IFT-3101U2G

Ultra2-to-Ultra2 SCSI RAID Controller

Instruction Manual

Revision 1.1

FCC ID:LC83101U2G

Infortrend

Federal Communications Commission (FCC) Statement

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

-- Reorient or relocate the receiving antenna.

-- Increase the separation between the equipment and receiver.

-- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.

-- Consult the dealer or an experienced radio/TV technician for help.

Warning: A shielded-type power cord is required in order to meet FCC emission limits and also to prevent interference to the nearby radio and television reception. It is essential that only the supplied power cord be used.

Use only shielded cables to connect I/O devices to this equipment.

You are cautioned that changes or modifications not expressly approved by the party responsible for compliance could void your authority to operate the equipment.

Chapter 1 Introduction

The IFT-3101U2G is an Ultra2-to-Ultra2 SCSI RAID controller specifically designed to provide RAID 0, 1(0+1), 3, 5 or JBOD capability to any host system equipped with a SCSI interface. It is totally independent of the host system's operating system. All RAID functions are performed by an AMD 5x86 CPU coupled with high-speed DRAMs and firmware in flash memory. In effect, it endows the host system with a high-speed and fault-tolerant disk storage operation using the RAID technology. It is an ideal solution for weaving several hard disks into one contiguous volume.

The controller has comprehensive drive failure management that allows automatic reassignment of reserved blocks when a bad sector is encountered during a write. Hot-swapping is supported through automatic disconnection of a failed drive and detection of a reserved drive followed with background rebuilding of data. The controller also supports spare drive operation. What's remarkable is all these failure recovery procedures are transparent to the host system.

The controller has been designed with ease of integration and maintenance in mind. All major features are described in the next chapter. The IFT-3101U2G already includes all the major operational requirements of a RAID subsystem. The overall features of a fully-built RAID subsystem will, however, depend on the actual components used and the creativity of the integrator.

Chapter 2 Features

Five operating modes:

JBOD	Just a Bunch of Drives
Non-RAID	Disk Spanning
RAID-0	Disk Striping
RAID-1	Disk Mirroring and Striping (RAID 0+1)
RAID-3	Disk Striping with Dedicated Parity
RAID-5	Multiple Block Striping with Interspersed Parity

- ✓ Comprehensive failure management including:
  - Automatic bad sector reassignment
  - Hot-swapping
  - Spare drive operation (Supports both Global Spare and Local Spare)
  - Background rebuilding (Rebuild priority selectable)
  - Verify-after-Write supported on normal writes, rebuild writes and/or RAID initialization writes
- ✓ Works with any operating system without additional software drivers.
- ✓ 3.5" drive profile allows easy integration into external subsystem enclosures or directly into the host system's drive bay.
- ✓ Up to three drive channels (optional 2-channel upgrade) for a total of 45 connected drives.
- ✓ Supports up to 15 SCSI IDs per channel.
- ✓ Three optional upgrade daughterboards:
  - IFT-90U: 3 Ultra Wide, single-ended, SCSI channels
  - IFT-90U: 2 Ultra Wide, differential, SCSI channels
  - IFT-90U: 3 Ultra2 Wide SCSI channels
- ✓ Up to 8 logical drives, each with independent RAID modes; up to 32 LUN's per SCSI ID (multiple SCSI ID's per channel are supported)
- ✓ Up to 8 partitions per logical drive
- ✓ Logical drive can be assigned a name for ease of identification
- ✓ Number of drives for each logical drive has no limitation
- ✓ Dynamic mapping of LUNs to logical drives. Two or more LUNs can be mapped to the same logical drive for redundant host operation.

- ✓ Concurrent/Background logical drive initialization.
- ✓ Performance optimization for Sequential or Random I/O.
- ✓ Allows multiple drive failure and concurrent multiple drive rebuild of a RAID (0+1) logical drive.
- ✓ Configuration of individual SCSI target parameters.
- ✓ Controller can be assigned a name for ease of identification.
- ✓ Prior to first disk access, it allows adjustment of delay time during controller initialization to enhance compatibility with slow-initial drives
- ✓ All channels are Ultra2-Wide SCSI-2 (backward compatible to SCSI-1) and can be configured as either a host or drive interface
- ✓ Two or more SCSI channels can be simultaneously set as host interface for redundant host system operation
- ✓ Compatible and will automatically match any SCSI hard disks with SCSI-1, SCSI-2 or (Ultra)-Wide-SCSI (1 or 2) specification.
- ✓ Full Ultra2-Wide-SCSI-2 implementation including Tagged Command Queuing and Multi-Threaded I/O
- ✓ Uses AMD 5x86 CPU with all executable firmware downloaded into high-speed DRAM
- ✓ EDO DRAM supported for enhanced performance.
- ✓ Up to 128 Mbytes of intelligent Read-Ahead/Write-Back cache
- ✓ Firmware resides in easy-to-update flash memory
- ✓ Front panel LCD and push buttons for configuration and message display
- ✓ Supports TELNET with PPP protocol for remote administration
- ✓ GUI RAID Manager and RS-232 terminal interface for RAID management
- ✓ SAF-TE support
- ✓ Supports Fault-bus for enclosure management.

Chapter 3 Configuration Options

The advantages of RAID are: Availability, Capacity and Performance. Choosing the right RAID level and drive failure management can increase Availability, subsequently increasing Performance and Capacity. The IFT-3101U2G RAID controller provides complete RAID functionality and enhanced drive failure management.

3.1 RAID Management

RAID stands for Redundant Array of Inexpensive Drive. The advantages of using a RAID storage subsystem are:

- Provides disk spanning by weaving all connected drives into one single volume.
- Increases disk access speed by breaking data into several blocks when reading/writing to several drives in parallel. With RAID, storage speed increases as more drives are added.
- Provides fault-tolerance by mirroring or parity operation.

What are the RAID levels?

RAID Level	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
NRAID	Non-RAID	1		Drive	Drive
RAID 0	Disk Striping	N	==NRAID	R: Highest W: Highest	R: High W: Highest
RAID 1 (0+1)	Mirroring Plus Striping (if N>1)	N+1	>>NRAID ==RAID 5	R: High W: Medium	R: Medium W: Low
RAID 3	Striping with Parity on dedicated disk	N+1	>>NRAID ==RAID 5	R: High W: Medium	R: Medium W: Low
RAID 5	Striping with interspersed parity	N+1	>>NRAID ==RAID 5	R: High W: Medium	R: High W: Low

NRAID

Disk Spanning

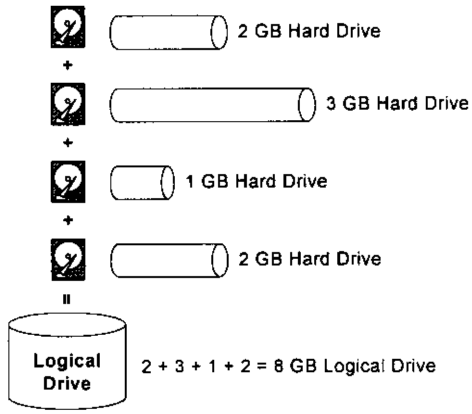

NRAID	
Minimum Disks required	1
Capacity	N
Redundancy	No

NRAID stands for Non-RAID. The capacity of all the drives are combined to become one logical drive (no block striping). In other words, the capacity of the logical drive is the total capacity of the physical drives. NRAID does not provide data redundancy.

JBOD

Single Drive Control

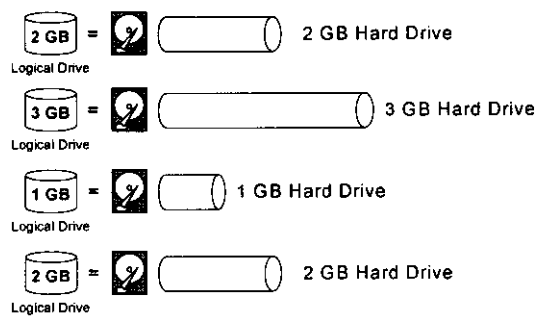

JBOD	
Minimum Disks required	1
Capacity	1
Redundancy	No

JBOD stands for Just a Bunch of Drives. The controller treats each drive as a stand-alone disk, therefore each drive is an independent logical drive. JBOD does not provide data redundancy.

RAID 0
Disk Striping

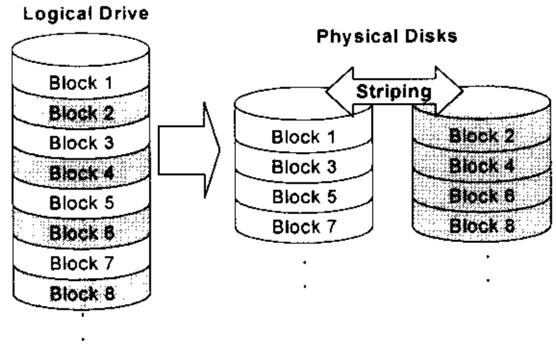

RAID 0	
Minimum Disks required	2
Capacity	N
Redundancy	No

RAID 0 provides the highest performance but no redundancy. Data in the logical drive is striped (distributed) across several physical drives.

RAID 1
Disk Mirroring

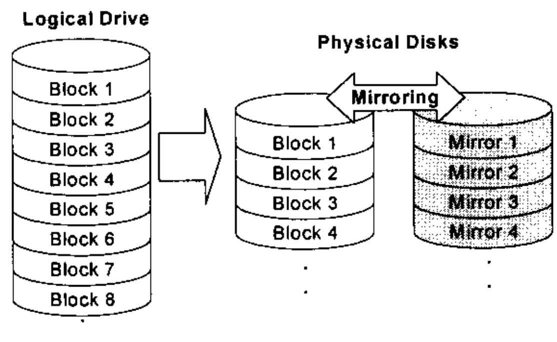

RAID 1	
Disks required	2
Capacity	N/2
Redundancy	Yes

RAID 1 mirrors the data stored in one hard drive to another. RAID 1 can only be performed with two hard drives. If there are more than two hard drives, RAID (0+1) will be performed automatically.

RAID (0+1)
Disk Striping with Mirroring

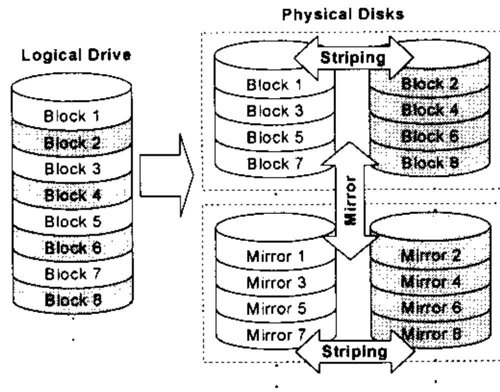

RAID (0+1)	
Minimum Disks required	4
Capacity	N/2
Redundancy	Yes

RAID (0+1) combines RAID 0 and RAID 1 - Mirroring and Striping. RAID (0+1) allows multiple drive failure because of the full redundancy of the hard drives. If there are more than two hard drives assigned to perform RAID 1, RAID (0+1) will be performed automatically.

IMPORTANT:
"RAID (0+1)" will not appear in the list of RAID levels supported by the controller. If you wish to perform RAID 1, the controller will determine whether to perform RAID 1 or RAID (0+1). This will depend on the drive number that has been selected for the logical drive.

RAID 3
Disk Striping with Dedicated Parity Disk

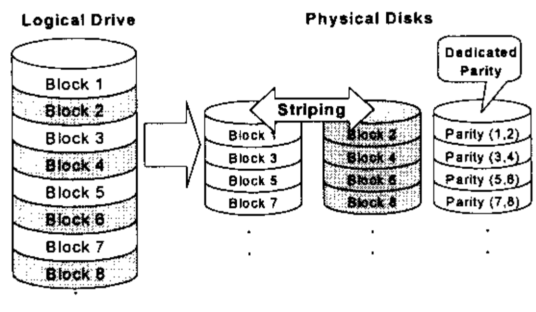

RAID 3	
Minimum Disks required	3
Capacity	N-1
Redundancy	Yes

RAID 3 performs Block Striping with Dedicated Parity. One drive member is dedicated to storing the parity data. When a drive member fails, the controller can recover/regenerate the lost data of the failed drive from the dedicated parity drive.

RAID 5
Striping with
Interspersed Parity

RAID 5	
Minimum Disks required	3
Capacity	N-1
Redundancy	Yes

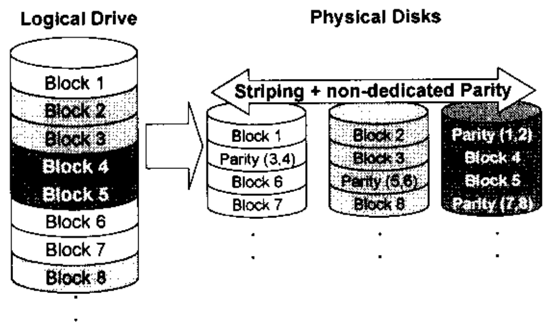

RAID 5 is similar to RAID 3 but the parity data is not stored in one dedicated hard drive. Parity information is interspersed across the drive array. In the event of a failure, the controller can recover/regenerate the lost data of the failed drive from the other surviving drives.

3.2 *Drive Failure Management*

3.2.1 Global and Local Spare Drive

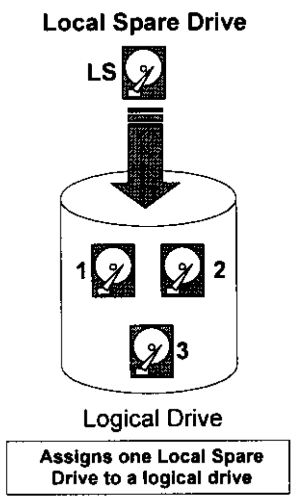

Local Spare Drive is a standby drive assigned to serve one specified logical drive. When a member drive of this specified logical drive fails, the Local Spare Drive becomes a member drive and automatically starts to rebuild.

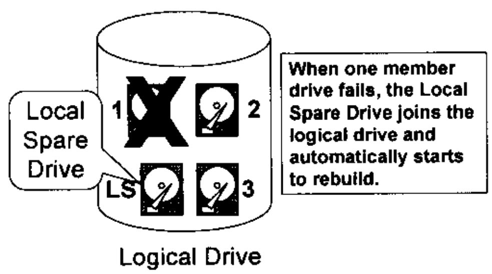

Global Spare Drive does not only serve one specified logical drive. When a member drive from any of the logical drive fails, the Global Spare Drive will join that logical drive and automatically starts to rebuild.

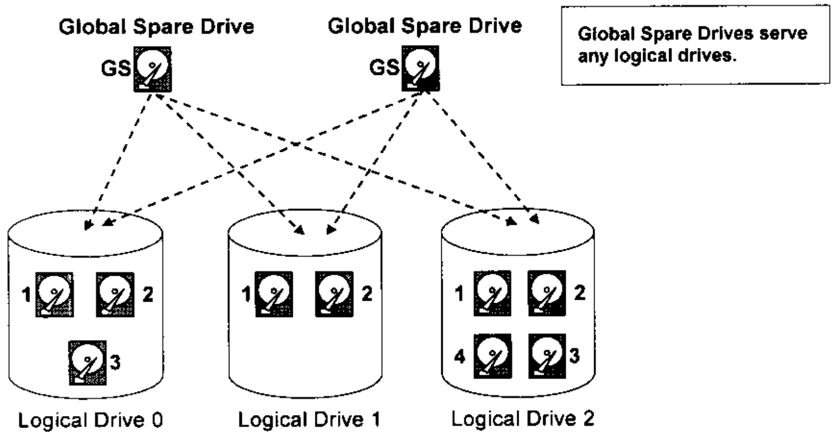

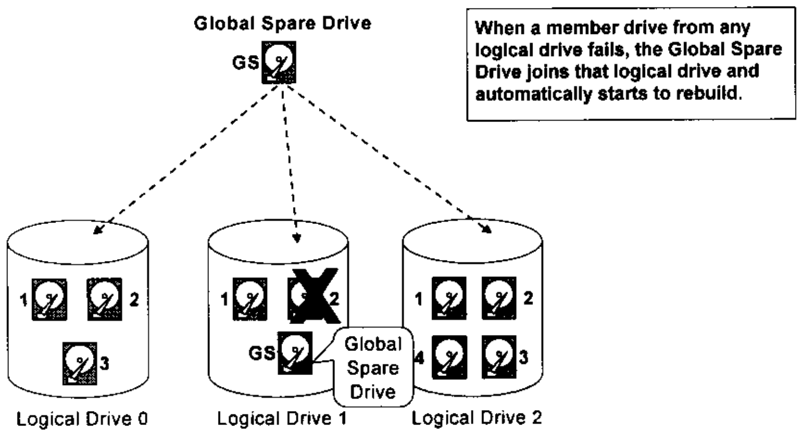

The IFT-3101U2G RAID controller provides both Local Spare Drive and Global Spare Drive functions. On certain occasions, applying these two functions together will better fit various needs. Take note though that the **Local Spare Drive always has higher priority than the Global Spare Drive.**

In the example shown below, the member drives in Logical Drive 0 are 9 GB drives, and the members in Logical Drives 1 and 2 are all 4 GB drives. It is not possible for the 4 GB Global Spare Drive to join Logical Drive 0 because of its insufficient capacity. However using a 9GB drive as the Global Spare drive for a failed drive that comes from Logical Drive 1 or 2 will bring huge amount of excess capacity since these logical drives require 4 GB only. In the settings below, the 9 GB Local Spare Drive will aid Logical Drive 0 once a drive in this logical drive failed. If the failed drive is in Logical Drive 1 or 2, the 4 GB Global Spare drive will immediately give aid to the failed drive.

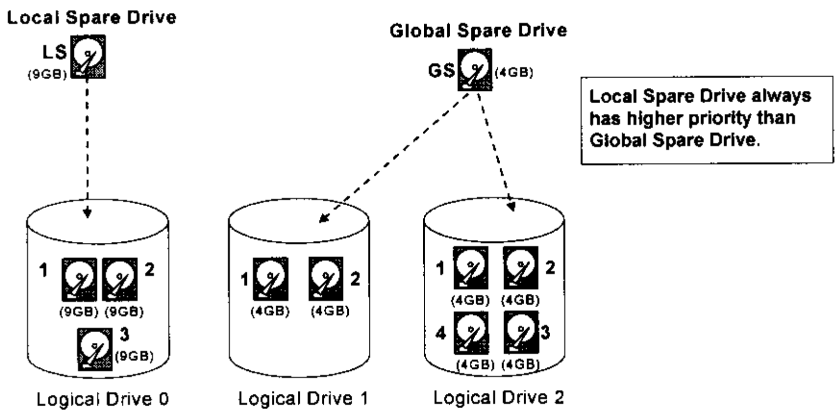

3.2.2 Identifying Drives

Assuming there is a failed drive in the RAID 5 logical drive, make it a point to replace the failed drive with a new drive to keep the logical drive working.

When trying to remove a failed drive and you mistakenly removed the wrong drive, you will no longer be able to read/write the logical drive because the two drives may have already failed.

To prevent this from happening, the controller provides an easy way of identifying for the failed drive. That is, the read/write LED of the failed hard drive will light. This LED will prevent you from removing the wrong drive, and is also helpful when locating for a drive.

Flash Selected SCSI Drive

The Read/Write LED of the drive you selected will light steadily for about one minute.

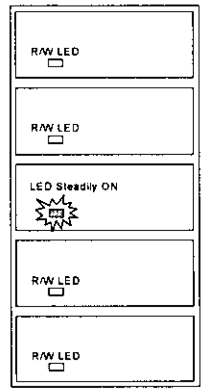

Flash All SCSI Drives

The Read/Write LED of all connected drives will light for about one minute. If the LED of the defective drive did not light on the "Flash Selected SCSI Drive" function, use "Flash All SCSI Drives". The "Flash All SCSI Drives" function will light LEDs of all the drives except the defective one.

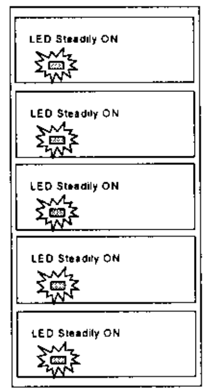

3.2.3 Automatic Rebuild and Manual Rebuild

Automatic Rebuild

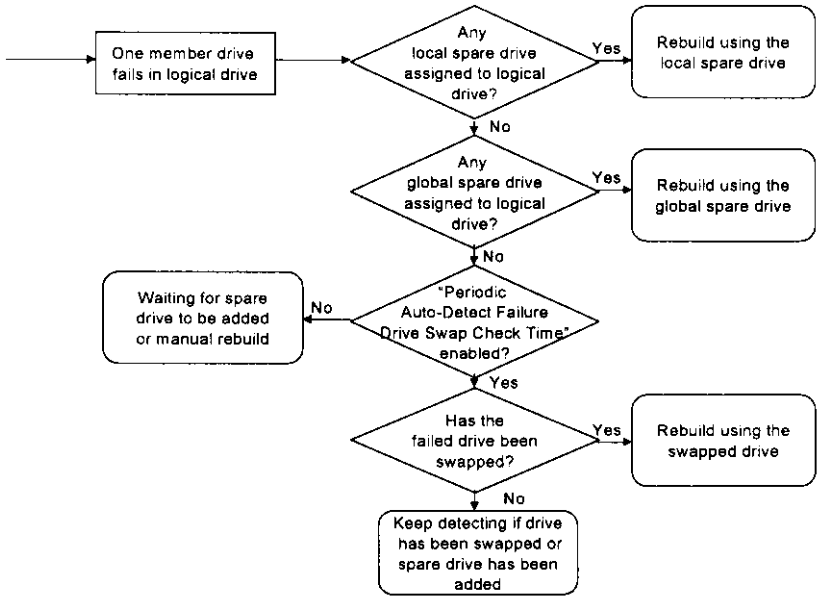

When a member drive in the logical drive failed, the controller will first check whether there is a Local Spare Drive assigned to this logical drive. If yes, it will automatically start to rebuild.

If there is no Local Spare Drive available, the controller will search for a Global Spare Drive. If there is a Global Spare Drive, it will automatically rebuild the logical drive.

If neither Local Spare Drive nor Global Spare Drive is available, the controller will not try to rebuild unless the user applies a forced-manual rebuild.

Manual Rebuild

When a user applies forced-manual rebuild, the controller will first check whether there is any Local Spare Drive assigned to this logical drive. If yes, it will automatically start to rebuild.

If there is no Local Spare Drive available, the controller will search for a Global Spare Drive. If there is a Global Spare Drive, it will automatically rebuild the logical drive.

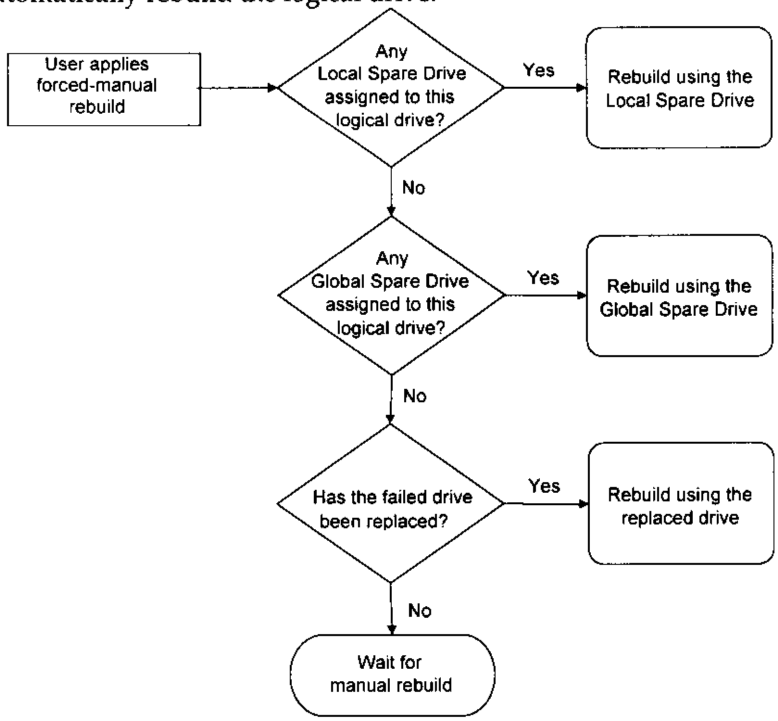

If neither Local Spare Drive nor Global Spare Drive is available, the controller will detect the SCSI channel and ID of the failed drive. Once the failed drive has been replaced by a new drive/used drive, it starts to rebuild using the replaced drive. If there is no available drive for rebuilding, the controller will not try to rebuild again until the user applies another forced-manual rebuild.

3.2.4 Concurrent Rebuild in RAID (0+1)

RAID (0+1) allows multiple drive failure and concurrent multiple drive rebuild. Newly replaced drives must be scanned and set as Local Spare Drives. These drives will be rebuilt at the same time (you do not need to repeat the rebuilding process for each drive).

3.3 Disk Array Parameters

3.3.1 Rebuild Priority

Rebuilding time will depend on the capacity of the logical drive. The IFT-3101U2G RAID controller provides background rebuilding ability. Meaning, the controller is able to serve other I/O requests while rebuilding the logical drives. The rebuilding process is totally transparent to the host computer or the operating system.

The background rebuild process has four priority options:

- Low
- Normal
- Improved
- High

The default priority is "Low" which uses the controller's minimum resources to rebuild. Choosing "Normal" or "Improved" will speedup the rebuilding process and choosing "High" will use the controller's maximum resources to complete the rebuilding process at the shortest time.

Rebuild priority can be configured through the RS-232C Terminal Interface, GUI RAID Manager or the front panel.

3.3.2 Verify-after-Write

The controller has the ability to force the hard drives to verify after data has been written to the media of the HDD. There are three selectable methods:

- Verification on LD Initialization Writes
  Performs Verify-after-Write while initializing the logical drive.
- Verification on LD Rebuild Writes
  Performs Verify-after-Write during the rebuilding process.
- Verification on LD Normal Drive Writes
  Performs Verify-after-Write during normal I/O requests.

Each method can be enabled or disabled individually. Hard drives will perform Verify-after-Write according to the selected method.

IMPORTANT:
The "Verification on LD Normal Drive Writes" method will affect "write" performance during normal use.

3.4 *Cache Parameters*

3.4.1 Optimization for Sequential or Random I/O

When using RAID with applications such as video or image oriented applications, the application reads/writes from the drive using large-block, sequential files instead of small-block, random access files. The IFT-3101U2G RAID controller provides the options to optimize for large-sequential I/O or optimize for small-random I/O access.

"Optimization for Sequential I/O" provides larger stripe size (block size, also known as Chunk size) than "Optimization for Random I/O". A lot of the controller's internal parameters will also be changed to optimize for sequential or random I/O. The change will take effect after the controller reboots.

If the existing logical drives were built with "Optimization for Random I/O", these logical drives will not read/write when using "Optimization for Sequential I/O" (shows "INVALID") and vice versa because the stripe size is different. Change it back to the original setting and reset the controller to make available the logical drive data again.

IMPORTANT:
Changing the setting to "Optimization for Sequential I/O" or "Optimization for Random I/O" should be performed only when no logical drive exist. Otherwise, you will not be able to access the data in the logical drive later on.

3.5 *Drive-Side SCSI Parameters*

3.5.1 SCSI Motor Spin-up

When the power supply is unable to provide sufficient current for all the hard drives and controllers that are powered-up at the same time, spinning-up the hard drives serially is one of the best way of consuming lower power-up current.

By default, all hard drives will spin-up when powered-on. These hard drives can be configured so that all of them will not spin-up at power-on. There are 3 methods of spinning-up the hard drive's motor: Spin-up at power-on, Spin-up serially in random sequence or Spin-up by SCSI command. Please refer to the hard drive's user's manual for instructions on configuring the hard drive using the "Spin-up by SCSI command". The procedure for each brand/model of hard drive should vary.

Configure all the hard drives as above and enable "SCSI Motor Spin-Up" in Drive-Side SCSI Parameters. Power off all hard drives and controller, and power them on again. All the hard drives will not spin-up at this time. The controller will then spin-up the hard drives one by one at four seconds interval.

IMPORTANT:
If the drives are configured as "Delay Motor Spin-up" or "Motor Spin-up in Random Sequence," some of these drives may not be ready yet for the controller to access when the system powers up. Increase the disk access delay time so that the controller will wait a longer time for the drive to be ready.

3.5.2 SCSI Reset at Power Up

By default, when the controller is powered up, it will send a SCSI bus reset command to the SCSI bus. When disabled, it will not send a SCSI bus reset command on the next power-up.

When connecting dual host computers to the same SCSI bus, the SCSI bus reset will interrupt all the read/write requests that are being performed. This may cause some operating systems or host computers to act abnormally. Disable the "SCSI Reset at Power-up" to avoid this situation.

3.5.3 Disk Access Delay Time

Sets the delay time before the controller tries to access the hard drives after power-on. The default is 15 seconds.

3.5.4 SCSI I/O Timeout

The "SCSI I/O Timeout" is the time interval that the controller waits for a drive to respond. If the controller attempts to read data from or write data to a drive, but the drive does not respond within the SCSI I/O timout value, the drive will be judged to be a failed drive.

When the drive itself detects a media error while reading from the drive platter, it will retry the previous reading or recalibrate the head. When the drive has encountered a bad block on the media, it has to reassign the bad block to another spare block. However, all of this takes time. The time to perform these operations can vary between different brands and models of drives.

During SCSI bus arbitration, a device with higher priority can utilize the bus first. A device with lower priority will sometimes get a SCSI I/O timeout when higher priority devices keep utilizing the bus.

The default setting for "SCSI I/O Timeout" is 7 seconds. It is highly recommended not to change this setting. Setting the timeout to a lower value will cause the controller to judge a drive as failed a drive is still retrying or while a drive is unable to arbitrate the SCSI bus. Setting the timeout to a greater value will cause the controller to keep waiting for a drive, and it may sometimes cause a host timeout.

3.5.5 Maximum Tag Count

The maximum number of tags that can be sent to each drive at the same time. A drive has a built-in cache that is used to sort all of the I/O requests ("tags") which are sent to the drive, allowing the drive to finish the requests faster. The cache size and maximum number of tags varies between different brands and models of drive. Using the default setting - "32" - is highly recommended. Changing the maximum tag count to "Disable" will cause the internal cache of the drive to be ignored (i.e., not used).

3.5.6 Periodic Drive Check Time

The "Periodic Drive Check Time" is an interval for the controller to check all of the drives that were on the SCSI bus at controller startup (a list of all the drives that were detected can be seen under "View and Edit SCSI Drives"). The default value is "Disabled". "Disabled" means that if a drive is removed from the bus, the controller will not be able to know - so long as no host accesses that drive. Changing the check time to any other value allows the controller to check - at the selected interval - all of the drives that are listed under "View and Edit SCSI Drives." If any drive is then removed, the controller will be able to know - even if no host accesses that drive.

3.5.7 SAF-TE Enclosure Monitoring

What is SAF-TE?

SAF-TE stands for SCSI Accessed Fault-Tolerant Enclosures. It is an enclosure management technology that uses the SCSI bus to interact with the controller. A SAF-TE-compliant enclosure monitors the fan temperature, power supply, UPS and also provides drive status LED's.

How does it work?

The SAF-TE device, which is often a back-plane within a drive-bay enclosure, must occupy a connector on one of the drive channels' SCSI cables. The presence of a SAF-TE device will be detected and its presence will be displayed in the BIOS configuration utility, Text RAID Manager and the GUI RAID Manager programs. The RAID controller communicates with the SAF-TE enclosure with standard SCSI commands, polling the device in order to get SAF-TE information.

The default value for "Periodic SAF-TE Device Check Time" is "Disabled". If the enclosure does have a SAF-TE device and features, enable the controller to poll the device by selecting a time interval. The RAID controller will then check the SAF-TE device status at that interval.

SAF-TE Support

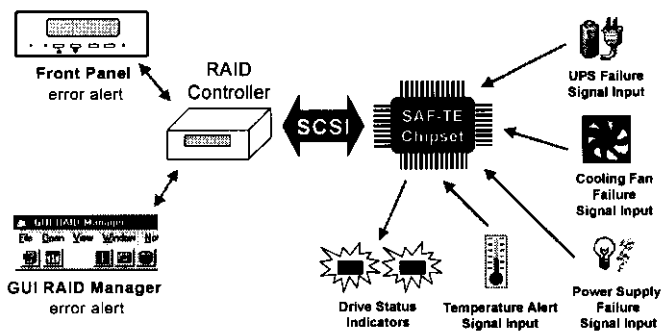

- SAF-TE chipset connects to the drive channel of the controller together with the other SCSI drives.

3.5.8 Periodic Auto-Detect Failure Drive Swap Check Time

The "Drive-Swap Check Time" is the interval at which the controller checks to see whether a failed drive has been swapped. When a logical drive's member drive fails, the controller will detect the failed drive (at the selected time interval). Once the failed drive has been swapped with a drive that has adequate capacity to rebuild the logical drive, the rebuild will begin automatically.

The default setting is "Disabled," meaning that that the controller will not Auto-Detect the swap of a failed drive. To enable this feature, select a time interval.

3.6 Dynamic Logical Drive Expansion

3.6.1 What Is It and How Does It Work?

Before Dynamic Logical Drive Expansion, increasing the capacity of a RAID system using traditional methods meant backing up, re-creating and then restoring. Dynamic Logical Drive Expansion allows users to add new SCSI hard disk drives and expand a RAID 0, 3 or 5 Logical Drive without powering down the system.

3.6.2 Two Modes of Dynamic Logical Drive Expansion

There are two modes of Dynamic Logical Drive Expansion: Mode 1 and Mode 2.

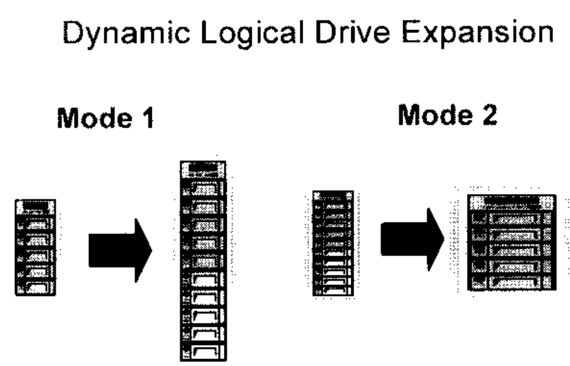

Mode 1 Expansion involves adding more SCSI hard disk drives to a logical drive, which may require that the user obtain an enclosure with more drive bays. The data will be re-striped onto the original and newly added disks.

RAID Expansion - Mode 1

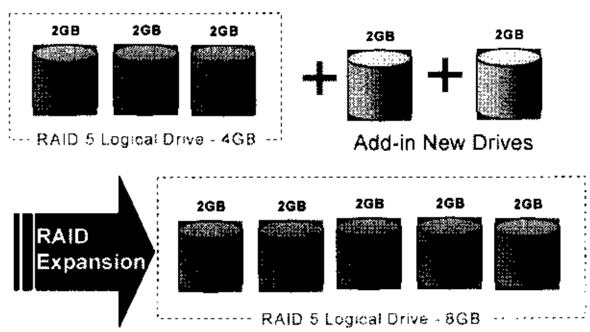

In the figure above, new drives are added to increase the capacity of a 4-Gigabyte RAID 5 logical drive. The two new drives increase the capacity to 8 Gigabytes.

Mode 2 Expansion, on the other hand, requires the same number of higher-capacity SCSI hard disk drives for a given logical drive.

RAID Expansion - Mode 2 (1/3)

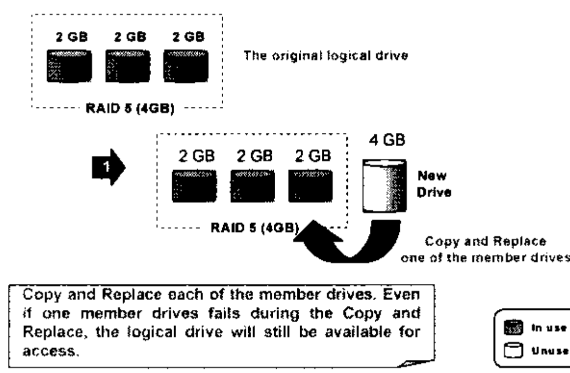

The figure above illustrates expansion of the same 4-Gigabyte RAID 5 logical drive using Mode 2 Expansion. Drives are copied and replaced, one by one, onto three higher-capacity drives.

RAID Expansion - Mode 2 (2/3)

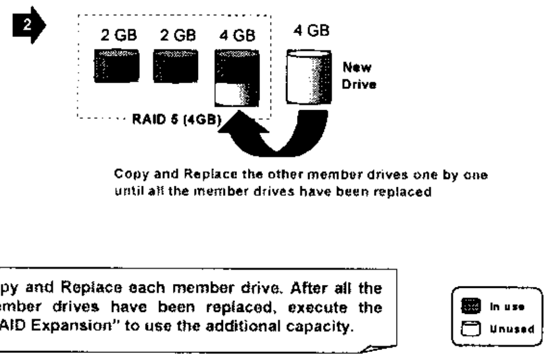

This results in a new 4-Gigabyte, RAID 5 logical drive composed of three physical drives. The 4 Gigabytes of increased capacity is in a new partition.

RAID Expansion - Mode 2 (3/3)

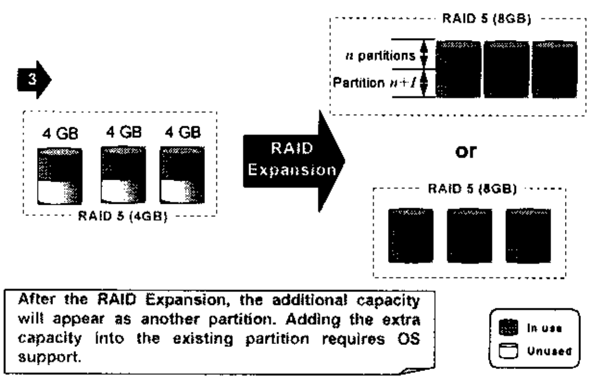

IMPORTANT:

- *The increased capacity from Mode 1 Expansion of a logical drive will be a new partition.*
- *At the time of this printing, the firmware does not support the "Copy and Replace" function that is required for Mode 2 Expansion. Third-party hard disk utilities may be used for Mode 2 Expansion of logical drives. Later versions of the*

firmware will support "Copy and Replace."

3.6.3 Example: RAID Expansion in Windows NT® Server

Limitations When Using Windows NT® 4.0

1. Only the Windows NT® Server Disk Administrator includes the Extend Volume Set function; Windows NT® Workstation does not have this feature.
2. The system drive (boot drive) of a Windows NT® system cannot be extended.
3. The drive that will be extended should be using the NTFS file system.

The Example:

The following example demonstrates the expansion of a 900MB RAID 0 logical drive. The HyperTerminal emulation software that comes with Windows 95®/Widnows NT® is used to connect to the RAID controller via RS-232.

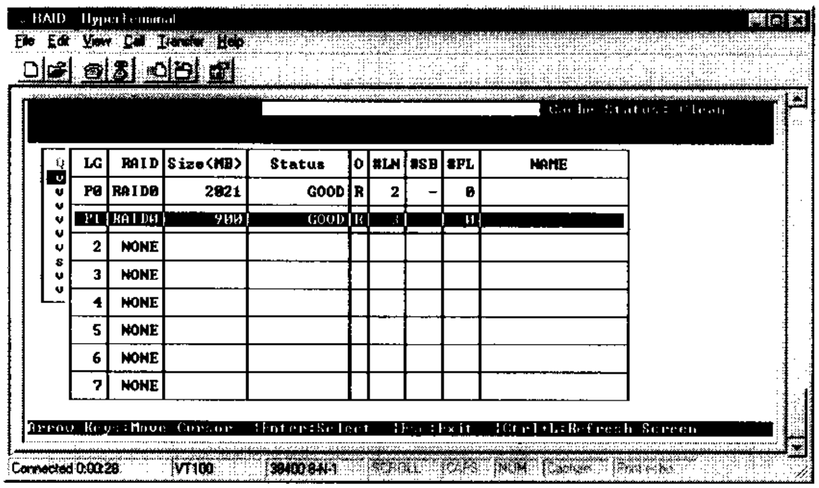

You can view information about this drive in the Windows NT® Server's Disk Administrator.

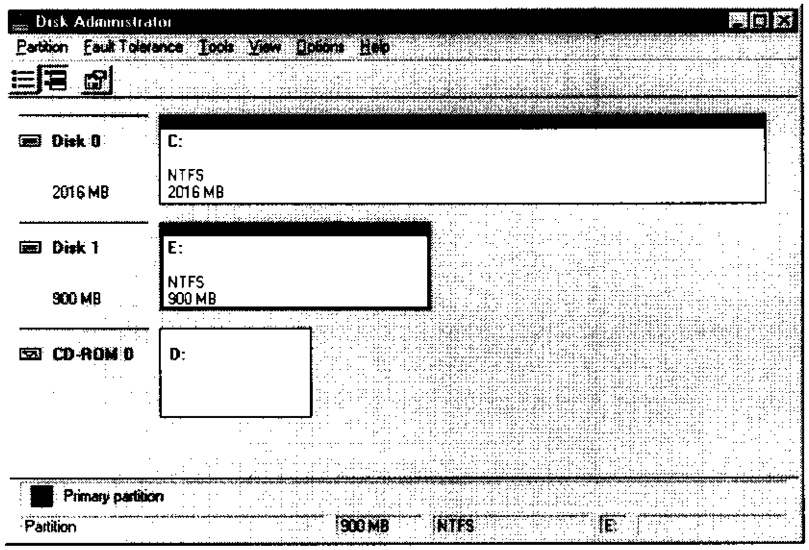

Place the cursor on Disk 1, right-click your mouse, and select "Properties." You will see that the total capacity for the Drive E: is just under 900MB.

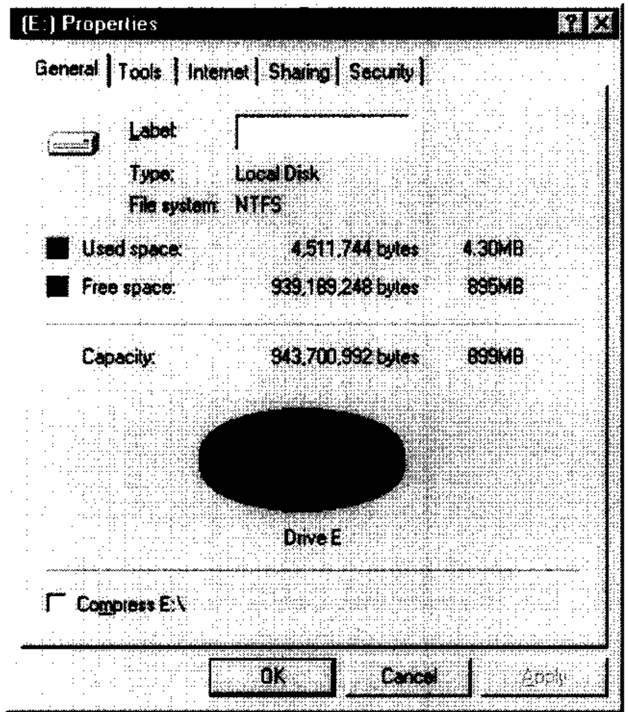

Follow the steps described in section 7.2.8 to add SCSI disk drives and perform Mode 1 Dynamic Logical Drive Expansion.

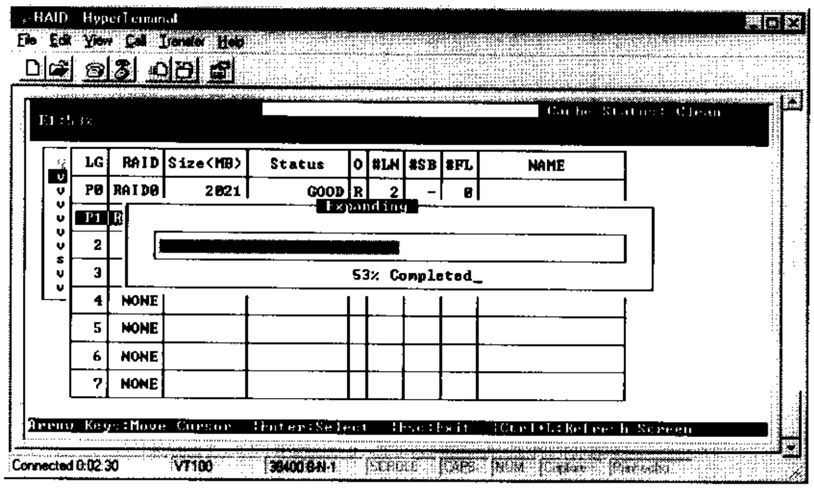

The 900MB logical drive has become a 1800MB logical drive. Place the cursor on that logical drive, and then press <Enter>.

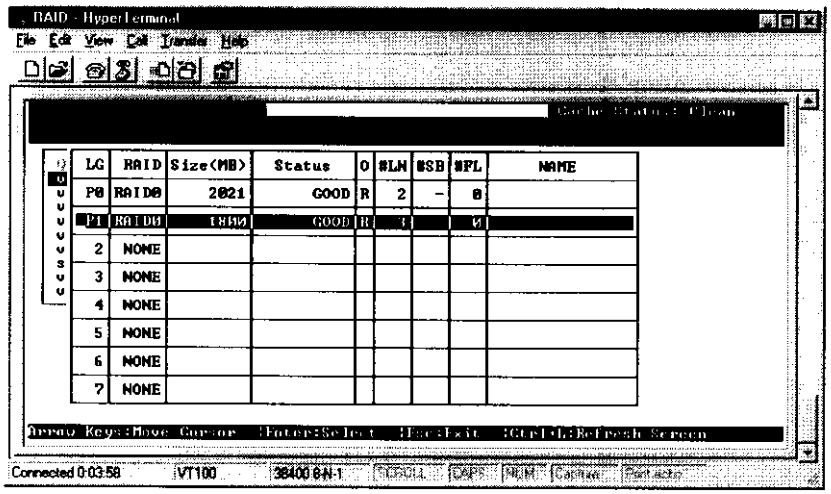

From the menu, select Partition Logical Drive. You will see that the 1800MB logical drive is composed of two 900MB partitions.

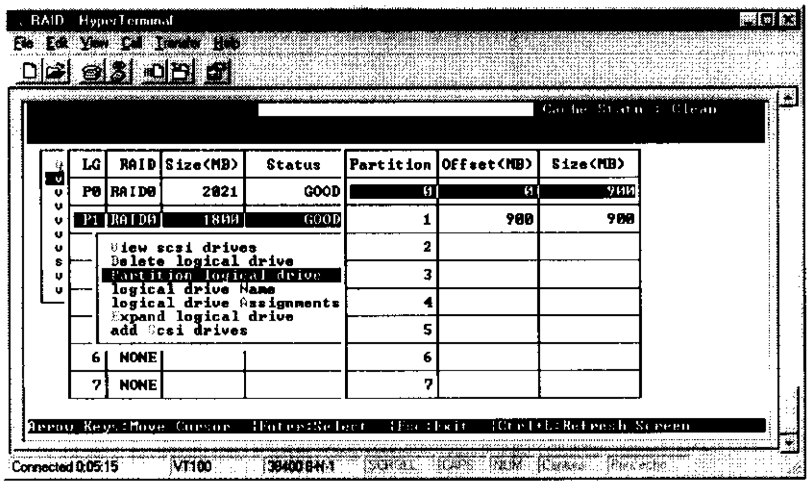

Follow the directions in section 7.3.1 to map the new partition to a Host LUN. The new partition must be mapped to a host LUN in order for the HBA (host-bus adapter) to see it. Once you have mapped the partition, reboot Windows NT®. The HBA should be able to detect an additional "disk."

Return to Windows NT® Server's Disk Administrator. There now exists a Disk 2 with 900MB of free space. Click on Disk 2 to select it.

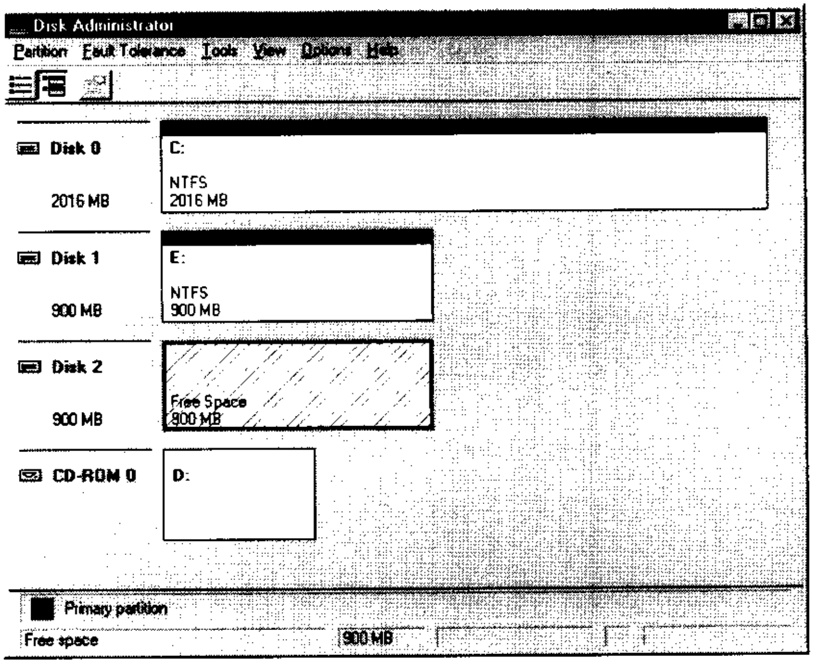

From the "Partition" menu, select "Extend Volume Set."

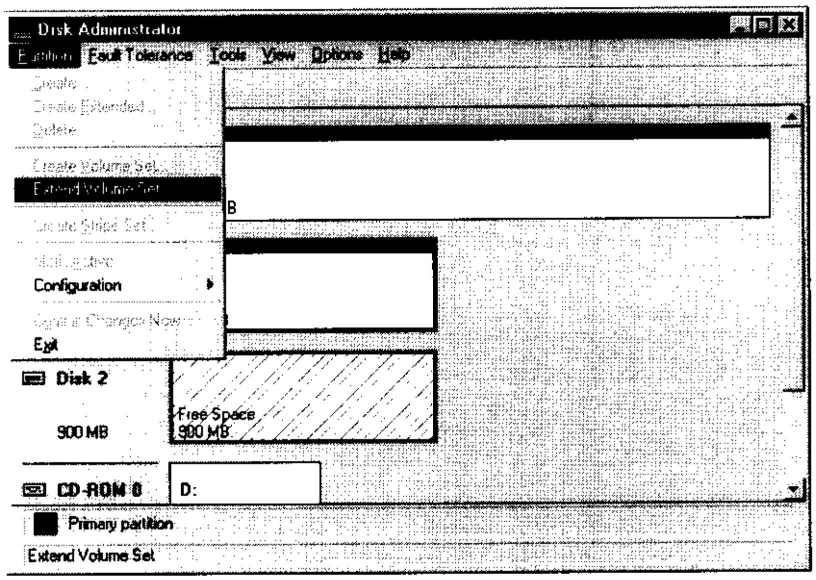

The screen will display that volume set of Drive E: has been extended by the 900MB in Disk2. Move the cursor to "Commit Changes Now" to confirm that you want the free space to become a part of the same logical drive.

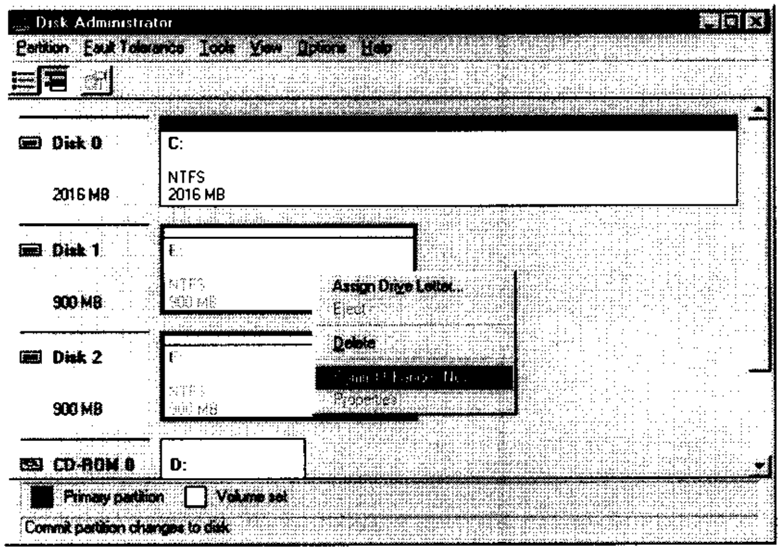

Logical Drive E: is now composed of two 900MB partitions with a total volume of 1800MB. To see this, hold down on the <Ctrl> key and select both Disk 1 and Disk2; then right-click your mouse and select "Properties."

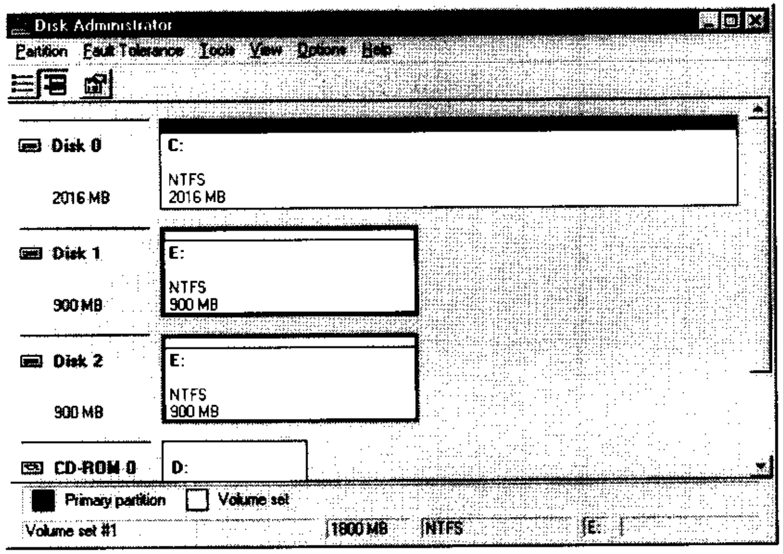

Drive E: now has a capacity just under 1800MB.

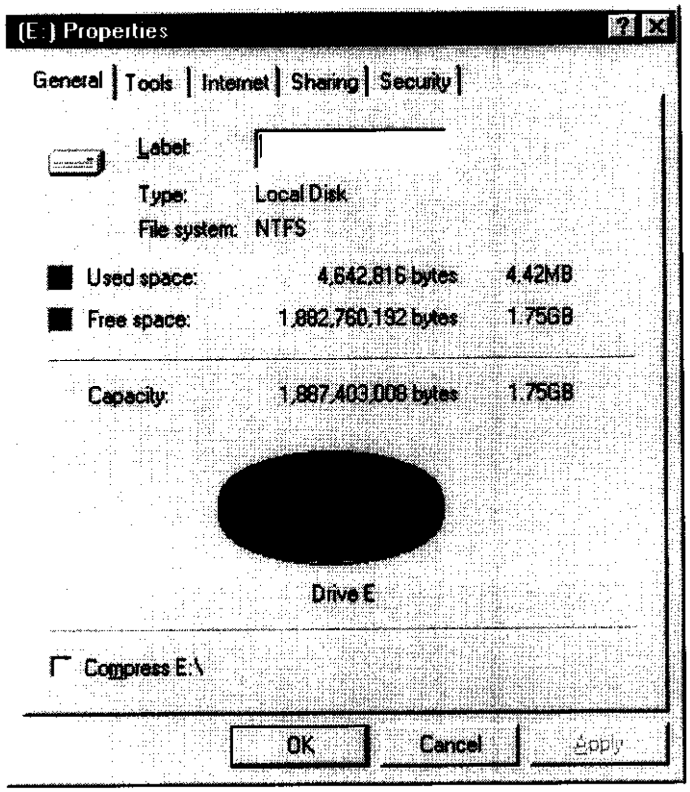

Chapter 4 Hardware Installation

4.1 Locations of the Parts

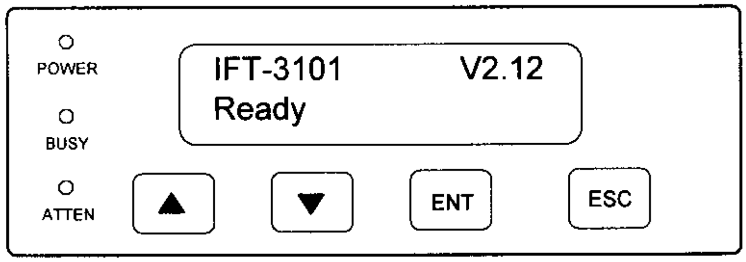

4.1.1 Front View

POWER	Lighted LED indicates power is on.
BUSY	Unlit indicates no activity. Blinking indicates data is being accessed. Lighted LED indicates unprocessed cached data is still in the memory.
ATTEN	Lights when an error message appears or service is required, e.g., when a drive fails and needs to be replaced.
▼ ▲ buttons	Scroll through available options.
ENT button	Choose or executes an option.
ESC button	Returns to previous menu or cancel selection.
2 x 16 LCD	Displays throughput during normal operation, approximately 256Kbytes/sec per division. Displays message for configuration and management.

4.1.2 The Main Board

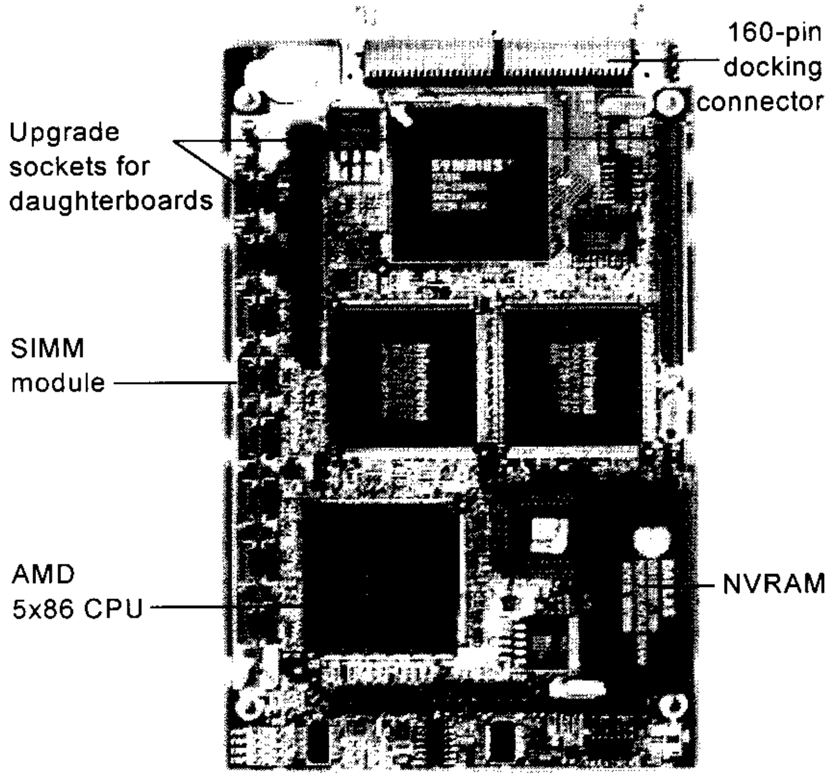

4.1.3 The Daughterboards

Top View of Ultra2 Wide SCSI Daughterboard (IFT-9152)

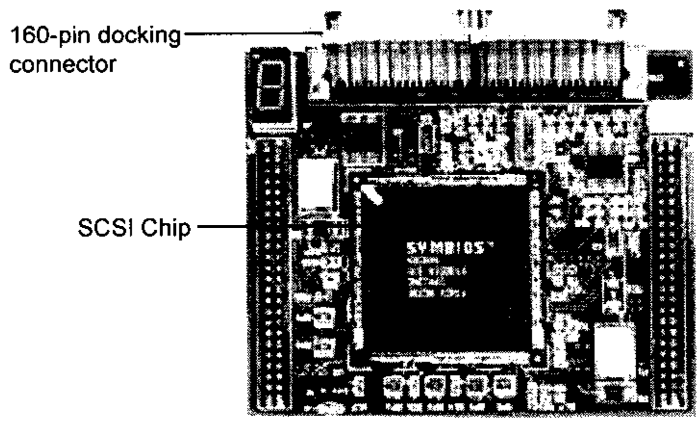

Top View of Ultra Wide Differential Daughterboard (IFT-9152D)

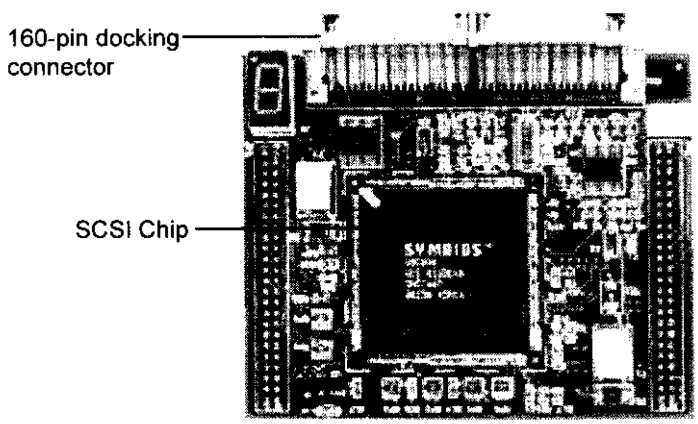

Top View of Fibre Daughterboard (IFT-9152F)

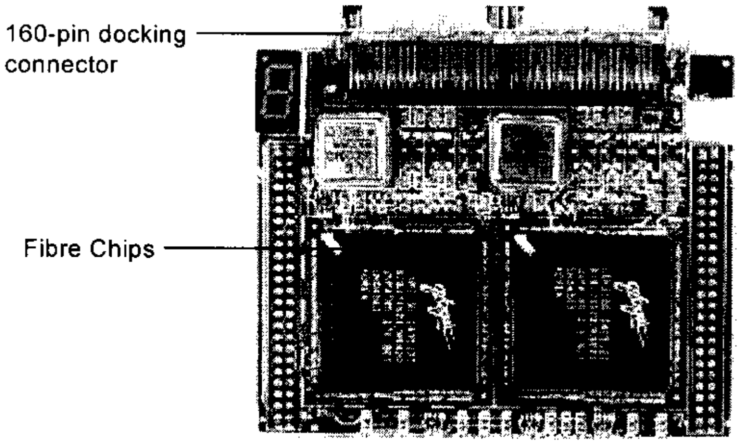

4.1.4 The Back-plane Boards

Back-plane Board for Main Board (IFT-9158B2A)

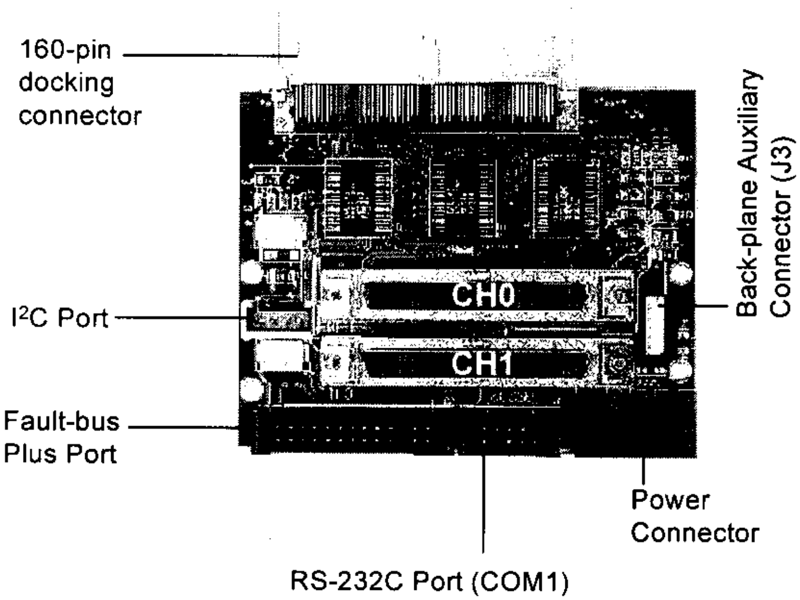

Back-plane Board for Ultra2 Wide SCSI Daughterboard (IFT-9158B2B)

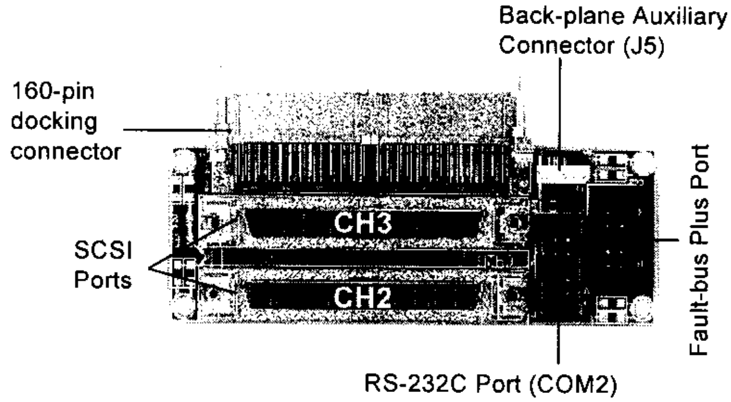

Back-plane Board for Ultra Wide Differential Daughterboard (IFT-9158B2D)

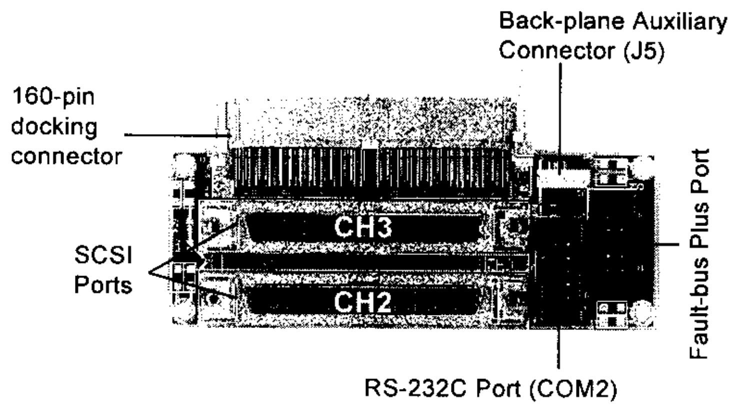

Back-plane Board for Fibre Channel Daughterboard (IFT-9158B2F)

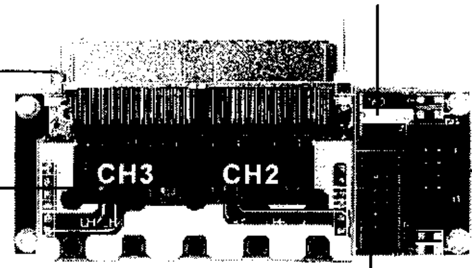

4.2 Installing DRAM SIMM

IMPORTANT:
The IFT-3101U2G controller requires a minimum of 8 Mbytes DRAM SIMM (with or without parity function) installed in a SIMM socket in order for it to operate. The controller is normally delivered without any DRAM installed.

The following are guidelines on using DRAMs:

- ✓ **Use 72-pin 60ns DRAM or 60ns EDO RAM SIMM module**. EDO RAM is recommended for better performance.
- ✓ A SIMM, with or without parity, can be auto-detected by the controller. A SIMM with parity is recommended for security.
- ✓ The minimum DRAM required is 8 MBytes installed in a SIMM socket, however 16 Mbytes is recommended.
- ✓ The controller supports 8 MB, 16 MB, 32 MB, 64 MB and 128MB DRAM SIMM modules. Maximum DRAM size is 128 MBytes.

To install the DRAM SIMM:

1. Power off the system and disconnect the power connector.
2. If a daughterboard has been installed, remove it prior to installing a SIMM module (or replacing an existing SIMM module).

3. Insert the DRAM SIMM vertically into the socket making sure the key is on the left side (1). Now push the module backward until the hooks on both sides of the socket snap into place (2).

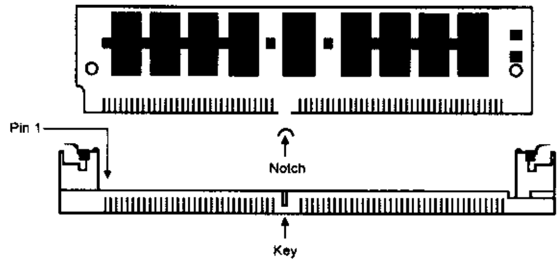

4.3 *Installing the Fibre/SCSI Channel Upgrade Daughterboard*

The IFT-3101U2G's base module has 2 Ultra2 Wide SCSI channels. Installing a SCSI channel upgrade daughter board (IFT-9152, IFT-9152D or IFT-9152F) onto the base module allows you to expand up to a total of 4 SCSI channels. (Only one daughterboard can be mounted at a time).

IFT-9152:	2 Ultra2 Wide SCSI channels
IFT-9152D:	2 Ultra Wide Differential SCSI channels
IFT-9152F:	2 single-loop Fibre channels

To install the SCSI channel upgrade daughterboard:

1. Make sure the power of the host system and drives are off.
2. Position the daughter board so that the SCSI connectors are facing the rear of the controller.
3. While at it, make sure the connector pins on the daughter board are aligned with the two header connectors on the controller's main board.
4. Press both sides of the daughter board downward so that the connector pins on the daughter board insert into the header connectors on the main board. Make sure the daughter board is seated properly.

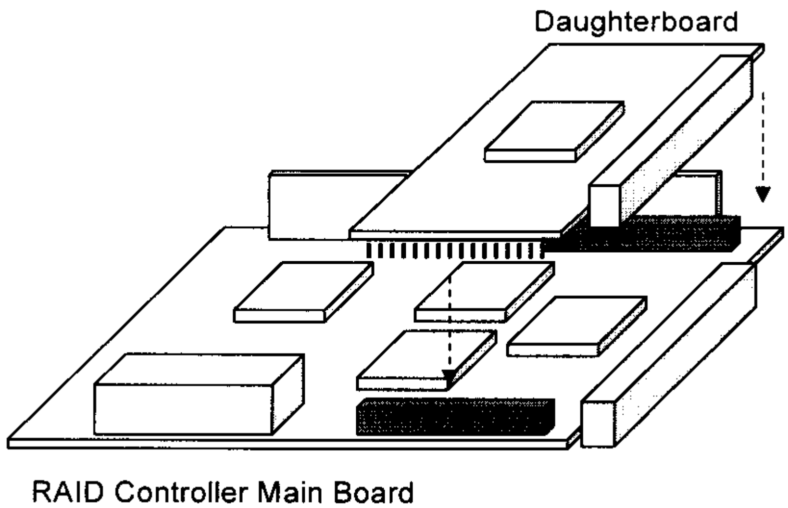

4.4 Power Connection

The power input and connection of the IFT-3102U controller is exactly the same as those for hard disk drives. The power connection is shown below.

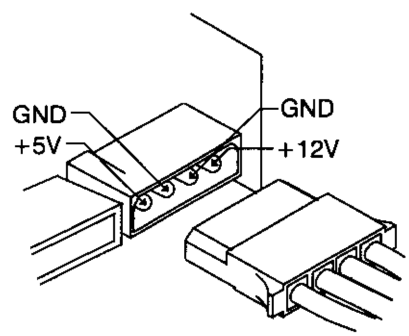

4.5 Daughterboards and Backplane Boards

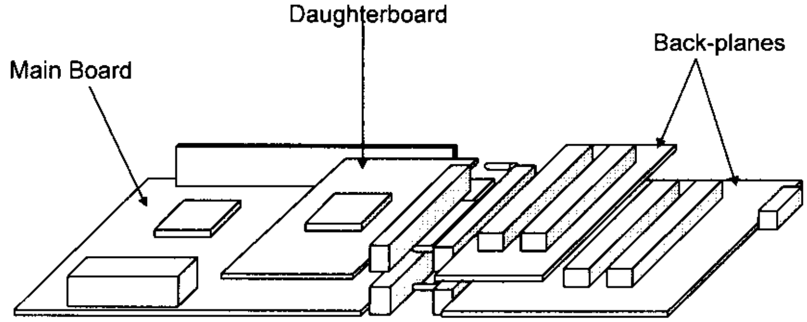

The connector on the rear side of the backplane board, which is used for connecting the 160-pin docking connector of the mainboard, also serves as the rear panel of the controller.

The backplane board is equipped with:

- Fault-bus, I^2C port, redundant controller port, RS-232C port and power connector
- SCSI ports (refer to the table below)

The table below lists the types of channels supported by the backplane boards and the mainboard or daughterboard that corresponds to the backplane boards.

Backplane Boards	Ultra2 SCSI	Differential	Fibre	Connects to...
IFT-9158B2A	2	-	-	IFT-3101U2G mainboard
IFT-9158B2B	2	-	-	IFT-9152 daughterboard
IFT-9158BD	-	2	-	IFT-9152D daughterboard
IFT-9158B2F	-	-	2	IFT-9152F daughterboard

* Channel 0 on each backplane board can be configured as either a SCSI port or Synchronized Cache port for controller redundancy.

4.5.1 Selecting an Appropriate Backplane Board

The table below lists the back-plane boards needed for various main board and daughterboard configurations.

Main Board	Daughterboard	Backplane Board	Ultra2 Wide	Ultra Wide Differential	Fibre
IFT-3101U2G	none	IFT-9158B2A	2	-	-
	IFT-9152	IFT-9158B2A + IFT-9158B2B	4	-	-
	IFT-9152D	IFT-9158B2A + IFT-9158B2D	2	2	-
	IFT-9152F	IFT-9158B2A + IFT-9158B2F	2	-	2

4.5.2 Back-plane Auxiliary Connector Cable

IMPORTANT:
The following applies when a daughterboard and daughterboard back-plane are installed.

The three daughterboard back-planes (IFT-9158B2B, IFT-9158B2D and IFT-9158B2F) include a back-plane auxiliary connector cable. The cable supplements power to the daughterboard back-plane and must be installed for proper operation. To install, connect one end of the cable to J3 on the daughterboard back-plane and the other end to J5 on the main board back-plane (IFT-9158B2A).

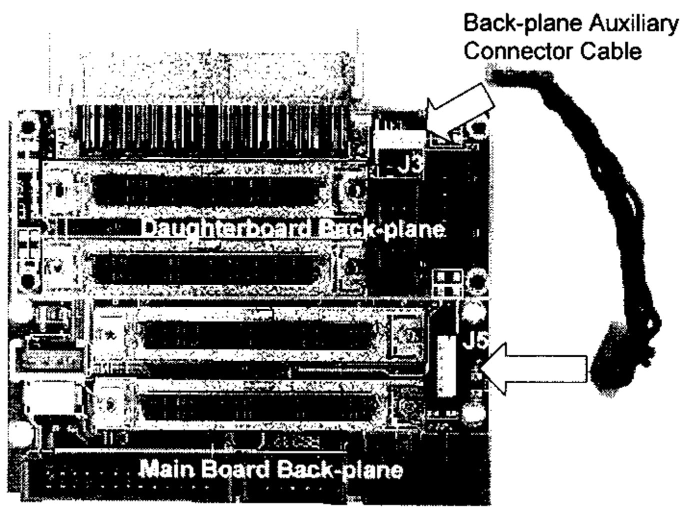

4.6 Serial Port Connection and Setup

The IFT-3101U2G controller can be configured via a PC running a VT-100 terminal emulation program, or a VT-100 compatible terminal. The provided interface cable converts the RS-232 signal of the 10-pin header connector on the controller into a 9-pin D-Sub male connector. The pin layout of the 9-pin D-Sub male connector is similar to that of a PC's serial port and is set as a DTE device. The proper connection of the cable and pin layout is shown below.

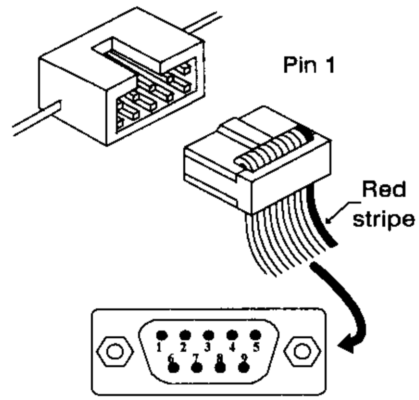

The following are guidelines on using the serial port:

- The serial port's default is set at 9,600 baud, 8 bit, 1 stop bit and no parity. Use the COM1 serial port of the controller.
- In most cases, connecting RD, TD and SG are enough to establish communication with a terminal.
- If you are using a PC as a terminal, any of the VT-100 terminal emulation software will suffice. Microsoft® Windows includes a terminal emulation program as presented with the "Terminal" icon in the Accessories window.

The baud rate can be changed using the front panel. To change the baud rate:

Press **ENT** for two seconds to enter the Main Menu. Press ▼ or ▲ to select "View and Edit Configuration ..", then press **ENT**.

```
View and Edit
Configuration ..
```

Select "Communication Parameters ..", then press **ENT**.

```
Communication
Parameters     ..
```

Select "RS-232 Configuration ..", then press **ENT**.

```
RS-232
Configuration ..
```

Select "COM1 Configuration ..", then press **ENT**.

```
COM1
Configuration ..
```

Select "Baud-rate 9600 ..", then press **ENT**.

```
Baud-rate  9600
               ..
```

Press ▼ or ▲ to select the baud rate, then press **ENT** for 2 seconds to confirm the selected baud rate.

```
Baud-rate 9600
Change to 38400?
```

* The available baud rates are: 2400, 4800, 9600, 19200 and 38400.

4.6 *Basic Operational Setup*

An example of the operational setup is shown below:

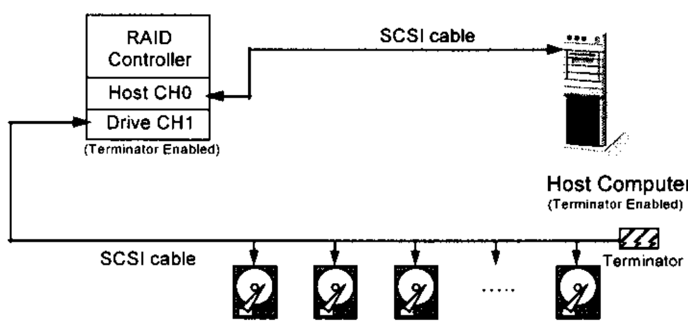

- The SCSI cable must be shorter than 12 meters using high-quality SCSI cable and terminators.
- Channel 0 is connected to the Host system.
- Drives are connected to drive channels 1 and/or 2.

- SCSI nodes on the same channel have their own unique ID number.
- Both the Host and drive SCSI cables are properly terminated.
- The terminator of all the hard drives must be disabled.
- The power supply is attached.
- All operation parameters are properly set.

To connect the components:

1. Make sure power is off or the power connector is disconnected.
2. Connect channel 0 of the controller to the Host system's SCSI port using a suitable SCSI cable.

> NOTE:
> *Channel 0 is the default Host interface using ID number '0'. Any of the channels can also be set as the Host interface. More than one channel can be set as the Host interface when operating with redundant Host or multiple Host systems.*

3. Make sure the host side of the SCSI cable is properly terminated.
4. Assign a unique SCSI ID for every hard disks that are to be connected on the same SCSI cable; between ID numbers '0' and '6' and '8' and '15'. The default ID of the controller's channel 0 is '7'.
5. Connect the other end of the drive SCSI cable to one of the remaining channel on the controller.
6. Connect the connectors located at the middle of the drive SCSI cable to the hard disk(s).
7. Terminate the SCSI cable by installing an external terminator on the last connector. Terminators on all the hard drives must be removed so that removing a hard drive will not affect cable termination.

4.10 In-band SCSI

4.10.1 What is it and why do you need it?

These days more and more external devices require communication with the host computer for device monitoring and administration. This is usually done through RS-232C ports.

Infortrend now offers an alternative means of communication for its RAID controllers—In-band SCSI. The traditional way for SCSI controllers to communicate with the host computer has been via software (such as the GUI RAID Manager) using an RS-232C connection. With In-band SCSI, integrators have more flexibility. They may use RS-232C or the existing SCSI cable instead.

How does it use the SCSI cable? In-band SCSI technology translates the original commands into standard SCSI commands. These SCSI commands are then sent to and received from the SCSI raid controller. The GUI RAID Manager can administrate the RAID controller just as it could before via RS-232C. (Note: It is assumed that users of In-band SCSI possess the following: a third-party SCSI adapter and a channel on their Infortrend RAID controller that can be designated as a host channel.) Both of these are required for In-band SCSI communication between the host and the RAID controller.)

4.10.2 How Do You Configure the GUI RAID Manager to Use In-band SCSI?

4.10.2.1 RAID Controller Adjustments

Don't disconnect the RS-232C cable yet! It is required for another 10 minutes or so. Some adjustments must be made to the RAID controller and to the host computer's SNMP settings before the two can communicate using SCSI commands. (Note: The SNMP settings must be changed prior to installation of the GUI RAID Manager. See *SNMP Settings* below for a detailed explanation.) The RAID controller settings can be changed using the Front Panel. (The front panel may be different in appearance from the one shown in this example.)

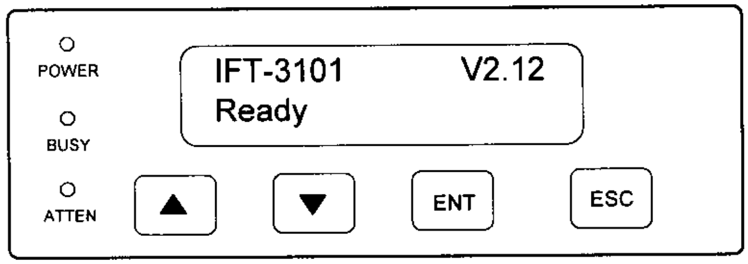

From the Main Menu, press ▼ or ▲ to select "View and Edit Configuration Parameters."

```
View and Edit
Config Parms
```

Press <Enter>; and then use the ▼ or ▲ to select "Host-side SCSI Parameters." Then press <Enter>.

```
Host-side SCSI
Parameters  ..
```

You will need to make adjustments in the following four submenu's: Peripheral Device Type, Peripheral Device Qualifier, Device Support for Removable Media, and LUN Application. Different host operating systems require different adjustments. Look at the table below to find the proper settings for your host operating system.

Peripheral Device Type Parameters Reference for Various Operating Systems:

Operating System	Peripheral Device Type	Peripheral Device Qualifier	Device Support for Removable Media	LUN Applicability
Windows NT® 4.0	0x1f	connected	disabled	All Undefined LUNs
Windows NT® 5.0	0x03	connected	enabled	All Undefined LUNs
NetWare® 4.x	0x03	connected	disabled	All Undefined LUNs
SCO Unix Ware 5.0x	0x7f	connected	either is okay	All Undefined LUNs
SCO OpenServer 2.1x	0x03	connected	either is okay	All Undefined LUNs
Solaris™ 2.5.x/2.6 (x86 and SPARC)	0x7f	connected	either is okay	All Undefined LUNs
Linux	0x1f	connected	enabled	All Undefined LUNs

Peripheral Device Type Settings:

Device Type	Setting
No Device Present	7f
Direct-access Device	0
Sequential-access Device	1
CD-ROM Device	5
Scanner Device	6
MO Device	7
Unknown Device	1f
Processor Type	3

Example: Settings for Windows NT® 4.0

The settings for Windows NT® 4.0 are provided here as an example. For the settings for other operating systems, please refer to the table above, *Peripheral Device Type Parameters Reference for Various Operating Systems.*

On the front panel, use ▼ or ▲ to select "Peripheral Device Type Parameters"; and then press <Enter>.

```
Periph        Dev
Type Parameters
```

(For this example, we assume that there are currently no peripheral devices.)

```
Device Type -
No Device (0x7f)
```

Press ▼ or ▲ to choose "Unknown Device - 1f".

```
Set Device Type?
Unknown (0x1f)
```

Press <Enter> to confirm the selection. Now that we have changed the Peripheral Device Type, let us set the Peripheral Device Qualifier. Press <Esc> to return to the sub-menu's mentioned above. Use the arrow keys to scroll down to Device Qualifier., press ▼ or ▲ to select "Device Qualifier Connected. "

The default setting is "Connected." If your Front Panel reads "Disconnected," press <ENT> and you will be prompted to change to "Connected". If your Device Qualifier setting reads "Connected," press <Esc> to return to the Host-side SCSI submenu's.

```
Device Qualifier
Connected
```

Use the ▼ or ▲ to select Support for Removable Media. The default setting is "Disabled." If the LCD reads "Enabled," press <Enter> and you will be prompted to accept a change. If the screen reads "Disabled," press <Esc> to return to the Host-side SCSI submenu's.

```
SupportRemovable
Media - Disabled
```

Press ▼ or ▲ to select "LUN Application"; and then press <Enter>. The default setting is "All Undefine LUN."

```
LUN Application-
All Undefine LUN
```

Press <Enter> and use ▼ or ▲ to select "Undefine LUN-0's."

```
Applies to      ?
Undefine LUN-0's
```

Press <Enter> to accept. The screen should display the following message.

```
LUN Application-
Undefine LUN-0's
```

The RAID controller adjustments necessary to use In-band SCSI have been completed. For locally accessing the host computer (see section 4.10.3.1, *Local Connection – SNMP Not Required*), all steps have been completed. For remotely accessing the host computer, further adjustments must be made (see section 4.10.3.2, *Remote Connection – SNMP Required*).

4.10.3 Using In-band SCSI in GUI RAID Manager

4.10.3.1 Local Connection — SNMP Not Required

If you are using the GUI RAID Manager on the host computer that is using In-band SCSI -- ie., 'local access' -- SNMP service is not required. You may now connect by going to FILE>CONNECT and selecting the port address.

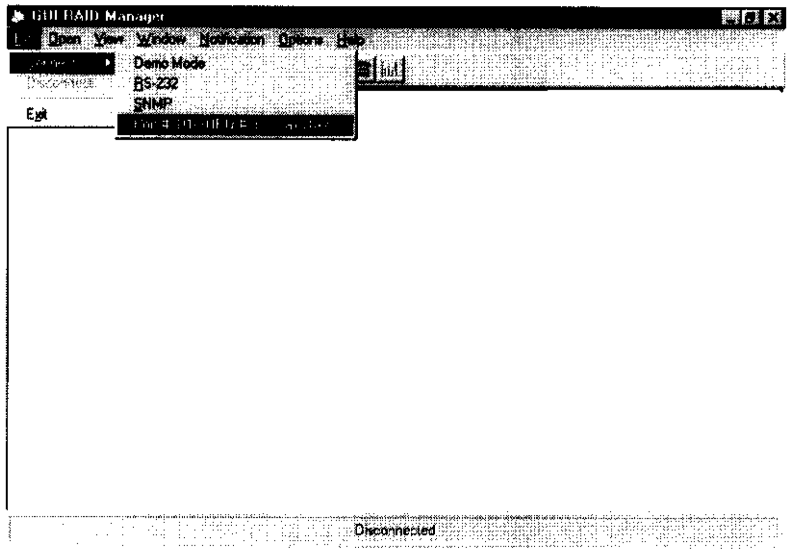

4.10.3.2 Remote Connection — SNMP Required

SNMP Service is required to administrate a RAID controller installed in a remote computer. This passage describes how to establish a connection from the GUI RAID Manager to a RAID controller that is connected to a remote host via In-band SCSI. For information regarding other operations and of the GUI RAID Manager, please refer to the *GUI RAID Manager User Guide.*

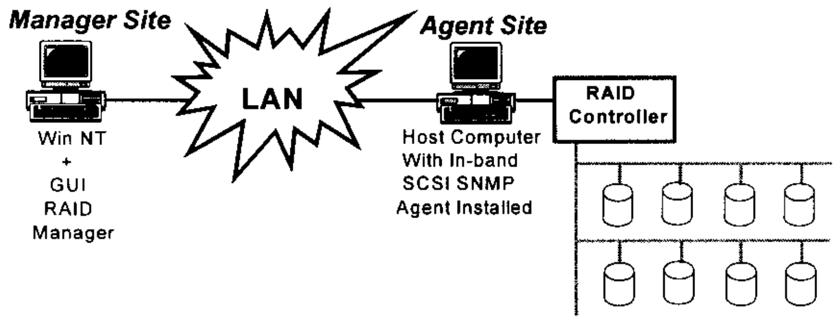

In the figure above, the "Agent Site" is a host computer connected to a RAID controller via a SCSI cable. The "Manager Site" is a Windows NT® system with the GUI RAID Manager. The Agent Site could be running an operating system other than Windows NT®. Infortrend provides the In-band SCSI SNMP agents for the following operating systems:

◊ Windows NT®
◊ NetWare
◊ SCO OpenServer
◊ SCO UnixWare
◊ Sun Solaris™

The Manager Site should be a Windows NT® Workstation or Server with SNMP service and the GUI RAID Manager installed.

Basic Procedures to Establish the Connection

The following criteria must be met for the Agent Site and Manager Site:

Checklist for Agent Site

1. The host computer is connected to the RAID controller via the host SCSI cable (the cable which is used to transfer data between the host computer and the RAID – there's no need for an extra SCSI cable.)

2. The host computer's operating system has SNMP service installed.

3. The host computer has the In-band SCSI SNMP agent installed for the corresponding operating system. (The example described herein is Windows NT®.)

4. The host computer is up and running.

IMPORTANT:
This following applies to Windows NT® 4.0 with Service Pack 3: if the "SNMP Service" is installed after the Service Pack 3 has been installed, Service Pack 3 must be re-installed in order for the SNMP service to work properly.

Checklist for Manager Site

1. The system is running Windows NT® (Workstation or Server) and has SNMP Service installed.

2. The Infortrend GUI RAID Manager was installed with the "SNMP Manager Site" option selected.

3. The GUI RAID Manager is running.

IMPORTANT:
This following applies to Windows NT® 4.0 with Service Pack 3: if the "SNMP Service" is installed after the Service Pack 3 has been installed, Service Pack 3 must be re-installed in order for the SNMP service to work properly.

Example Settings for Agent Site Using Windows NT

1. Install SNMP Service in Windows NT®. Look for the "Network" icon in the Control Panel. Double click on the "Network" icon to open it.

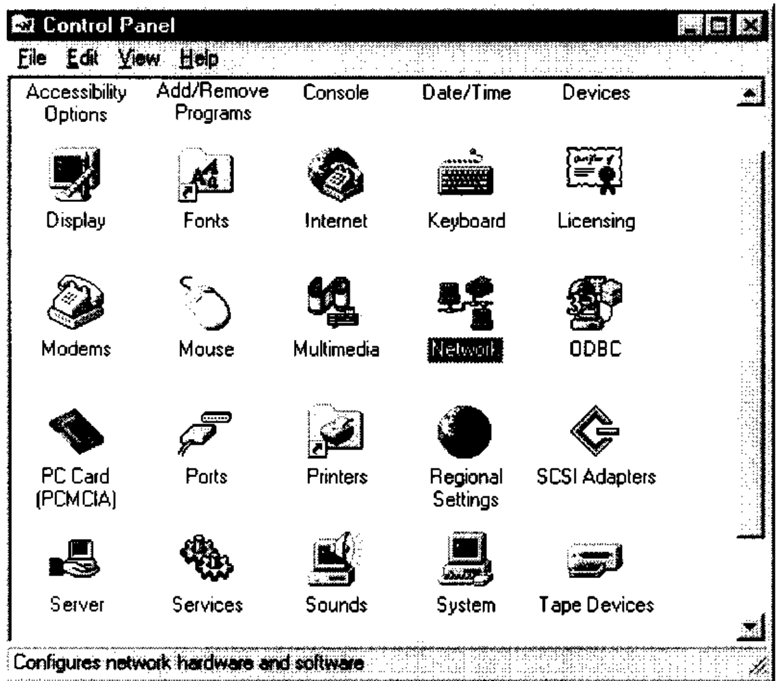

The Properties window appears. Choose the "Service" tab. If the SNMP Service is already installed, please go ahead to step called "Install the SNMP Agent and GUI RAID Manager." If the SNMP Service is not yet installed, click on "Add" and choose "SNMP Service" to install.

IMPORTANT:

This following applies to Windows NT® 4.0 with Service Pack 3: if the "SNMP Service" is installed after the Service Pack 3 has been installed, Service Pack 3 must be re-installed in order for the SNMP service to work properly.

2. Install the SNMP Agents and GUI RAID Manager. The GUI RAID Manager can install the In-band SCSI SNMP Agent during installation. During GUI RAID Manager installation, be sure to select both the "SNMP Manager Site" and "SNMP Agent Site" options. Complete the installation by following the on-screen instructions.

Example Settings for Manager Site

1. Install the SNMP manager and GUI RAID Manager. During the installation of GUI RAID Manager, click to select the option "SNMP Manager Site." Complete the installation by following the on-screen instructions.

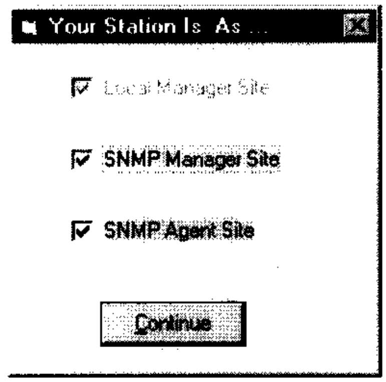

2. Run the GUI RAID Manager. Start the connection by choosing FILE>CONNECT>SNMP.

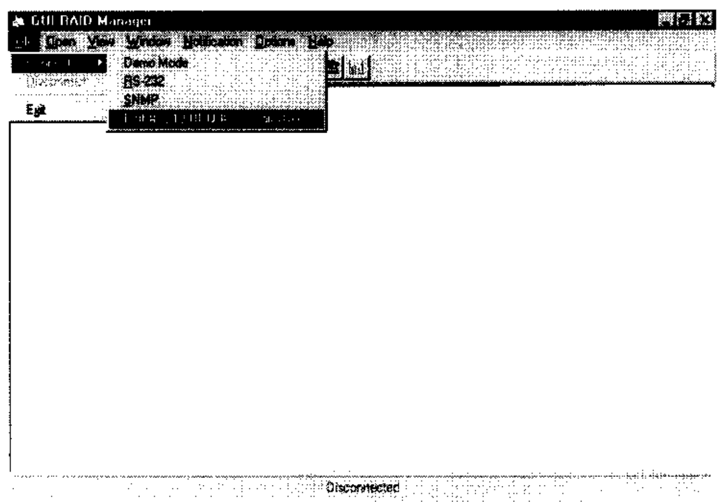

3. The "Setting of SNMP Interface" window appears. Enter the IP address and community name of the Agent Site. For "Agent-site

Communication Type," choose "SCSI Bus Interface." The "Controller Index" refers to the sequence of the RAID controller which is going to be administrated. If only one RAID controller is installed in the agent site computer, choose "0". If there is more than one RAID controller installed in the agent site computer, choose "1" to administrate the second RAID controller. Choose "2" to administrate the third RAID controller, etc.

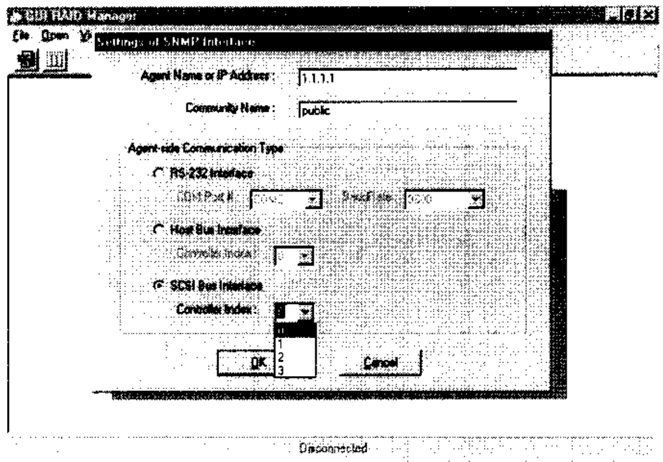

4. After the connection is established, all of the operations in the GUI RAID Manager are the same as before (please refer to the *GUI RAID Manager User Guide* for complete details on its operation.)