the netmask for Network Interface e0b : 255.255.0.0
Please enter media type for e0b {100tx-fd, 100tx, auto (100/1000)}
[auto]:
Please enter flow control for e0b {none, receive, send, full} [full]:
Do you want e0b to support jumbo frames? [n]:
Would you like to continue setup through the web interface? [n]:
Please enter the name or IP address of the default gateway : 9.26.96.1
Please enter the name or IP address of the administration host : ibappl
Please enter the IP address for ibappl : 9.26.166.128
Please enter timezone [EST]: EST
Where is the filer located? : HEAVY-C2
What language will be used for multi-protocol files (Type ? for list)? :
Use vol lang command to set filer language after files installed
What language will be used for multi-protocol files (Type ? for list)? :
Enter the root directory for HTTP files : /vol/vol0/home/http
Do you want to run DNS resolver? [y]:
Please enter DNS domain name : torolab.ibm.com
You may enter up to 3 nameservers
Please enter the IP address for first nameserver : 9.26.4.5
```

```
Do you want another nameserver? [n]:
```

```
Do you want to run NIS client? [n]:
```

As you can see from the preceding example, the default values are presented in brackets. To use a default value, simply press Enter without typing anything.

When you have addressed all the questions, enter the following command:

```
turstore> ifconfig -a
```

You should see the IP address of the target configured during setup. The resulting output will be similar to the example below:

```
e0a:  flags=848043<UP,BROADCAST,RUNNING,MULTICAST>  mtu
1500
    inet 9.26.97.229 netmask 0xffffe00 broadcast 9.26.97.255
    ether 00:a0:98:01:28:ec (auto-100tx-fd-up) flowcontrol full
e0b:  flags=848043<UP,BROADCAST,RUNNING,MULTICAST>  mtu
1500
    inet 172.16.0.31 netmask 0xffff0000 broadcast 172.16.255.255
    ether 00:a0:98:01:28:ed (auto-1000t-fd-up) flowcontrol full
lo:
flags=1948049<UP,LOOPBACK,RUNNING,MULTICAST,TCPCKSUM>
mtu 9188
    inet 127.0.0.1 netmask 0xff000000 broadcast 127.0.0.1
```

Enter **'reboot'** for the changes to take effect.

## 4. LINUX iSCSI CONFIGURATION

To be able to use Linux iSCSI, you must configure both the initiator server and the storage unit.

### 4.1 Linux iSCSI configuration of the Initiator Server

There are 14 blades in the BladeCenter. Depending on your database needs, you can use as many as you want among these 14 blade servers for Linux iSCSI. To be able to run Linux iSCSI, you **must make** the iSCSI initiator name unique on each blade server.

To configure the chosen blades as iSCSI initiators, complete the following steps in the same order on all of the blades.

**1-** Enter the following commands to verify the iSCSI module:

```
turtle01:~ # lsmod | grep iscsi
```

The resulting screen should be similar to the example below:

```
iscsi_mod      201340 2
```

```
turtle01:~ # modinfo iscsi iscsi_mod
```

The resulting output should be similar to the example below:

```
filename: /lib/modules/2.4.19-64GB-SMP/misc/iscsi_mod.o
description: "iSCSI initiator"
author: "Cisco Systems, Inc."
license: "GPL"
parm:      force_tcq int, description "when non-zero, force tagged
command queueing for all devices"
parm:      untagged_queue_depth int, description "queue depth to use
for devices that don't support tagged command queueing"
parm:      translate_deferred_sense int, description "translate deferred
sense data to current sense data in disk command responses"
parm:      iscsi_reap_tasks int
```

If you don't see the iSCSI module information after entering the commands above, you must re-install the iSCSI driver. Download the iSCSI driver from <http://sourceforge.net/projects/linux-iscsi/> and follow the installation instructions included in the driver documentation, which is also available on that Web site.

**2-** Enter the following commands to check or configure the iSCSI initiator name of the blade:

```
turtle01:~ # cat /etc/initiatorname.iscsi
```

The resulting output should be similar to the example below:

```
## DO NOT EDIT OR REMOVE THIS FILE!  
## If you remove this file, the iSCSI daemon will not start.  
## If you change the InitiatorName, existing access control lists  
## may reject this initiator. The InitiatorName must be unique  
## for each iSCSI initiator. Do NOT duplicate iSCSI InitiatorNames.  
InitiatorName=iqn.1987-  
05.com.cisco.01.5a8b71fe3ed4b4c327e18f27efd5686b
```

This initiator name will be used while registering to the iSCSI target. That's why it **must** be unique for each blade. If you need to use two or more iSCSI initiators, you must be sure that each blade has a different initiator name. Before registering the iSCSI target, you can change this name but **do not** change it after the iSCSI initiator blade has been registered with the iSCSI target. When changing the initiator name, any name can be chosen as long as format is preserved. As an example, written below are the initiator names of the blades in the iSCSI test environment.

```
iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e28f27efd5686b
```

The initiator name of the third blade; *turtle03* is:

```
iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e38f27efd5686b
```

The initiator name of the fourth blade; *turtle04* is:

```
iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e48f27efd5686b
```

**3-** Enter the following command in order to modify the iSCSI configuration file:

```
turtle01:~ # vi /etc/iscsi.conf
```

Add the following lines at the beginning of the file. If you do not add these lines, there will be some problems with connectivity and performance:

```
Continuous=no  
HeaderDigest=never  
DataDigest=never  
DiscoveryAddress=[TargetIPAddress]
```

Do not forget to write your IP address on the storage unit into the *DiscoveryAddress* field. Also note that you should add the *DiscoveryAddress=[FilerIPAddress]* entry into the file for each GbE interface used on the filer.

The resulting screen will be similar to the example below:

```
Continuous=no
HeaderDigest=never
DataDigest=never
DiscoveryAddress=172.16.0.31
# iSCSI configuration file - see iscsi.conf(5)
```

4. Enter the following command in order to start the iSCSI service:

```
turtle01:~ # /etc/init.d/iscsi start
```

Note that after setup is done, the iSCSI service starts automatically upon reboot.

Check the running iSCSI processes:

```
turtle01:~ # ps -edf | grep iscsi
```

The resulting output will be similar to the example below:

```
root  564  1 0 17:36 ?    00:00:00 iscsid -l /dev/iscsi -m 755
root  569  1 0 17:36 ?    00:00:00 [iscsi-timer]
root  570  1 0 17:36 ?    00:00:00 [iscsi-tx]
root  571  1 0 17:36 ?    00:00:00 [iscsi-rx]
```

## 4.2 Linux iSCSI Configuration of the Target Storage

In order to use iSCSI, you must configure the storage unit.. Since the network settings of the target are now configured, you can connect to the unit using telnet.

» By typing '?', you can see the commands that are available on the storage unit.

```
turstore> ?
```



The resulting output will be similar to the example below:

?	help	nfsstat	snmp
arp	hostname	nis	software
backup	httpstat	options	source
cf	ifconfig	orouted	storage
cifs	ifstat	partner	sysconfig
config	igroup	passwd	sysstat
date	ipsec	ping	timezone
df	ipspace	priv	traceroute
disk	iscsi	qtree	ups
disk_fw_update	iswt	quota	uptime
dns	license	rdate	useradmin
download	logger	reboot	version
dump	logout	restore	vfiler
echo	lun	rmc	vif
ems	man	route	vlan
environment	maxfiles	routed	vol
exportfs	mt	savecore	vscan
fc	nbtstat	secureadmin	wcc
fcstat	ndmpcopy	setup	ypcat
file	ndmpd	shelfchk	ypgroup
filestats	netdiag	snap	ypmatch
fpolicy	netstat	snapmirror	ypwhich
halt	nfs	snapvault	

To configure the storage unit as an iSCSI target, complete the following steps:

**1-** Check the iSCSI name of the target by entering the following command:

```
turstore> iscsi nodename
```

The resulting output will be similar to the example below:

```
iSCSI target nodename: iqn.1992-08.com.netapp:sn.84165423
```

**2-** Enter the following commands in order to verify that the storage device can see the blade configured as iSCSI initiator:

```
turstore> iscsi show initiator
```

Depending on the number of the blades configured as iSCSI initiators, the resulting output should be similar to the example below:

```
Initiators connected on adapter iswta:
```

```
Tgt_PG  iSCSI Initiator Name / ISID
  2  iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e18f27efd5686b /
00:02:3d:00:00:01
  2  iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e28f27efd5686b /
00:02:3d:00:00:01
  2  iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e38f27efd5686b /
00:02:3d:00:00:01
  2  iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e48f27efd5686b /
00:02:3d:00:00:01
```

**3-** Since the iSCSI target can see the iSCSI initiator, the next step is to create an igroup on the storage unit for the blade server. To create the needed iSCSI group on the iSCSI target, enter the following command:

```
igroup create -i -t linux [groupName] [initiatorName]
```

groupName: the name to assign to the igroup.

initiatorName: the name of the iSCSI initiator obtained in step two of the section 4.1.

For example, the following command creates an igroup named as turtles for the first blade whose iSCSI initiator name is iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e18f27efd5686b:

```
turstore> igroup create -i -t linux turtles iqn.1987-
05.com.cisco.01.5a8b71fe3ed4b4c327e18f27efd5686b
```

To verify the igroup that was created, enter the following command.

```
turstore> igroup show
```

Depending on your igroup name, the resulting output should be similar to the example below:

```
turtles (iSCSI) (ostype: linux):  
    iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e18f27efd5686b
```

To add another iSCSI initiator into the igroup created above, enter the following command:

```
igroup add [groupName] [initiatorName]
```

*groupName*: the name of the igroup.

*initiatorName*: the name of the iSCSI initiator to be added to the igroup.

For example, the following command adds the iSCSI initiator whose initiator name is iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e28f27efd5686b into the igroup named as **turtles**:

```
turstore> igroup add turtles iqn.1987-  
05.com.cisco.01.5a8b71fe3ed4b4c327e28f27efd5686b
```

After adding the iSCSI initiator into the igroup, verify it was added by entering the following command:

```
turstore> igroup show
```

The resulting output will be similar to the example below:

```
turtles (iSCSI) (ostype: linux):  
    iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e18f27efd5686b  
    iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e28f27efd5686b
```

**4-** After creating the igroup, the next step is to create the LUNs on the storage unit and to map them to the igroup created in step 3 in order to make them accessible to the iSCSI initiator.

Enter the following command to create the lun:

```
lun create -s [size] -t linux [LUN_name]
```

*Size:* The size of the LUN that is to be created.

*LUN\_name:* The name to assign to the LUN.

Enter the following command to create the LUN **tur01** whose size is 55 GB on a volume named turtles:

```
lun create -s 55g -t linux /vol/turtles/tur01
```

To verify the LUN that was created above, enter the following command:

```
turstore> lun show
```

Depending on your volume name and LUN name, the output should be similar to the example below:

```
/vol/turtles/tur01    55.0g (59057897472) (r/w, online)
```

By changing the parameters above, you can create a LUN that has the size that you need, and by repeating you can continue creating LUNs as needed.

For example, the following command creates another LUN named **tur02** whose size is 55 GB on the same volume named turtles:

```
lun create -s 55g -t linux /vol/turtles/tur02
```

To verify all LUNs that are created for the iSCSI initiators, enter the following command:

```
turstore> lun show
```

Depending on the number of the created LUNs, the output should be similar to the example below:

```
/vol/turtles/tur01    55.0g (59057897472) (r/w, online)
/vol/turtles/tur02    55.0g (59057897472) (r/w, online)
```

Execute the following command to map the LUNs to the igroup:

```
lun map [LUN_name] [igroupname] [LUN_ID]
```

*LUN\_name*: The name of the LUN that is to be mapped to the igroup.

*igroupname*: The name of the igroup that the LUN is to be mapped to.

*LUN\_ID*: The number that is to be assigned to the mapped LUN.

Enter the following command to map the LUN named **tur01** to the igroup named **turtles** with the LUN\_ID **0**:

```
turstore> lun map /vol/turtles/tur01 turtles 0
```

To verify the mapped LUN, enter the following command:

```
turstore> lun show
```

The output should be similar to the example below:

```
/vol/turtles/tur01      55.0g (59057897472) (r/w, online, mapped)
```

Continue mapping all LUNs created on the previous step. For example, enter the following command in order to map the second LUN created on the previous step:

```
turstore> lun map /vol/turtles/tur02 turtles 1
```

To verify all LUNs that are created for the iSCSI initiators, enter the following command:

```
turstore> lun show -m
```

The output should be similar to the example below:

LUN path	Mapped to	LUN ID
-----		
/vol/turtles/tur01	turtles	0
/vol/turtles/tur02	turtles	1

**An alternative way to create & map a LUN:**

You can also create and map a LUN by entering the following command:

```
turstore> lun setup
```

The screen will be similar to the example below:

This setup will take you through the steps needed to create LUNs and to make them accessible by initiators. You can type ^C (Control-C) at any time to abort the setup and no unconfirmed changes will be made to the system.

Do you want to create a LUN? [y]: **y**

After entering the command and answering 'y' to the first question to create and map a LUN, you should enter the necessary information on the resulting screen:

OS type of LUN (default/solaris/windows/linux) [default]: **linux**

A LUN path must be absolute. A LUN can only reside in a volume or qtree root. For example, to create a LUN with name "lun0" in the qtree root /vol/vol1/q0, specify the path as "/vol/vol1/q0/lun0".

Enter LUN path: **/vol/turtles/tur01**

A LUN can be created with or without space reservations being enabled.

Space reservation guarantees that data writes to that LUN will never

fail.

Do you want the LUN to be space reserved? [y]: **y**

Size for a LUN is specified in bytes. You can use single-character multiplier suffixes: b(sectors), k(KB), m(MB), g(GB) or t(TB).

Enter LUN size: **55g**

You can add a comment string to describe the contents of the LUN.

Please type a string (without quotes), or hit ENTER if you don't want to supply a comment.

Enter comment string:

The LUN will be accessible to an initiator group. You can use an existing group name, or supply a new name to create a new initiator group. Enter '?' to see existing initiator group names.

Name of initiator group [turtles]: **turtles**

turtles (iSCSI) (ostype: linux):

iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e18f27efd5686b

iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e28f27efd5686b

iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e38f27efd5686b

iqn.1987-05.com.cisco.01.5a8b71fe3ed4b4c327e48f27efd5686b

The LUN will be accessible to all the initiators in the initiator group. Enter '?' to display LUNs already in use by one or more initiators in group "turtles".

LUN ID at which initiator group "turtles" sees "/vol/turtles/tur01" [0]: **0**

LUN Path : /vol/turtles/tur01

OS Type : linux

Size : 55g (59057897472)

Initiator Group : turtles

Initiator Group Type : iSCSI

Initiator	Group	Members	:	
				iqn.1987-
05.com.cisco.01.5a8b71fe3ed4b4c327e18f27efd5686b				

	:			iqn.1987-
05.com.cisco.01.5a8b71fe3ed4b4c327e28f27efd5686b				

	:			iqn.1987-
05.com.cisco.01.5a8b71fe3ed4b4c327e38f27efd5686b				

	:			iqn.1987-
05.com.cisco.01.5a8b71fe3ed4b4c327e48f27efd5686b				

Mapped to LUN-ID : 0

Do you want to accept this configuration? [y]: **y**

Do you want to create another LUN? [n]: **y**

Type 'y' to the last question and repeat the process for the other iSCSI initiators.

To verify all LUNs that have been created for the iSCSI initiator(s) by the alternative way, enter the following command:

```
turstore> lun show -m
```

The output below represents the mapped LUNs on the iSCSI target of the iSCSI test environment:

LUN path	Mapped to	LUN ID
/vol/turtles/tur01	turtles	0
/vol/turtles/tur02	turtles	1
/vol/turtles/tur03	turtles	2
/vol/turtles/tur04	turtles	3

### 4.3 Creating a File System on the iSCSI Initiator

This is the last section of the iSCSI configuration. We will create the appropriate file system on the BladeCenter using the LUNs of the iSCSI target. After completing the setup of the new file system, we will mount the LUNs of the target unit to the **/home** folder of the blades to be used by DB2 UDB:

1- Verify that the blades can see the LUNs by entering the following commands on each blade:

```
turtle01:~ # cat /proc/scsi/scsi
```

The resulting output should be similar to the example below:

```
Attached devices:
Host: scsi2 Channel: 00 Id: 00 Lun: 00
  Vendor: TEAC   Model: FD-05PUB   Rev: 2000
  Type: Direct-Access          ANSI SCSI revision: 02
Host: scsi3 Channel: 00 Id: 00 Lun: 00
  Vendor: TEAC   Model: CD-224E   Rev: 2.9B
  Type: CD-ROM          ANSI SCSI revision: 02
Host: scsi4 Channel: 00 Id: 00 Lun: 00
Vendor: NETAPP Model: LUN Rev: 0.2
```



```

Type: Direct-Access          ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 00 Lun: 01
Vendor: NETAPP Model: LUN      Rev: 0.2
Type: Direct-Access          ANSI SCSI revision: 03

```

The default location for the device files of the LUN devices on the iSCSI initiator server is `/dev/iscsi/bus0/target0/`.

Enter the following command on each blade to see the LUNs:

```
turtle01:~ # ls -l /dev/iscsi/bus0/target0/
```

Since we configured two LUNs, the resulting output should be similar to the example below:

```

drwxr-xr-x 2 root root 4096 ... lun0
drwxr-xr-x 2 root root 4096 ... lun1

```

**2-** After ensuring that each blade can see the LUNs defined on the iSCSI target, it is time to create the partition table on each LUN. We have two LUNs and two initiators, so we will configure them in the following order and match:

LUN ID	Host Name on Test Environment
0	Turtle01
1	Turtle02

As seen on the table, the first blade will use LUN 0. To create a partition table on any LUN, enter the following command:

```
fdisk [Devicename]
```

DeviceName: The device name assigned to the LUN. As an example, the default device names for the LUNs created in step 4 are:

```

LUN 0: /dev/iscsi/bus0/target0/lun0/disk
LUN 1: /dev/iscsi/bus0/target0/lun1/disk

```

For example, to create a partition table on LUN 0, enter the following command:

```
turtle01:~ # fdisk
/dev/iscsi/bus0/target0/lun0/disk
```

In the resulting screen, shown below, you will be prompted to answer a series of questions. (The appropriate responses are shown in **bold**.)

```
Command (m for help): n
Command action
  e  extended
  p  primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-56322, default 1): [enter]
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-56322, default 56322):
[enter]
Using default value 56322

Command (m for help): w
```

Repeat this process on each blade until the partition tables have been created on all present LUNs.

**3-** To create and configure a file system on each LUN, enter the following commands:

```
mkfs [DeviceName]
```

*DeviceName*: The device name assigned to the LUN.

For example, to create a new file system on LUN 0, enter the command below:

```
turtle01:~ # mkfs /dev/iscsi/bus0/target0/lun0/disk
```

The resulting output will be similar to the example below:

```
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
7223040 inodes, 14417920 blocks
720921 blocks (5.00%) reserved for the super user
First data block=0
440 block groups
32768 blocks per group, 32768 fragments per group
16416 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736,
1605632, 2654208, 4096000, 7962624, 11239424
Writing inode tables: done
Writing superblocks and filesystem accounting information: done
```

To convert the filesystem to a journaling filesystem, enter the following command:

```
tune2fs -j [DeviceName]
```

*DeviceName:* The device name assigned to the LUN.

As an example, the command below converts the filesystem (ext2) of the LUN 0 to the journaling file system (ext3):

```
turtle01:~ # tune2fs -j
/dev/iscsi/bus0/target0/lun0/disk
```

The resulting output will be similar to the example below:

```
Creating journal inode: done
This filesystem will be automatically checked every 22 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override
```

Repeat this process on each blade until the file systems have been created on the other present LUNs.

**4-** The last step is mounting the LUN devices to the */home* folder on each blade. Using the appropriate LUNID, enter the following command on each blade server:

```
mount [DeviceName] /home
```

*DeviceName:* The device name assigned to the LUN.

For example to mount the LUN 0 to the *'/home'* folder on the first blade, enter the following command:

```
turtle01:~ # mount /dev/iscsi/bus0/target0/lun0/disk /home
```

To check the mount point, enter the following command:

```
turtle01:~ # mount | grep home
```

The resulting output should be similar to the example below:

```
/dev/sdb on /home type ext3 (rw)
```

To mount the LUNs permanently, we will use the iSCSI *fstab* file present on each blade server. For example, edit ***/etc/fstab.iscsi*** using a text editor such as *'vi'* and add the following line to the end of the file on the first blade server:

```
dev/iscsi/bus0/target0/lun0/disk /home ext3 rw 0 0
```

To check the file, enter the following command:

```
turtle01:~ # vi /etc/fstab.iscsi
```

The resulting output should be similar to the example below:

```
# /etc/fstab.iscsi file for filesystems built on iscsi devices.  
# A typical entry here would look like:  
# /dev/iscsi/bus0/target0/lun0/disk /mnt ext2 defaults 0 0  
dev/iscsi/bus0/target0/lun0/disk /home ext3 rw 0 0
```

Repeat this process on each blade until the */etc/fstab.iscsi* file on each server has the appropriate mounting information for mounting point ***/home***.

## 5. CONFIGURING DB2 UDB FOR LINUX iSCSI

This section includes the key points of configuration while installing DB2 UDB on Linux iSCSI environment.

### 5.1 Installing DB2 UDB

To install DB2 UDB, use the standard installation process. No extra tool or component is needed for the Linux iSCSI environment.

While installing DB2 UDB, the key point depends on your database needs. If you configured Linux iSCSI on more than one blade server to use a partitioned database environment, DB2 UDB settings on each blade server must be identical.

Generally, while creating DB2 UDB instances on the blades, the Group ID (gid) and the User ID (uid) values must be the same on all of the blades configured as iSCSI initiators. Otherwise, Linux iSCSI for DB2 UDB will not work because of permission errors.

The following tables demonstrate the group IDs (gid) and user IDs (uid) of DB2 UDB on four blade servers configured as iSCSI initiators on the DB2 test environment for Linux iSCSI:

Group ID	Group Name	Group Members
101	Dasadm1	dasusr1,db2inst1
102	Db2fgrp1	db2fenc1
103	Db2grp1	db2inst1

User ID	User Name
105	Dasusr1
106	Db2inst1
107	Db2fenc1

To learn more about installing DB2 UDB with different options, refer to the following links:

<http://www.ibm.com/db2/linux>

<http://www.ibm.com/software/data/db2/linux/papers.html>

### 5.2 Testing DB2 UDB

To test IBM DB2 UDB on the Linux iSCSI environment, we will create a database called db2iscsi and generate a simple transaction.

**1-** To create the test database, log on as the db2 instance user:

```
turtle01:~ # su - db2inst1
```

**2-** Enter the following commands to activate DB2 UDB:

```
db2inst1@turtle01:~> db2start
```

You should get one of the two following messages:

```
0 0 SQL1063N DB2START processing was successful.  
SQL1063N DB2START processing was successful.
```

```
0 0 SQL1026N The database manager is already active.  
SQL1026N The database manager is already active.
```

Any other message will probably indicate that your instance is not correctly configured. In that case, check the log files for errors and re-configure your DB2 UDB instance again.

**3-** Create the folder 'db2test' on /home/db2inst1 (assuming that the name of the instance user is *db2inst1*), which will contain the test database:

```
db2inst1@turtle01:~> mkdir db2test
```

**4-** Create the test database 'db2iscsi':

```
db2inst1@turtle01:~>db2 create database db2iscsi on  
/home/db2inst1/db2test
```

The resulting output should be similar to the example below:

```
DB20000I The CREATE DATABASE command completed successfully.
```

**5-** Connect to the database:

```
db2inst1@turtle01:~>db2 connect to db2iscsi
```

#### Database Connection Information

```
Database server      = DB2/LINUX 8.1.2  
SQL authorization ID = DB2INST1  
Local database alias = DB2HA
```


**6- Check the status of the database:**

```
db2inst1@turtle01:~>db2 list active databases
```

```
Active Databases

Database name           = DB2ISCSI
Applications connected currently = 1
Database path           =
/db2iscsi/db2inst1/NODE0000/SQL00001/
```

If you are able to see the result above, you have successfully completed the test. To stay updated about iSCSI, you can find additional resources by searching linux iSCSI on the web. Also you can find the most recent versions of the iSCSI drivers at '<http://sourceforge.net/projects/linux-iscsi/>' .



© Copyright IBM Corporation 2004  
IBM Canada  
8200 Warden Avenue  
Markham, ON  
L6G 1C7  
Canada

Printed in United States of America 8-03  
All Rights Reserved.

IBM, DB2, DB2 Universal Database, and the ebusiness logo are trademarks or registered trademarks of the International Business Machines Corporation in the United States, other countries, or both.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

The information in this white paper is provided AS IS without warranty. Such information was obtained from publicly available sources, is current as of 07/30/2003, and is subject to change. Any performance data included in the paper was obtained in the specific operating environment and is provided as an illustration. Performance in other operating environments may vary. More specific information about the capabilities of products described should be obtained from the suppliers of those products.