Infortrend®

Infortrend EonStor® Storage System

Configuration Planning Guide

Jan. 09, rev. 1.1b

Firmware Revision: 3.64.x and above

Contact Information

Asia Pacific
(International Headquarters)
Infortrend Technology, Inc.
8F, No. 102 Chung-Shan Rd., Sec. 3
Chung-Ho City, Taipei Hsien, Taiwan
Tel: +886-2-2226-0126
Fax: +886-2-2226-0020
sales.ap@infortrend.com
support.ap@infortrend.com
http://esupport.infortrend.com.tw
http://www.infortrend.com.tw

Americas
Infortrend Corporation
2200 Zanker Road, Unit D,
San Jose, CA. 95131
USA
Tel: +1-408-988-5088
Fax: +1-408-988-6288
sales.us@infortrend.com
http://esupport.infortrend.com
http://www.infortrend.com

China
Infortrend Technology, Limited
Room 1210, West Wing, Tower One, Junefield Plaza, No. 6 Xuanwumen Street, Xuanwu District, Beijing, China
Post code: 100052
Tel: +86-10-6310-6168
Fax: +86-10-6310-6188
sales.cn@infortrend.com
support.cn@infortrend.com
http://esupport.infortrend.com.tw
http://www.infortrend.com.cn

Europe (EMEA)
Infortrend Europe Limited
1 Cherrywood, Stag Oak Lane
Chineham Business Park
Basingstoke, Hampshire
RG24 8WF, UK
Tel: +44-1256-707-700
Fax: +44-1256-707-889
sales.eu@infortrend.com
support.eu@infortrend.com
http://esupport.infortrend-europe.com/
http://www.infortrend.com

Japan
Infortrend Japan, Inc.
6F, Okayasu Bldg.,
1-7-14 Shibaura Minato-ku,
Tokyo, 105-0023 Japan
Tel: +81-3-5730-6551
Fax: +81-3-5730-6552
sales.jp@infortrend.com
support.jp@infortrend.com
http://esupport.infortrend.com.tw
http://www.infortrend.co.jp

Germany
Infortrend Deutschland GmbH
Werner-Eckert-Str.8
81829 Munich Germany
Tel: +49 (0)89 45 15 18 7 - 0
Fax: +49 (0)89 45 15 18 7 - 65
sales.de@infortrend.com
support.eu@infortrend.com
http://www.infortrend.com/germany

Copyright 2008

This Edition First Published 2008

Disclaimer

Infortrend Technology makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Furthermore, Infortrend Technology reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation to notify any person

of such revisions or changes. **Product specifications are also subject to change without notice.**

Trademarks

Infortrend, Infortrend logo, EonStor and SANWatch are all registered trademarks of Infortrend Technology, Inc. Other names prefixed with "IFT" and "ES" are trademarks of Infortrend Technology, Inc.

All other names, brands, products or services are trademarks or registered trademarks of their respective owners.

Table of Contents

Organization of this Guide

Chapter 1 Provides background information about host interface features and describes the major storage configuration types.

Chapter 2 Describes RAID levels and logical drives (also termed as RAID groups or arrays) and how they provide fault tolerance and combined performance.

Chapter 3 Gives you description of the basic steps for storage configuration and LUN mapping.

Appendix 1. Tunable firmware parameters and firmware limitations.

Appendix 2. Using hot spares.

Revision History

- Rev. 1.0: Initial release

- Rev. 1.1: - Removed JBOD from the RAID level introduction. NRAID provides similar functionality.

  - Added definitions for the Active and Passive data paths in a scenario involving redundant controllers, redundant paths, and the EonPath multi-pathing driver.

- Rev. 1.1a: Dynamic switch of LD ownership in the event of external link failure is now supported by firmware release 3.64h.

- Rev. 1.1b: Replaced the sample drawings for the Dynamic LD Assignment feature.

Related Documentations

- Firmware Operation Manual

- SANWatch User's Manual
- EonPath User's Manual
- Embedded RAIDWatch User's Manual
- Installation and Hardware Reference Manual
- Quick Installation Guide
- Rackmount Rail Installation Guide (for later models, rackmounting details are described in its Quick Installation Guide)
- System Troubleshooting Guide
- LCD Keypad Navigation Map

These documents can be found in the product utility CD included with your system package and are continuously updated according to the progress of technologies and specification changes.

Chapter

1

Host Interface and Storage Configuration Basics

1-1. Host Interface Types:

The EonStor series storage systems are equipped with prevalent types of host link interfaces including:

1. Fibre Channel,
2. Serial Attached SCSI (SAS),
3. Internet SCSI (iSCSI).

Parallel SCSI is gradually being replaced by SAS and is not included in the following discussion.

1-2. Storage Configuration Types:

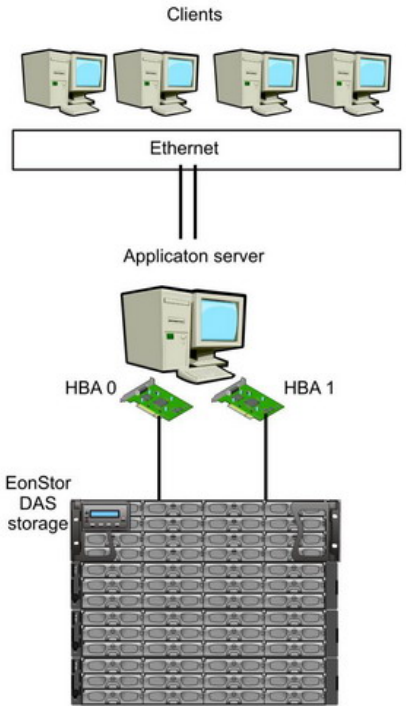

DAS

Direct-attached Storage. Refers to storage that is directly attached to application servers using the SAS data links. DAS configuration is characterized by its limited cabling distance and the inability to share storage for multiple servers.

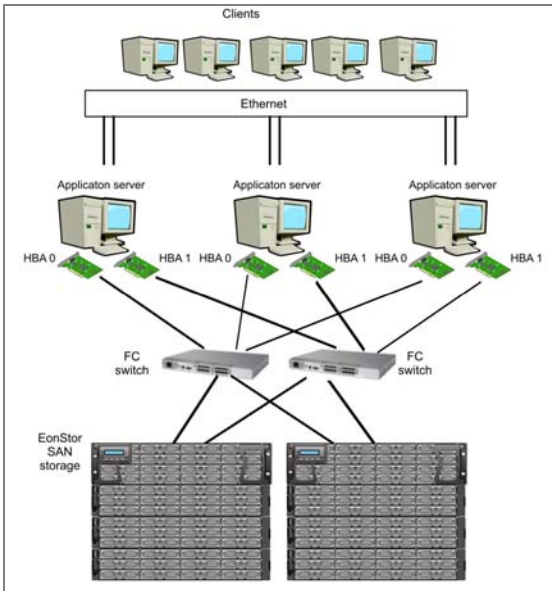

SAN

Storage Area Network. Refers to configurations that include storage systems connected through Fibre Channel data links. SAN configurations often include interconnect hardware such as fabric switches. Fibre Channel SAN can span across an entire enterprise or further and enable the connections to almost limitless number of application servers in a storage network.

IP SAN

Often considered as a cost-down alternative to Fibre Channel SAN. Refers to the configurations with iSCSI storage that attaches to an existing Ethernet network. iSCSI storage reduces the implementation cost by exchanging SCSI commands over the TCP/IP infrastructure.

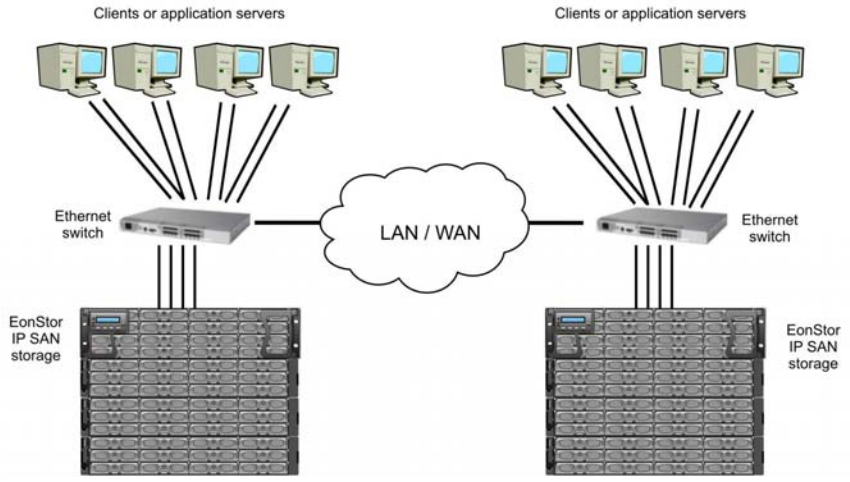

1-3. Host Link Components:

Storage-side components:

Host ports:

1. SAS links for DAS:

There are two different kinds of SAS ports: SFF-8088 and SFF-8470; both are multi-lane wide ports.

1-1. DAS Host Port Example: **EonStor B12S**

1-2. DAS Host Port Example: **EonStor S16S**

Host Link Cables:

SAS cable with SFF-8088 connector	SAS cable with SFF-8470 connector
	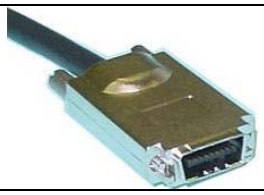

One 120cm host link cable (with SFF-8088 or SFF-8470 connectors) is shipped with the EonStor DAS series. A 50cm version is also available. Other SAS link cables are separately purchased.

2. FC links for SAN:

SAN Host Port Example: **EonStor B12F**

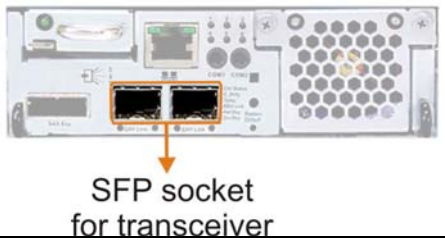

Fibre Channel host ports are SFP sockets that receive separately purchased Fibre Channel transceivers. The transceiver converts electrical signals into optical signals and transmits data over fiber optical links.

Fibre Channel optical transceiver:	Fiber optical cable (LC-to-LC):
	

3. Ethernet links for IP SAN:

IP SAN Host Port Example: **EonStor S16E**	
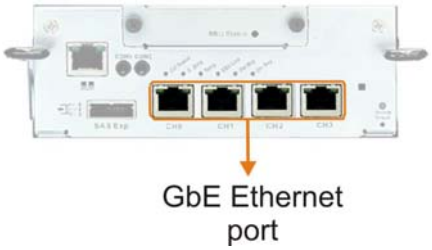	

Host Links cables:

The Ethernet cables are user-supplied. Use Cat5e or better performance cables for cabling iSCSI storage to an IP SAN.

1-4. Cabling Host Ports and Theories behind Topologies:

Shown below are the basics about cabling systems with the single- and redundant-controller configurations.

There are two reasons for having a redundant controller system:

1. Double the system performance by combining the computing power of the partner controllers.
2. To continue service in the event of a single controller failure.

Cabling and configuring a storage system powered by redundant controllers can be tricky because attentions must be paid to prepare fault-tolerant paths as a precaution for device failure. For a mission-critical application, down time can be very costly.

Shown below are sample topologies that help you design your own configurations. There are more connection samples in the EonStor series' Hardware manual. The key elements in each topology are briefly described.

<table>
<tr><th colspan="4">Legends</th></tr>
<tr><td>HBA:</td><td>Host bus adapter</td><td>CH0:</td><td>Host channel 0</td></tr>
<tr><td>LD:</td><td>Logical drive; logical group of 6, 8, or other number of disk drives.</td><td>CH1:</td><td>Host channel 1</td></tr>
<tr><td>AID:</td><td>e.g., A112; a host ID managed by controller A</td><td>RCC:</td><td>The communications paths between controllers</td></tr>
<tr><td>BID:</td><td>e.g., B113; a host ID managed by controller B</td><td>FC switch</td><td>Fibre Channel switch that provides intermediate connectivity to form a storage area network. FC switches also provide access control such as zoning.</td></tr>
<tr><td>LUN Mapping:</td><td>Host LUN mapping is presented by the encircled numbers either placed by the LD or on the data paths.</td><td colspan="2" rowspan="2">NOTE:
1. The samples below are made with the Fibre Channel connectivity.
2. The default host IDs can vary on the EonStor models:</td></tr>
<tr><td>Controller</td><td>The RAID controllers within storage system. Controllers are identified as controller A or controller B.</td></tr>
<tr><td colspan="2"></td><td>FC</td><td>112 and 113</td></tr>
<tr><td colspan="2"></td><td>SAS</td><td>0, 1 (single controller)
6, 7 (dual-controller)</td></tr>
<tr><td colspan="2"></td><td>iSCSI</td><td>0, 1 (single controller)
6, 7 (dual-controller)</td></tr>
</table>

1-4-1. Calculating an Approximate Storage Performance:

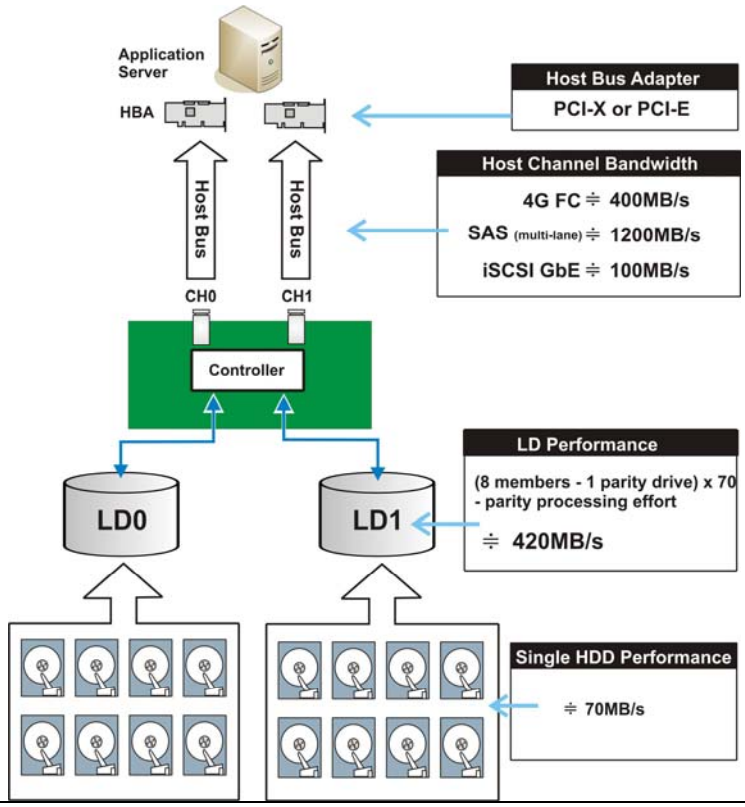

An optimal system performance depends on a careful planning with the concerns for various component factors.

HDD Speed:

Today's HDD can deliver a throughput speed between 70MB/s and 100MB/s, and 150 IOPS. You can use the performance data by disk vendor as a basis for speculating an optimal deployment.

LD:

Logical drives provide combined performance by grouping multiple hard drives. For a logical drive composed of RAID3, 5, and 6, parity or spare drives do not contribute to RAID performance.

LD Size (Stripe width):

Combine a reasonable number of hard disks into a logical drive. A logical drive consisting of too many members will take a very long time to rebuild. A combination of 6 or 8 members can be the optimal. Of course, RAID0 provides the best performance but with no fault tolerance.

LD Performance:

With the above measures, we can come up with a rough LD performance by subtracting 20% off the combined performance because a certain amount of system resource has to be consumed for generating and distributing parity data.

Taking a RAID5 LD of 8 members as an example,

(8 - 1) x 70MB/s – parity handling efforts = 420MB/s

The LD performance can roughly fill a 4Gbps Fibre host channel.

Multi-pathing Driver:

With the EonPath multi-pathing driver, traffic on multiple host links can be balanced by presenting a logical drive on them.

1-4-2. System Overall Performance:

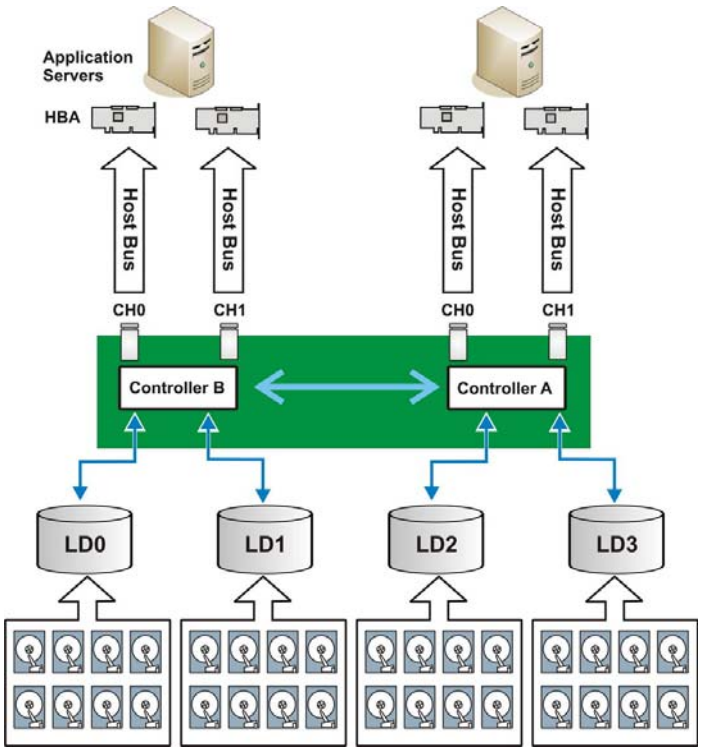

You can fully utilize the powerful engine in the EonStor series through the configuration means.

A combination of 32 HDD in a RAID and a JBOD can theoretically make a best use of the power of a 16-bay redundant controller system:

↓	There are 4 LDs: Each of 8 members; 2 in the RAID enclosure and 2 in the JBOD.
↓	Each LD delivers a 420MB/s performance (see previous description.
↓	Each RAID controller manages 2 LDs (LD assignment).
↓	There are 4 host channels (2 per controller).
↓	4 LDs deliver a total of 1600MB/s performance, which is slightly lower than the approximate of system capability.

	For the fact that your application servers may not always generate I/Os that fully stress the arrays, more disk drives can be attached. In a storage configuration, logical drives, host LUN mapping, and other configurations can be re-arranged, if the nature of host applications and data has been changed throughout the time of use.

	Other Considerations: ▪ For high-speed I/O channels, use host bus adaptors that are at least with a PCI-X x8 lane. Using outdated HBAs on a narrow bus can hinder the best host-storage performance. ▪ For a higher level of fault tolerance, say, if you connect 4 host links from redundant RAID controllers, use dual-ported HBAs for making the connections instead of linking all 4 ports to a quad-ported HBA. ▪ Perform throughput testing on the whole deployment before starting your applications. ▪ Understand and fine-tune your I/Os. Create logical drives to your needs for performance, fault tolerance, or for both. Some minor details, such as HBA BIOS settings and queue depth configurations, can be important but are easily ignored.

1-4-3. Single-controller storage:

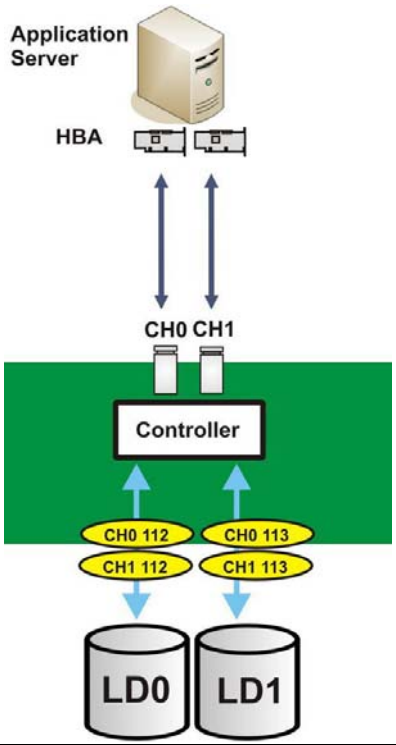

Preparing a single-controller storage system is comparatively simple.

Elements in this drawing are:

LD: Logical drives are configured by grouping physical drives.

IDs: Infortrend firmware comes with 1 host ID on each channel. Other IDs are manually created.

ID Mapping: Logical drives are mapped to IDs on both host channels. Mapping a logical drive to IDs on different channels provides access from 2 data paths.

1-4-4. Redundant-controller storage in a switched fabric:

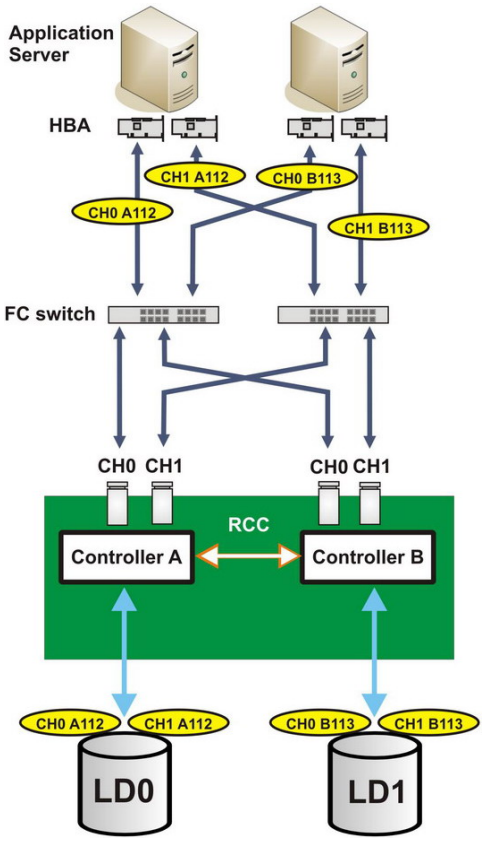

Preparing a redundant-controller system requires both AID and BID. Resource distribution is also determined by Logical Drive Assignment. If a logical drive is assigned to controller A, then controller A manages the I/Os to that logical drive.

Elements in this drawing are:

LD: Logical drives are configured by grouping physical drives.

LD assignment: Each logical drive is either assigned to controller A or to controller B.

ID Mapping: Logical drives are mapped to IDs on all host channels to leverage all host port bandwidth.

- Infortrend firmware comes with 1 host ID on each channel. You need to manually create more IDs.
- More IDs can be associated with each LD to provide more active paths.

Data Paths: Data paths are routed from different RAID controllers, between FC switches, and to different servers. This way, a server can still access data when a cabling failure occurs.

Multi-pathing: The EonPath software is necessary on the servers.

NOTE:

1. Multiple IDs on a Fibre Channel host channel is not allowed if they are configured into the “point-to-point” mode.

   The maximum number of LUN is:

   Point-to-point: 4 (host channels) x 1 (IDs per channel) x 32 (LUNs per ID) = 128

   FC-AL: 4 (host channels) x 8 (IDs per channel) x 32 (LUNs per ID) = 1024

   You can seldom use the maximum number, and having too many LUN can cause a performance drag.

2. It is recommended to set your storage and switch ports to the loop mode (FC-AL). In some circumstances with cabling/controller failures, a server may not regain the access to storage through a switch port configured in the fabric mode (point-to-point).

1-4-5. Redundant-controller storage for dedicated performance:

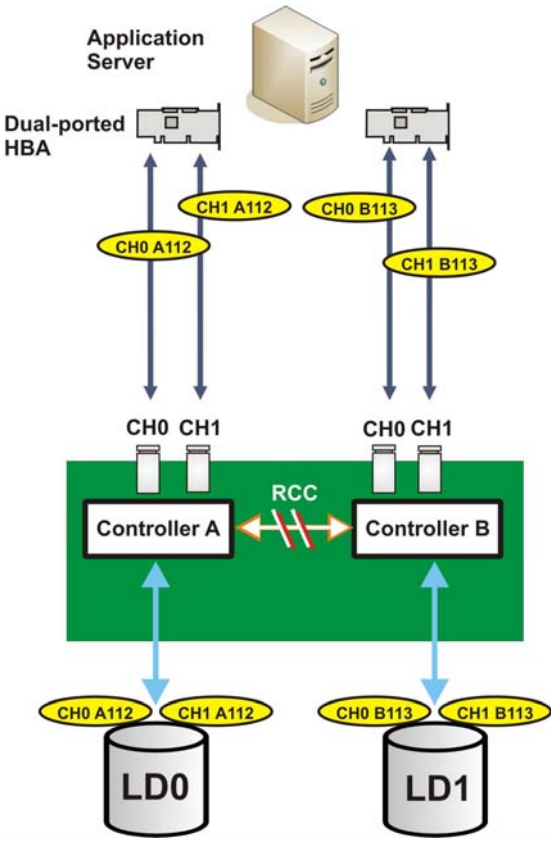

Some storage applications may not require high level of fault tolerance, e.g., AV post-production editing.

Elements in this drawing are:

LD: Logical drives are configured by grouping physical drives.

LD assignment: Each logical drive is either assigned to controller A or to controller B.

ID Mapping: Logical drives are mapped to IDs on all host channels to leverage all host port bandwidth.

- Infortrend firmware comes with 1 host ID on each channel. You need to manually create more IDs.
- More IDs can be associated with each LD to provide more active paths.

Data Paths: Data paths are directly routed to an application server. A special firmware is required to disable the RCC communications between controllers to conserve the most for I/O service.

Multi-pathing: The EonPath software is necessary on the servers.

NOTE:

The sample topologies in this document do not cover the cases of using the onboard hub (onboard FC bypass) such as those applied in the ASIC266 models. The onboard hub turns host ports of partner RAID controllers into a host loop.

1-4-6. Redundant-controller, high availability, for clustered servers:

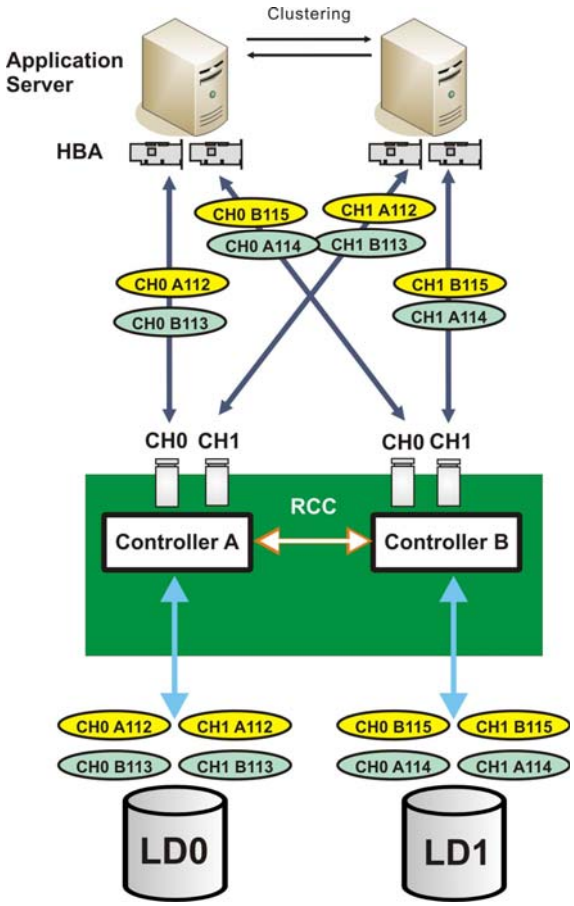

Provides shared storage for high availability clustered servers.

Elements in this drawing are:

LD: Logical drives are configured by grouping physical drives.

LD assignment: Each logical drive is either assigned to controller A or to controller B.

ID Mapping: Logical drives are mapped to IDs on all host channels to leverage all host port bandwidth.

The IDs in green circles are stand-by IDs. The stand-bys provide alternate access in the event when the controller having the original ownership fails.

- Infortrend firmware comes with 1 host ID on each channel. You need to manually create more IDs.
- More IDs can be associated with each LD to provide more active paths.

Data Paths: Data paths are directly routed to clustered servers so that both servers can access the LD.

Multi-pathing: The EonPath software is necessary on the servers.

1-4-7. One controller failed in a redundant-controller storage:

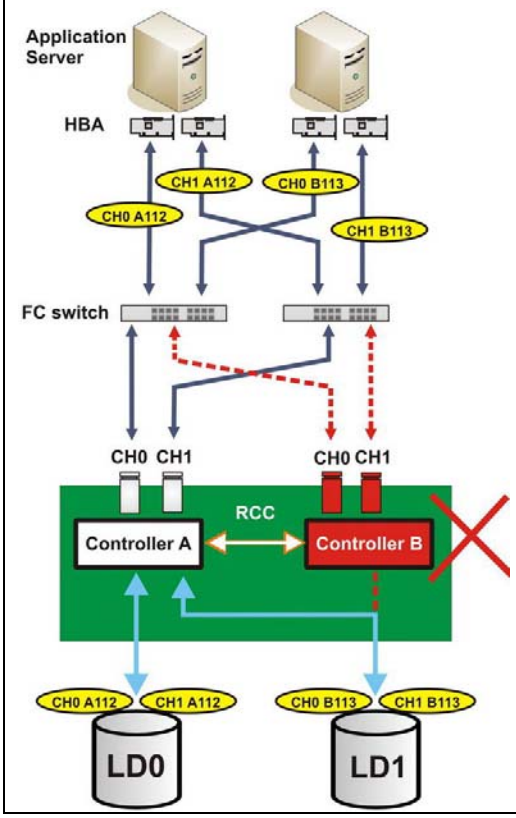

Elements in this drawing are:

Controller failure: Controller B fails. All AID and BID are taken over by controller A, the surviving controller.

Disk Access: LD1 is accessed through the alternate data paths on the backplane.

The failover process takes only a few seconds and is transparent to users.

1-4-8. Cable link failure. Before Dynamic LD Assignment with FW3.64J, a cabling failure can cause a degraded performance in the scenario diagrammed below.

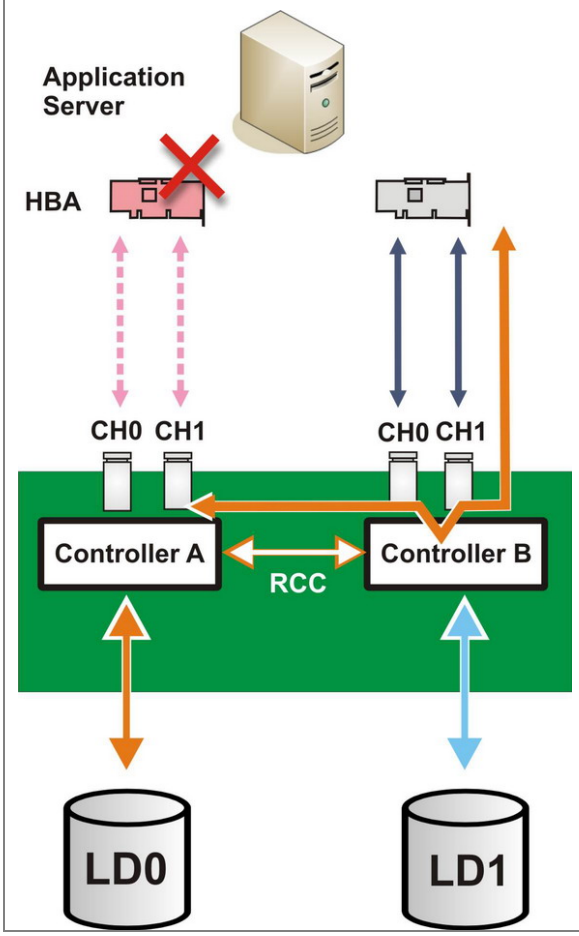

A cabling failure occurs, e.g., an HBA failure. If a data route is disconnected, I/Os will be directed through the RCC links between partner controllers.

Because it is a cabling failure, controller A still holds the ownership of LD0.

Re-directing I/Os through the alternate data paths and RCC links consumes considerable resources. Performance will be compromised although both controllers are still working normally.

1-4-9. Dynamic Switch of LD Ownership in a redundant-controller storage:

Dynamic LD Assignment can dramatically improve system performance in the same cabling failure scenario.

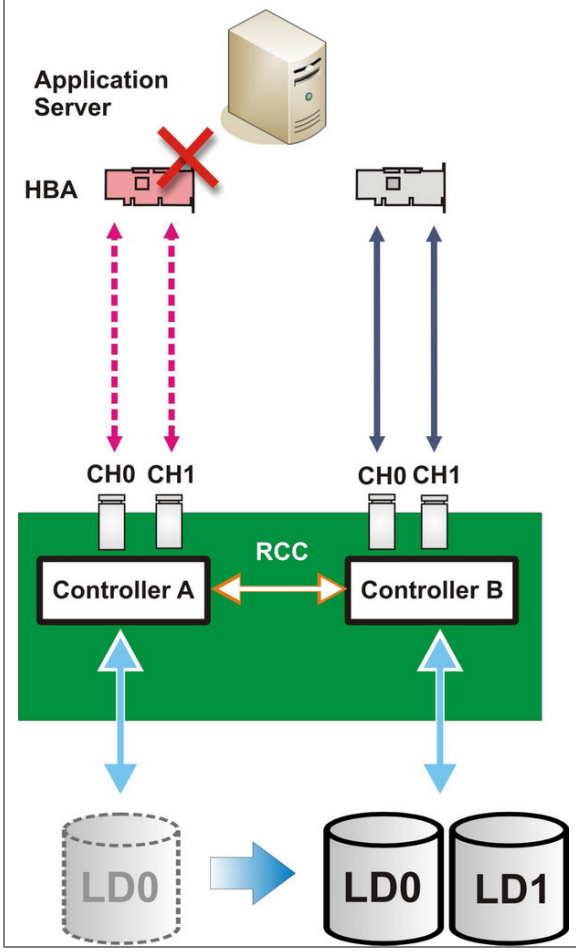

Since firmware revision 3.64J, LD ownership can be temporarily shifted to the partner controller to avoid the overhead of re-directing I/Os through the RCC links. The LD0 ownership is temporarily handed to Controller B.

As the result, the Dynamic Assignment approach can produce a 50% performance gain over the traditional approach of routing through the RCC links,

NOTE: The Dynamic LD Assignment complies with multi-pathing driver capable of TPGS (Target Port Group Service) so that path preference can be restored once the broken host links are restored.

1-4-10. The Active and Passive path mechanism to a redundant-controller storage:

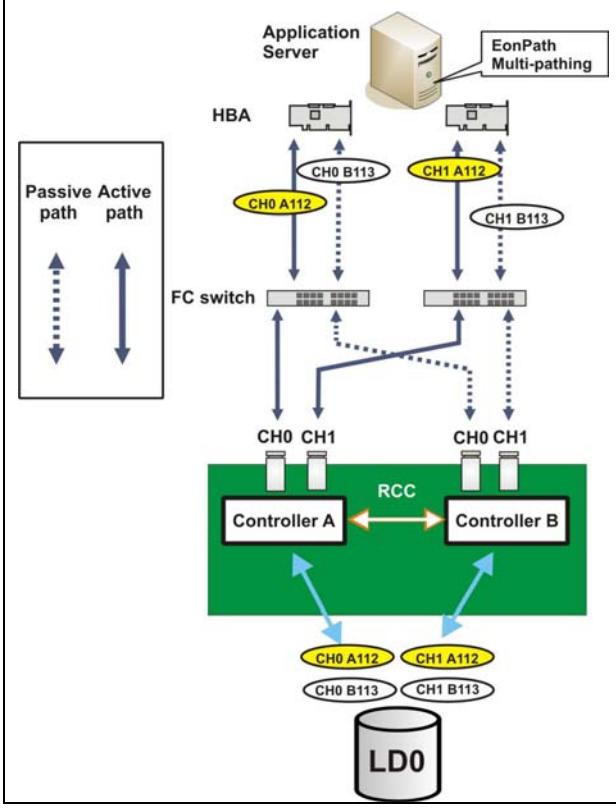

The data path's Active/Passive status is determined by the logical drive ownership. If a logical drive (LD0) is assigned to controller A, the data paths to controller A are considered as the Active or optimal paths for the access to LD0. I/Os will be distributed through the Active paths.

The path status is negotiated between firmware and the EonPath driver on the host side. In the event of Active path or controller failure, I/Os will be directed through the Passive paths.

Chapter

2

RAID Levels

Redundant Arrays of Independent Disks, or RAID, offers the following advantages: availability, capacity, and performance. Choosing the right RAID level and drive failure management can increase capacity and performance, subsequently increasing availability. Infortrend's external RAID controllers and subsystems provide complete RAID functionality and enhanced drive failure management.

A RAID storage delivers the following advantages:

- **Capacity**: Provides disk spanning by weaving multiple disk drives into one single volume.
- **Performance**: 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 as the host channel bandwidth allows.
- **Fault Tolerance**: Provides fault-tolerance by mirroring or distributing parity across member drives.

A Comparison of RAID Levels

RAID Level	Description	Capacity	Data Availability
NRAID	Non-RAID	N	None
RAID0	Disk Striping	N	Less than one single drive
RAID1 (0+1)	Mirroring Plus Striping (if N>2)	N/2	high >RAID5
RAID3	Striping with Parity on a dedicated parity disk	N-1	high ==RAID5
RAID5	Striping with interspersed parity	N-1	high ==RAID5
RAID6	Striping with P+Q (2^{nd} drive failure tolerance by redundantly distributed parity), interspersed parity	N-2	highest >>RAID5

NOTE:

Logical volumes, such as RAID50, can provide a higher level of fault tolerance than RAID5. However, the use of logical volumes is not always necessary. Using logical volumes can create the load on system hardware and may not be the optimal for most applications.

Performance Comparison of RAID Levels

RAID Level	Performance Sequential
NRAID	Drive performance
RAID0	R: Highest W: Highest
RAID1 (0+1)	R: High W: Medium
RAID3	R: High W: Medium
RAID5	R: High W: Medium
RAID6	R: High W: Slightly lower than RAID5

Sample Applications

RAID Level	Performance Sequential
RAID0	RAID0 can deliver the best performance, but please be reminded it provides no protection to your data. RAID0 is ideal for applications needing a temporary data pool for high-speed access.
RAID1 (0+1)	RAID1 is useful as a small group of drives pertaining high availability and fast write access although it is expensive in terms of its usable drive capacity.
RAID3	RAID3 works well with single-task applications featuring large transfers such as video/audio post-production editing, medical imaging, or scientific research requiring a purpose-oriented performance.
RAID5	RAID5 is most widely-used and is ideal for a media, legal, or financial database repository with lower write requests. RAID5 can adapt to multi-task applications with various I/O sizes. A RAID5 with an adequate stripe size is also applicable with large I/O transfers.
RAID6	RAID6 provides a high level of data availability, benefits of RAID5, with the minor trade-off of a slightly lower write performance. RAID6 can mend the defects of using cost-effective SATA drives where magnetic defects can cause problems if another member drive fails at the same time.

RAID Levels in Details

NRAID - Disk Spanning

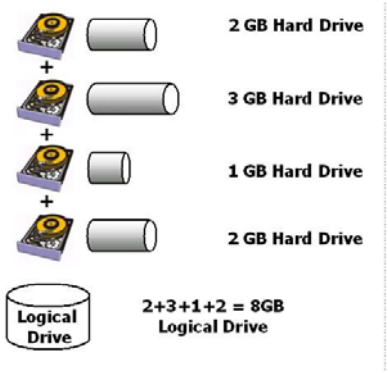

NRAID	
Minimum Disks required	1
Capacity	N
Redundancy	No

NRAID stands for Non-RAID. The capacity of all drives is 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 member drives. NRAID does not provide data redundancy.

Some vendors provide a self-defined RAID level, JBOD, as a way to concatenate disk drives into a volume. NRAID can be made of 1 or multiple disk drives in a way very similar to the use of JBOD.

RAID0 - Disk Striping

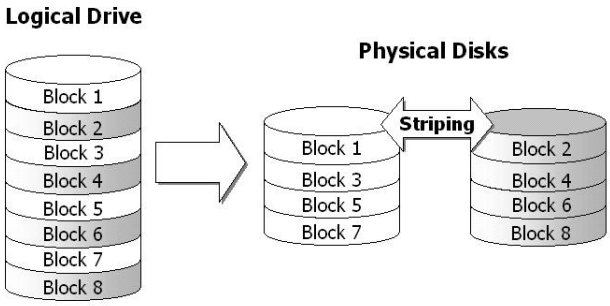

RAID0	
Minimum Disks required	2
Capacity	N
Redundancy	No

RAID0 provides the highest performance but no redundancy. Data in the logical drive is striped (distributed) across physical members.

RAID1 - Disk Mirroring

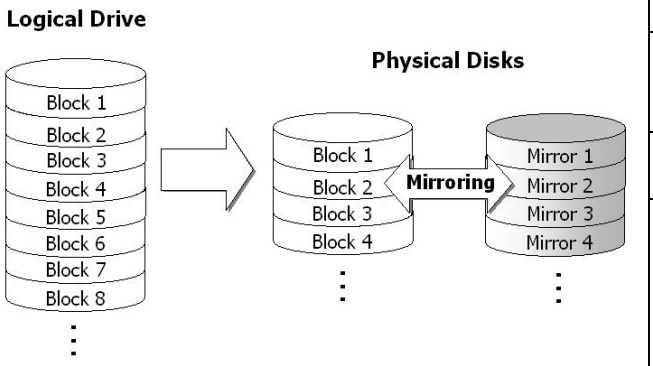

RAID1	
Minimum Disks required	2
Capacity	N/2
Redundancy	Yes

RAID1 mirrors the data stored in one hard drive to another. By Infortrend's definition, RAID1 can only be performed with two hard drives. If there are more than two hard drives, RAID (0+1) will be automatically applied.

RAID (0+1) - Disk Striping with Mirroring

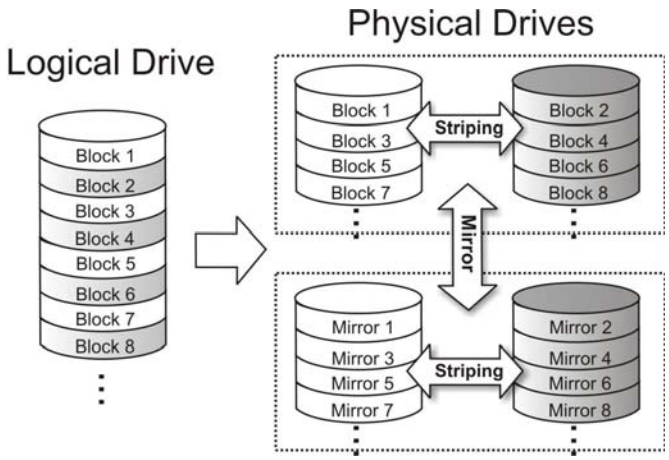

RAID (0+1)	
Minimum Disks required	4
Capacity	N/2
Redundancy	Yes

RAID (0+1) combines RAID0 and RAID1 - Mirroring and Striping. RAID (0+1) allows multiple drive failures because of the full redundancy of mirrored pairs. Multiple members can fail if they are not in a mirrored pair. If there are more than two hard drives included in a RAID1, RAID (0+1) will be automatically applied.

IMPORTANT!

"RAID (0+1)" will not appear in the list of RAID levels supported by the controller. If you wish to perform RAID1, the system firmware will determine whether to perform RAID1 or RAID (0+1). This will depend on the number of disk drives selected to compose a logical drive.

RAID3 - Disk Striping with Dedicated Parity Disk

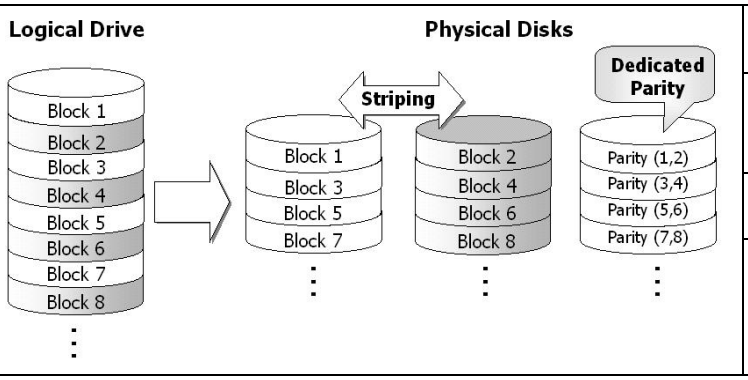

RAID3	
Minimum Disks required	3
Capacity	N-1
Redundancy	Yes

RAID3 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 or regenerate the lost data in the failed drive by comparing and re-calculating data on the remaining members.

RAID5 - Striping with Interspersed Parity

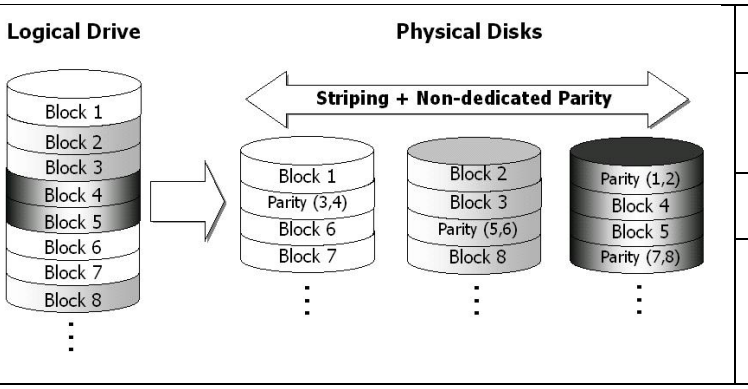

RAID5	
Minimum Disks required	3
Capacity	N-1
Redundancy	Yes

RAID5 is similar to RAID3 but the parity data is not stored in a dedicated hard drive. Parity information is interspersed across all members of the logical drive. In the event of a drive failure, the controller can recover or regenerate the lost data of the failed drive by comparing and re-calculating data on the remaining members.

RAID6 - Striping with Redundant (P+Q) Parity Scheme

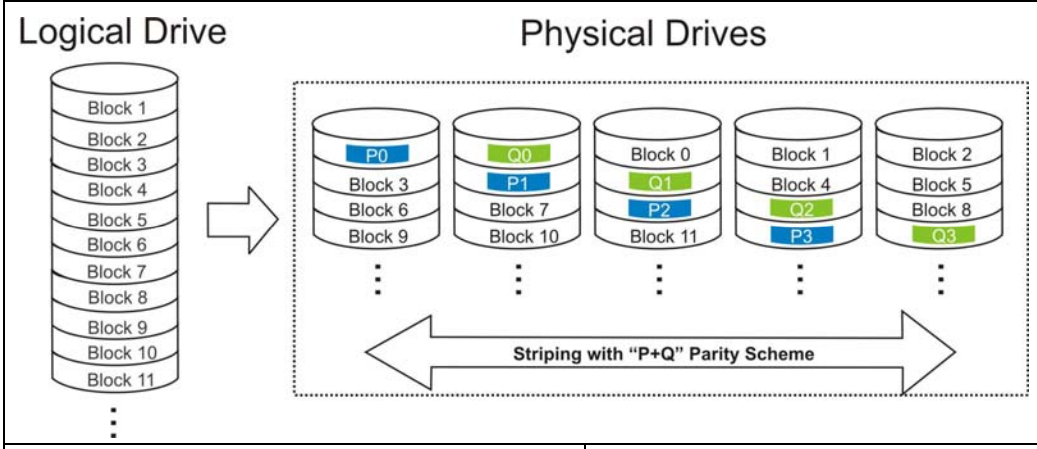

NOTE: A RAID6 array can withstand simultaneous failures of two disk drives, or one drive failure and bad blocks on another member drive.

RAID6	
Minimum Disks required	4
Capacity	N-2
Redundancy	Yes

RAID5 has been popular for it provides combined performance from its member drives and reasonable protection against a single disk failure. However, when storage systems grow larger and need to serve a wide variety of applications, the RAID5 protection can be insufficient. In the event of single drive failure, the occurrence of bad blocks on another member drive can render the affected data stripes unusable. RAID6 improves RAID5 and provides significantly higher redundancy level in terms of its ability to withstand two simultaneous drive failures.

RAID6 is similar to RAID5 but two parity blocks are available within each data stripe across the member drives. Each RAID6 array uses two (2) member drives for storing parity data. The RAID6 algorithm computes two separate sets of parity data and distribute them to different member drives when writing to disks. A RAID6 array requires the capacity of two disk drives for storing parity data.

Each disk drive contains the same number of data blocks. Parity information is consequentially interspersed across the array following the preset algorithms. A RAID6 array can tolerate the failure of more than one disk drive; or, in the degraded condition, one drive failure and bad blocks on the other. In the event of disk drive failure, the controller can recover or regenerate the lost data of the failed drive(s) without interruption to normal I/Os.

Other RAID Levels

RAID levels 10, 30, and 50 are implemented as logical volumes. RAID volumes are stripe sets of logical drives. If a logical volume consists of a RAID3 and a RAID5 logical drive, it is not easy to define its RAID level.

Chapter

3

Sample RAID Configuration Procedure

1. Planning Applications, LUNs, and Logical Drives

Planning helps you avoid configuration errors and is crucial for facilitating the process. Use two methods with the planning:

1. Sketch your connections and applications. You may refer to the samples in Chapter 1 or those in your system Hardware Manual. They can help linking the key elements.

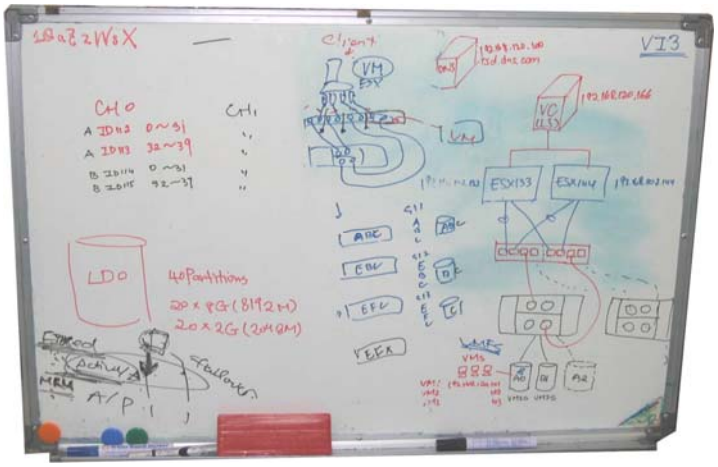

Use a notebook and sketch the planned application for future reference.

2. Use Worksheets to keep a hard record of how your storage is configured. An example is shown below:

Application	File system	RAID level of LUN	LUN ID	LUN capacity	Server details (OS)	Host links info. (HBA, switch, etc.)

You can expand the worksheet to include more details such as the disk drive channel on which the disks reside, JBOD enclosure ID, whether the LUNs are shared, and shared by which servers, etc.

2. Selecting Hard Drives

It is best to use hard drives of the same capacity and the same rotation speed in an enclosure.

You should not:

1. Include drives of different capacities in a logical drive.
2. The mixed use of SAS and SATA drives in an enclosure is allowed. However, including drives of different RPM and capacities in a logical drive should be avoided.
3. Use a smaller-size hard drive as a hot spare. See **Appendix 2** for details on using hot spares.

3. Drive Location:

Tray Numbering:

The same disk tray layout always applies to all Infortrend's storage enclosures. Trays are numbered, from left to right and then from top to bottom. It is advised you select members for a logical drive following the tray numbering rule, in order to avoid confusing yourself using the LCD keypad or the text-based firmware utility.

For example, a typical single enclosure configuration can look like this:

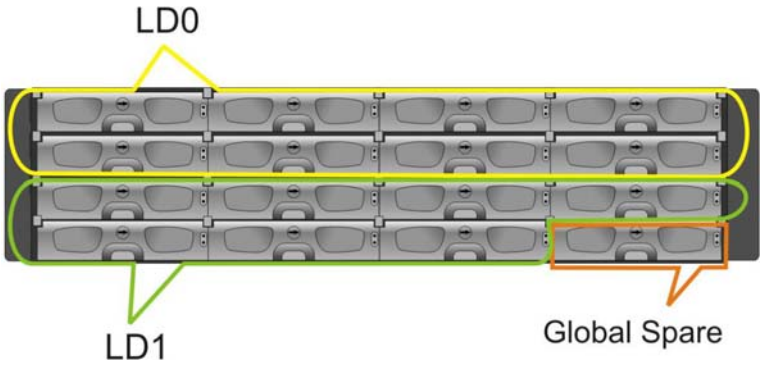

Disk drives in slots 1 to 8 are included in LD0, Logical Drive #0.

Disk drives in slots 9 to 15 are included in LD1, Logical Drive #1.

Slot 16 is configured as a Global Spare, which will participate in the rebuild of any logical drives.

A firmware utility screen to physical drive information looks like this. Following drive numbering sequence helps avoid configuration errors.

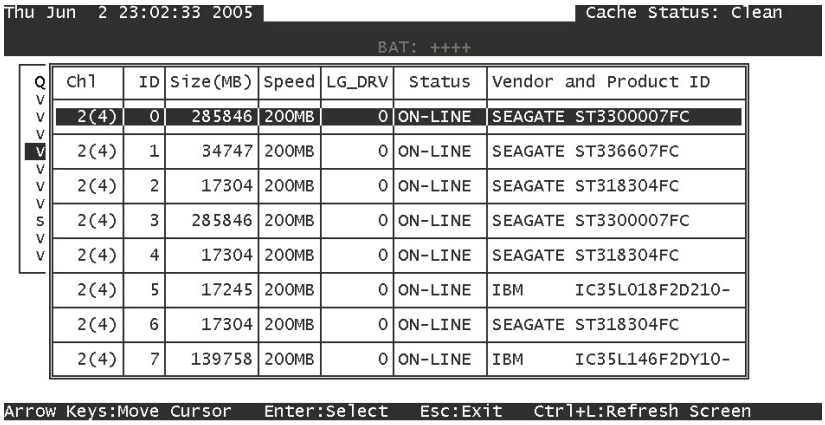

4. Connecting Management Interfaces:

Serial cables for the RS-232 COM1 serial port:

DB9-to-phone jack	DB9-to-phone jack, Y-cable
	

Step 1. Use the included serial cable to connect the COM1 serial ports. COM1 is always located on the RAID controllers.

Step 2. If your system is powered by a single RAID controller, connect the single end-to-end cable.

If your system is powered by redundant RAID controllers, use the Y-cable.

If you prefer a telnet console, connect Ethernet cables to the controllers' 10/100BaseT Ethernet ports.

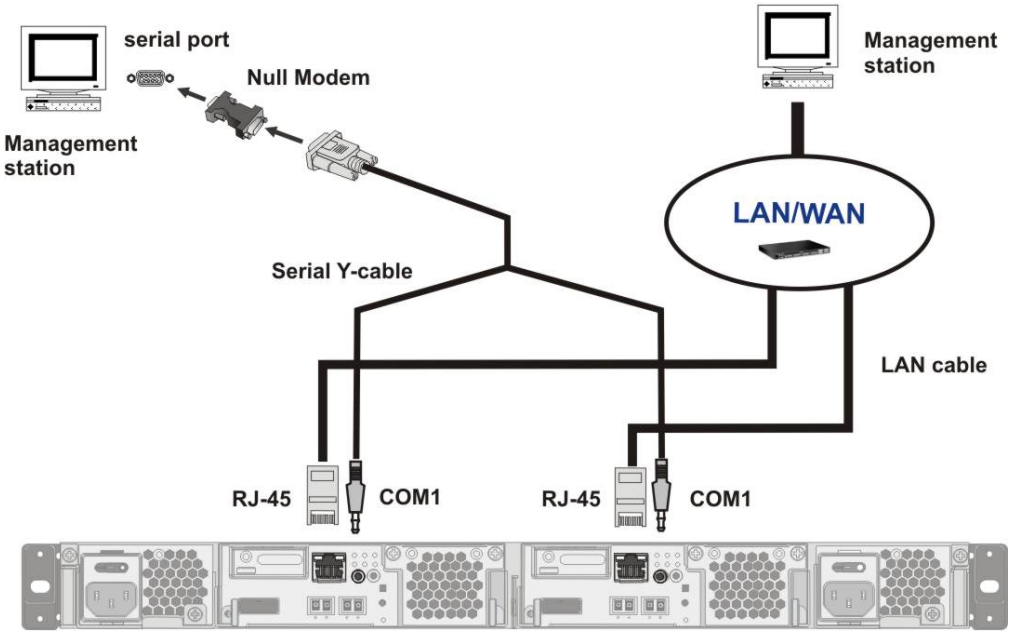

Step 3. If using the serial port connection for local management, attach a null modem to the DB9 end of the serial cable.

5. Opening a Management Console:

1. **Serial Console Using a Terminal Emulation Program**

   Step 1. Locate and start a hyper terminal program. For example, the Windows program menu: **Start** -> **Accessories** -> **Communications** -> **Hyper Terminal**.

   Step 2. The program starts with an input field requiring you to enter a name for the connection.

Step 3. The next screen requires you to select a serial port on your PC.

Step 4. Select appropriate baud rate and data/stop bit values (identical to those set for the COM1 port on your RAID subsystem). Click OK, and you should then be able to establish a management console. The firmware defaults are:

Baud rate	38400
Data bit	8
Parity	none
Stop bit	1
Flow control	Hardware

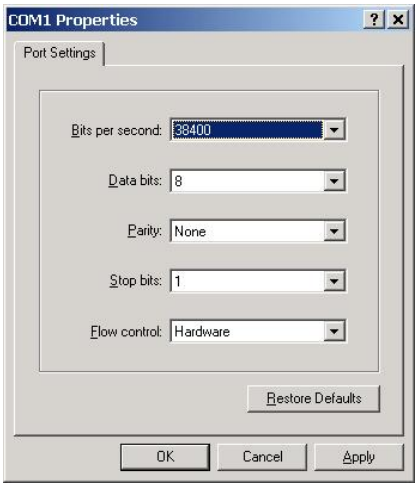

Step 5. The initial screen for the text-based utility should display.

Use the following keys to start using the utility:

←→↑↓	To move around menu options
[Enter]	To enter a sub-menu or to execute a selected option
[Esc]	To cancel an option or return to the previous menu
[Ctrl]+[L]	To refresh the screen information

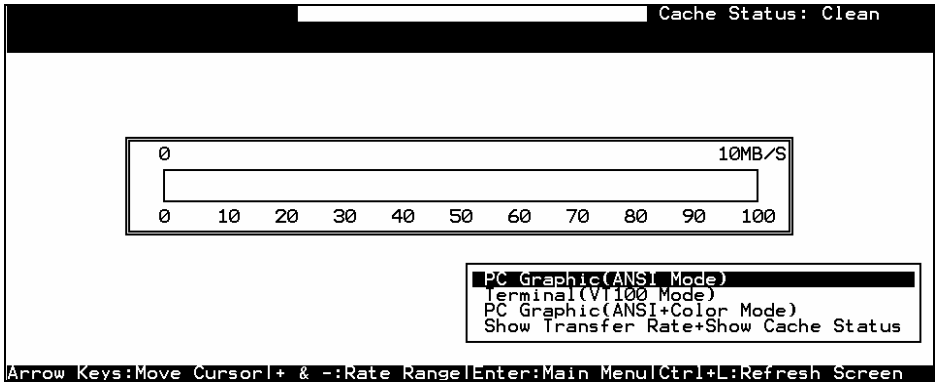

Step 6. Use the cursor keys to select a display mode. Press Enter to enter the main menu.

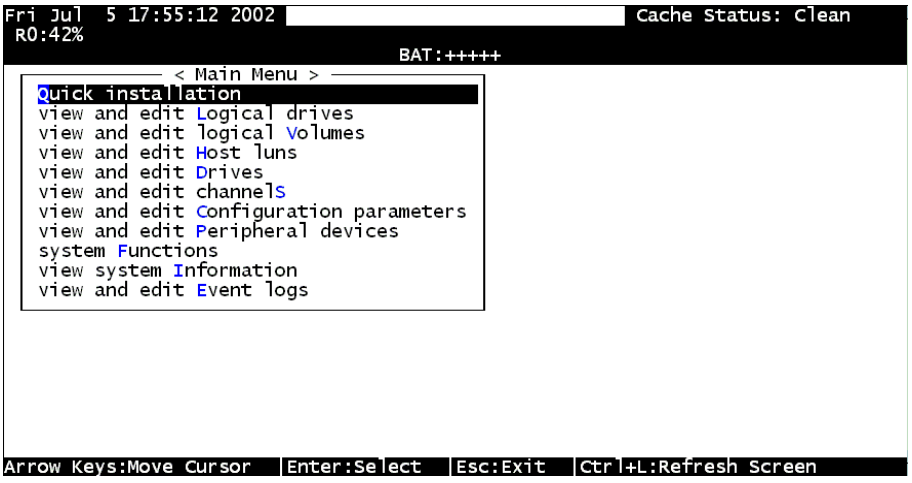

2. Telnet via Ethernet

Step 1. Use an Ethernet cable with RJ-45 phone jacks to connect the Ethernet port on the controller module.

Step 2. Connect the other end of the Ethernet cable to your local area network. An IP address should be acquired for the subsystem's Ethernet port. The subsystem firmware also supports automatic client configuration such as DHCP.

Step 3. Consult your network administrator for an IP address that will be assigned to the system Ethernet port.

Step 4. Use the LCD keypad or RS-232 console to select "**View and Edit Configuration Parameters**" from the Main Menu on the terminal screen. Select "**Communication Parameters**" -> "**Internet Protocol (TCP/IP)**" -> press **ENTER** on the chip hardware address -> and then select "**Set IP Address**."

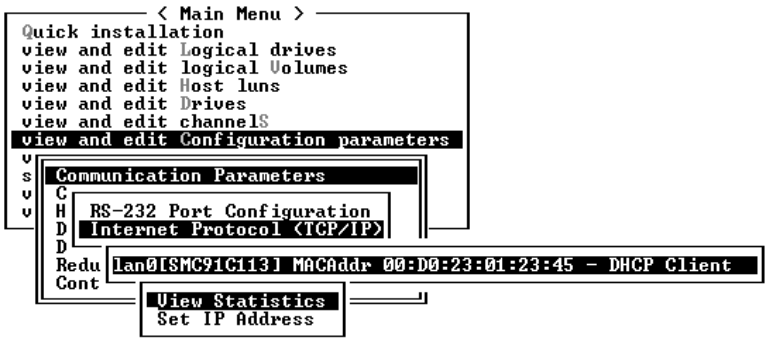

If DHCP is preferred, just key in “DHCP” in the Set IP Address field.

NOTE:

The IP default is “DHCP client.” However, if DHCP server can not be found within several seconds, a default IP address “10.10.1.1” will be loaded. This feature is available in the EonStor ASIC400 models.

Step 5. Provide the IP address, NetMask, and Gateway values accordingly.

Step 6. PING the IP address from your management computer to make sure the link is valid.

Step 7. Open a command prompt window and key in “telnet xxx.xxx.xx.xxx (controller IP address)” to access the embedded firmware utility.

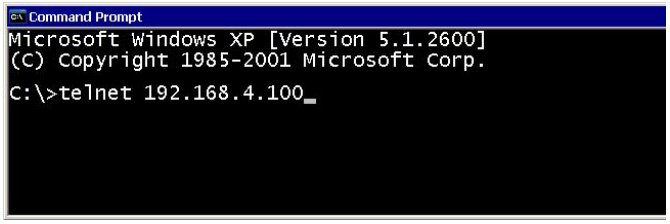

Step 8. Enter the preset password for accessing the storage system. If there is no preset password, press Enter to proceed.

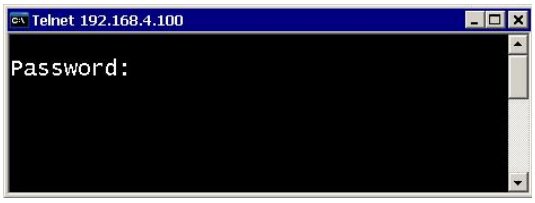

NOTE:

A management console using SANWatch or the web-based Embedded RAIDWatch is not the topic of this document. Please refer to their specific user documents for details.

3. Secure Link over SSH

Firmware supports remote management over the network connection and the security under SSH (Secure Shell) protection. SSH is widely used for its ability to provide strong authentication and secure communications over insecure channels. The SSH secure access can also be found as an option in the connection window of the **SANWatch** management software.

SSH is more readily supported by Linux- or Unix-based systems. The support for SSH on Microsoft Windows platforms can be limited.

For making SSH link using Windows, there are SSH tools such as the “PuTTY” shareware.

If a shareware is used, it may be necessary to configure the display options, e.g., the “Character set translation on received data” and “font type” setting in order for the terminal screen to be correctly displayed. The appearance settings may vary on different SSH tools.

Character set translation setting:

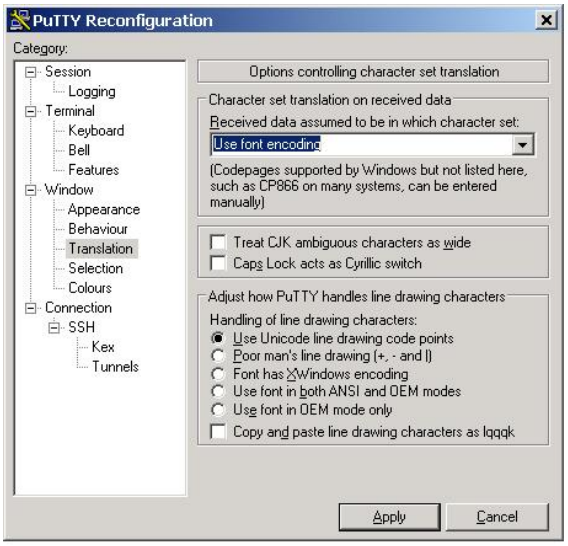

Appearance menu:

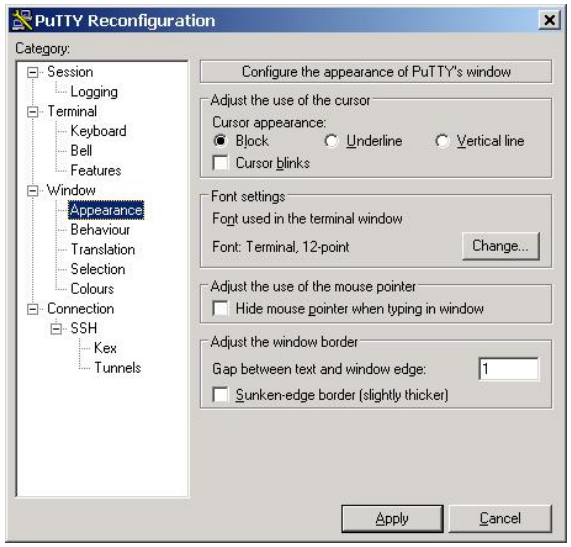

Font type menu:

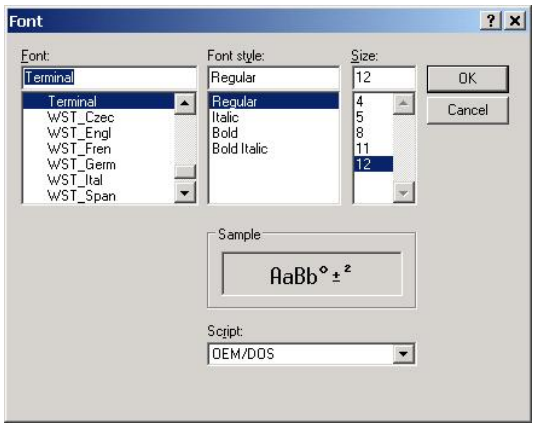

6. Creating RAID Elements

Step 1. Make sure all physical drives are properly installed by checking the View and Edit Drives menu.

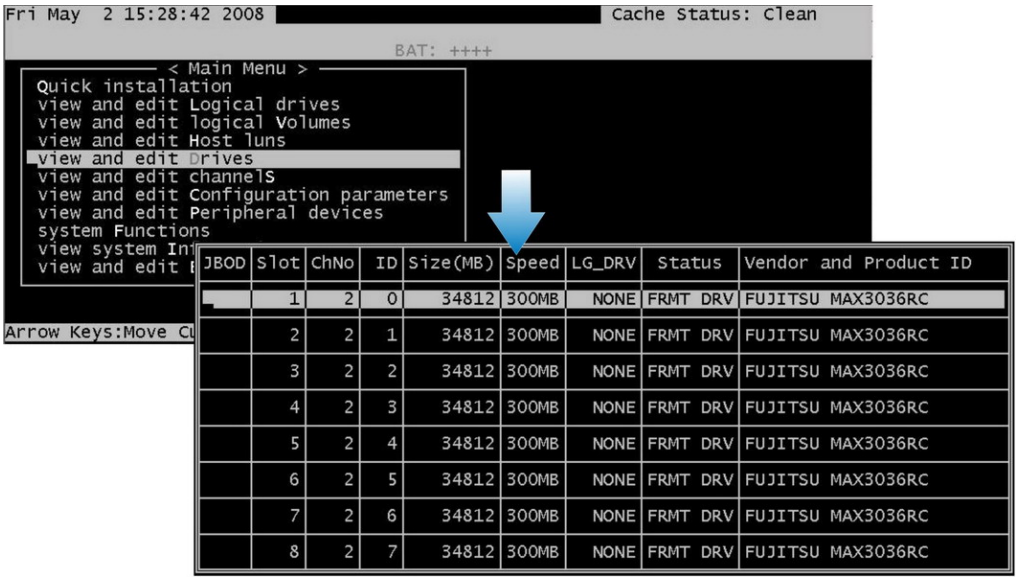

JBOD	Slot	ChNo	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
	1	2	0	34812	300MB	NONE	FRMT DRV	FUJITSU MAX3036RC
	2	2	1	34812	300MB	NONE	FRMT DRV	FUJITSU MAX3036RC
	3	2	2	34812	300MB	NONE	FRMT DRV	FUJITSU MAX3036RC
	4	2	3	34812	300MB	NONE	FRMT DRV	FUJITSU MAX3036RC
	5	2	4	34812	300MB	NONE	FRMT DRV	FUJITSU MAX3036RC
	6	2	5	34812	300MB	NONE	FRMT DRV	FUJITSU MAX3036RC
	7	2	6	34812	300MB	NONE	FRMT DRV	FUJITSU MAX3036RC
	8	2	7	34812	300MB	NONE	FRMT DRV	FUJITSU MAX3036RC

Use arrow keys to scroll down and make sure installed hard drives are all present. The list can be a long one if you attach expansion JBODs.

HDDs in a JBOD are identified by the number in the JBOD column.

JBOD	Slot	ChNo	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
		2	16				SES	VITESSE VSC7153 Eval Brd
0		3(4)	0	140016	300MB	NONE	FRMT DRV	FUJITSU MAX3147RC
0		3(4)	1	140016	300MB	NONE	FRMT DRV	FUJITSU MAX3147RC
0		3(4)	2	140016	300MB	NONE	FRMT DRV	FUJITSU MAX3147RC
0		3(4)	3	140016	300MB	NONE	FRMT DRV	FUJITSU MAX3147RC
0		3(4)	4	140016	300MB	GLOBAL	STAND-BY	FUJITSU MAX3147RC
0		3(4)	5	140016	300MB	GLOBAL	STAND-BY	FUJITSU MAX3147RC
0		3(4)	6	140016	300MB	NONE	FRMT DRV	FUJITSU MAX3147RC

Step 2. Use the ESC key to return to the Main Menu. Now you can go to the View and Edit Logical Drives menu to begin RAID configuration.

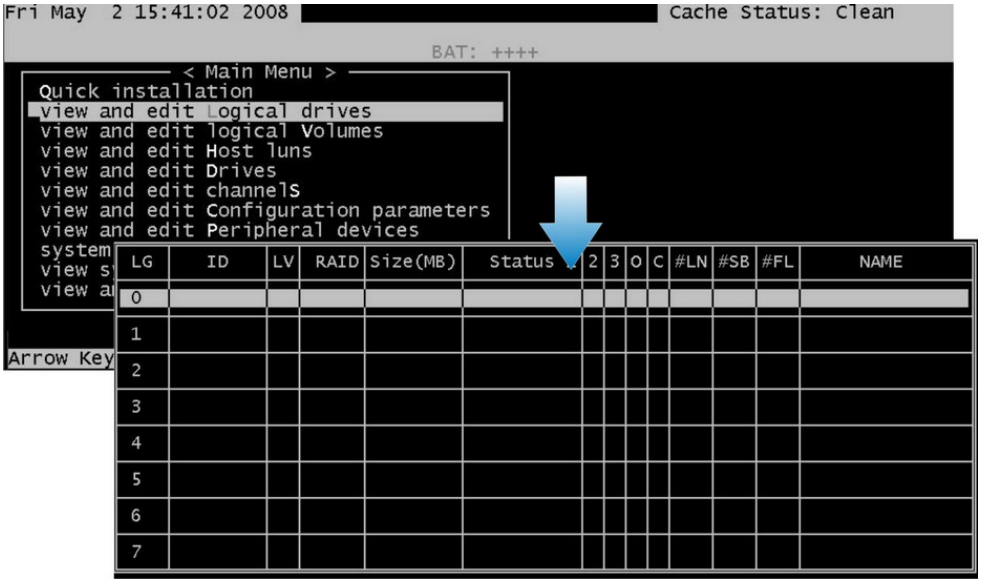

Step 3. Select an index number by pressing Enter on it, usually the configuration starts from LG0. Confirm your selection by moving highlighted area to Yes and press Enter.

Step 4. Select a RAID level.

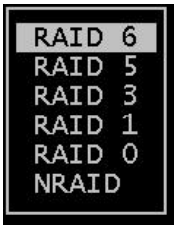

Step 5. Select members to be included in the logical drive by moving the highlighted color bar and pressing Enter on each drive. A selected member will be highlighted and its index number shown in the index column.

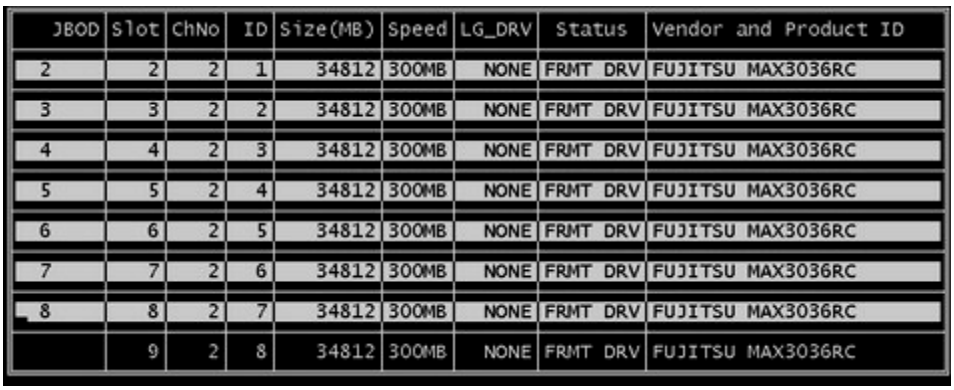

The above screen shows that 8 members have been selected. The number of members is determined by

the enclosure and also the performance concerns mentioned earlier in this document.

If you have a 24-bay enclosure, you might as well create 2 12-member LDs or 3 8-member LDs. With a 12-bay enclosure, you can compromise with 2 6-member LDs.

Step 6. Press the ESC key when you have selected all members. An LD parameters window will prompt.

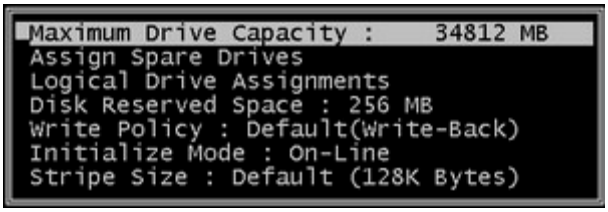

Step 6-1.

The first option, Maximum Drive Capacity, is useful if you suspect your drive members may have slightly different block numbers, which determines the actual drive capacity you can allocate from each drive.

Setting the Max. Drive Capacity slightly lower can get around the issue that one of the members can actually be slightly smaller. Chances are some blocks in some drives might have been marked as defective by drive manufacturers before shipping, and hence the usable number of blocks is reduced. For Infortrend's system firmware, all members in a logical drive must be of the same capacity and speed.

You can also specify half of the size. The unused capacity can later be utilized as a secondary RAID partition using the RAID expansion function.

Step 6-2.

This is where you specify a Local or Dedicated spare drive. For details, please refer to Appendix 2. The Dedicated spare only joins the rebuild of this logical drive.

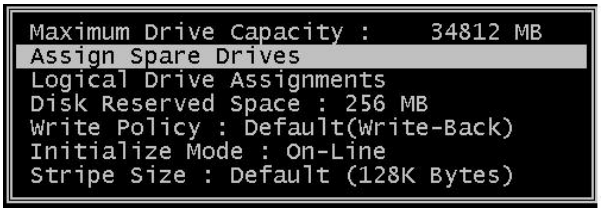

Step 6-3.

If you are configuring LDs for a redundant-controller system, you can equally assign LDs to both controllers so that the computing power of the partner controllers can be fully utilized.

For example, if you have 4 LDs, you can assign 2 LDs to controller A and another 2 to controller B.

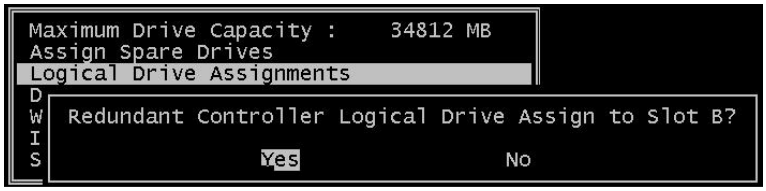

Step 6-4.

The Reserved Space option is view-only, skip this option. The space is automatically segregated for keeping logical drive configuration data.

Step 6-5.

Write-back caching can significantly enhance LD performance. Write-through is only selected if you do not have the protection of battery backup.

The “Default” option enables the LD’s caching policy to be automatically adjusted to a system-level caching policy, which is dynamically disabled in critical events such as component failures or thermal alarm. The system-level option is found in View and Edit Configuration Parameters -> Caching Parameters.

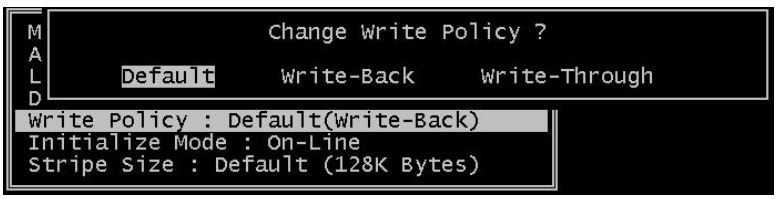

Step 6-6.

The Online Initialization Mode allows you to continue with the rest of the system setup steps without having to wait for the logical drive to be fully initialized. Initializing an LD terabytes in size can take hours.

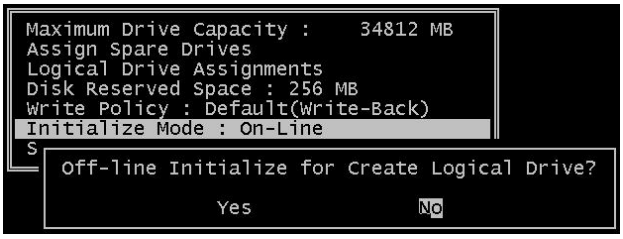

Step 6-7.

The default stripe size (128KB) is applicable to most applications. The stripe size can be adjusted in situations when the I/O characteristics are predictable and simple. For example, logical drives in a RAID system serving an AV stream editing application have a dedicated purpose. In such environment, you can match the size of host I/O transfers to the LD stripe size so that 1 or 2 host I/Os can be efficiently served within a parallel write.

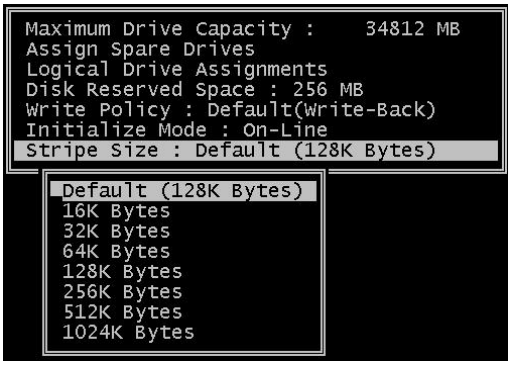

Step 7. Press the ESC key once you have set all configurable details. A confirm message box will prompt. Check the details before moving to the Yes option. Press Enter on Yes to begin the creation process.

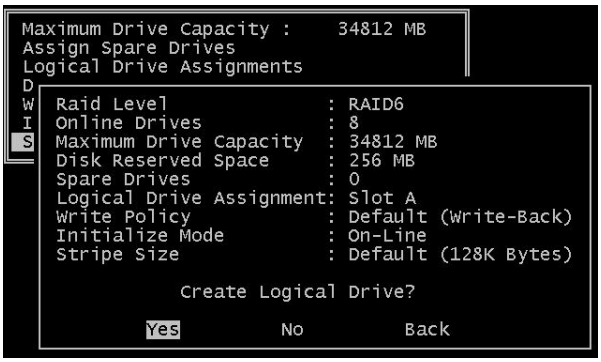

Step 8. A succession of events will prompt. Use the ESC key several times to skip them if no erroneous events occurred.

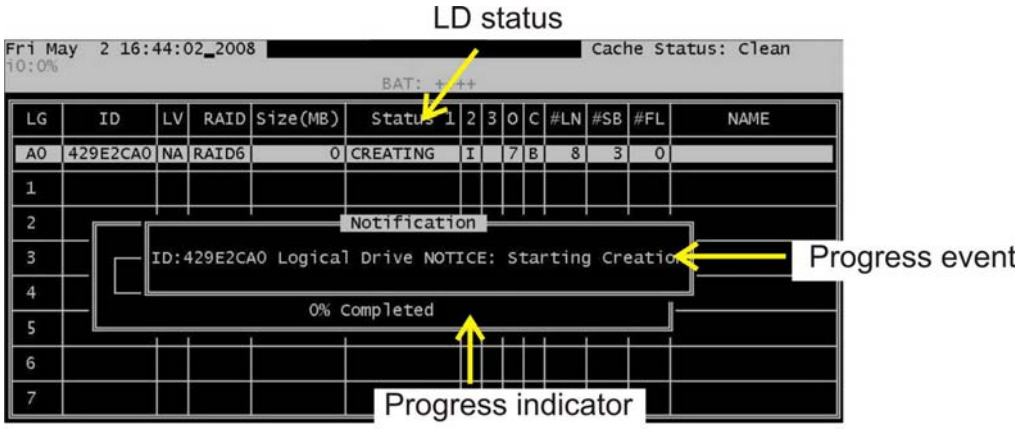

Step 9. Press ESC to hide this progress indicator. The progress bar will run in the background. If the online mode was selected, you can continue with the rest of the procedure, such as host LUN mapping.

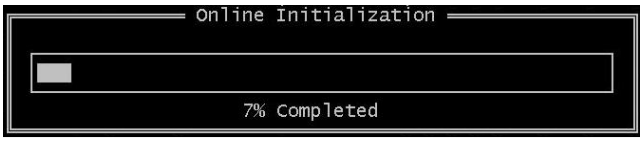

Step 10. You should return to the “View and Edit Logical Drives” screen. Press Enter on the LD you just created, and select “Logical Drive Name.” Enter a name for ease of identification, such as “ExchangeServer.”

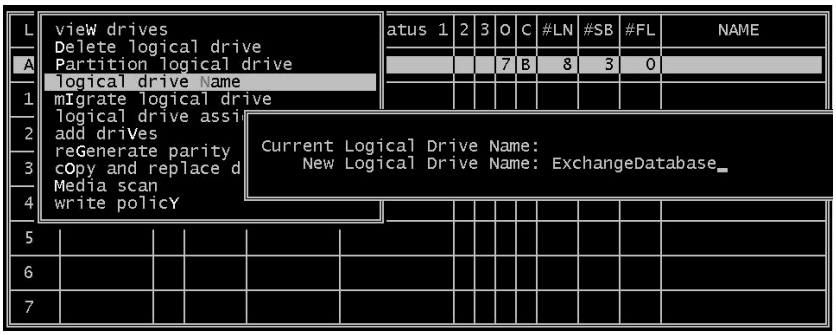

NOTE:

You may divide a logical drive or logical volume into partitions of desired capacity, or use the entire capacity as a single volume.

1. It is not a requirement to partition any logical configuration. Partitioning helps when multiple servers or applications need its disk space and you do not have the measures such as File Locking to prevent access contention.
2. With the concerns for the limited number of logical drives, partitioning can easily divide logical drives into volumes of the sizes you prefer.
3. You can not create partitions on a logical drive that already contains data. Partitioning will destroy data.

Step 11. Select another entry in the LD list to repeat the process to create more logical drives using the methods described above.

LG	ID	LV	RAID	Size(MB)	Status 1	2	3	0	C	#LN	#SB	#FL	NAME
A0	429E2CA0	NA	RAID6	208872	GOOD			7	B	8	3	0	ExchangeDatabas
1													
2													
3													
4													
5													
6													
7													

Step 12. Create more host IDs in the “View and Edit Channels” menu.

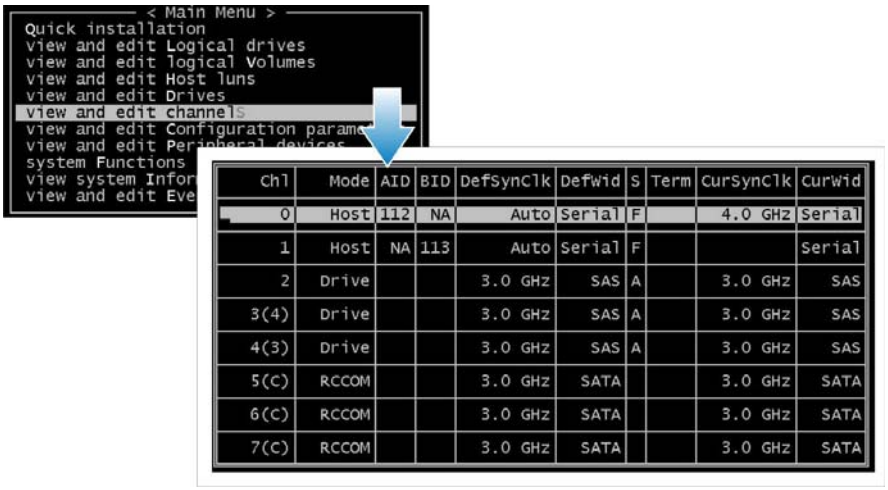

Step 12-1.

Press Enter to select a host channel.

Step 12-2.

Press Enter on View and edit SCSI ID.

Step 12-3.

Press Enter on any of the existing IDs.

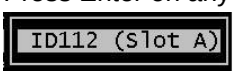

Step 12-4.

Press Enter to add host channel IDs.

Step 12-5.

Select Slot A or Slot B controller. Slot A and Slot B determines ownerships of logical drives. A logical drive associated with a Slot A ID will be managed by the Slot A controller (controller A); one associated with a Slot B ID by the Slot B controller.

Step 12-6.

Select an ID from the pull-down list.

Step 12-7.

Confirm the Add action by selecting Yes, and continue the Add ID process by selecting No. Repeat the process to create more AIDs or BIDs as is planned for your configuration.

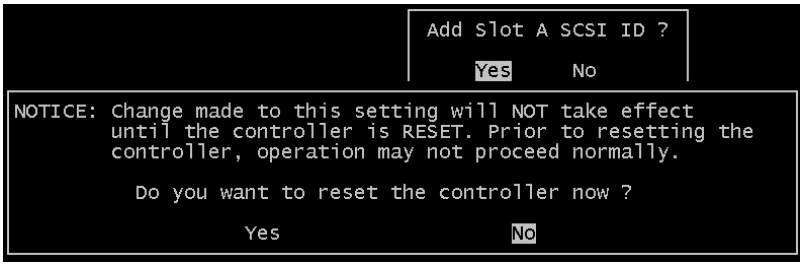

Step 14. Reset the controller after you created all the AIDs and BIDs you planned for your configuration.

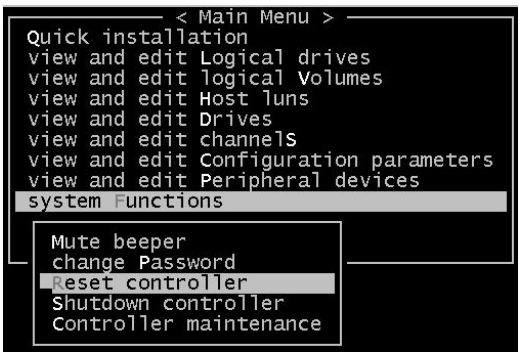

Step 15. A reset may take several minutes. Enter the View and Edit Host LUNs menu.

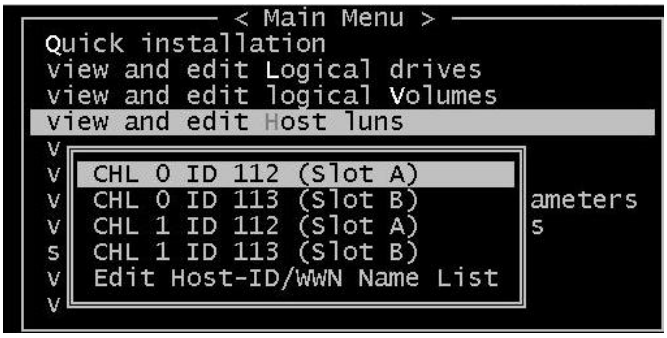

Step 16. Press Enter on a host ID. It is now necessary to refer to the topology plan you made previously. The below example makes for a dedicated DAS topology.

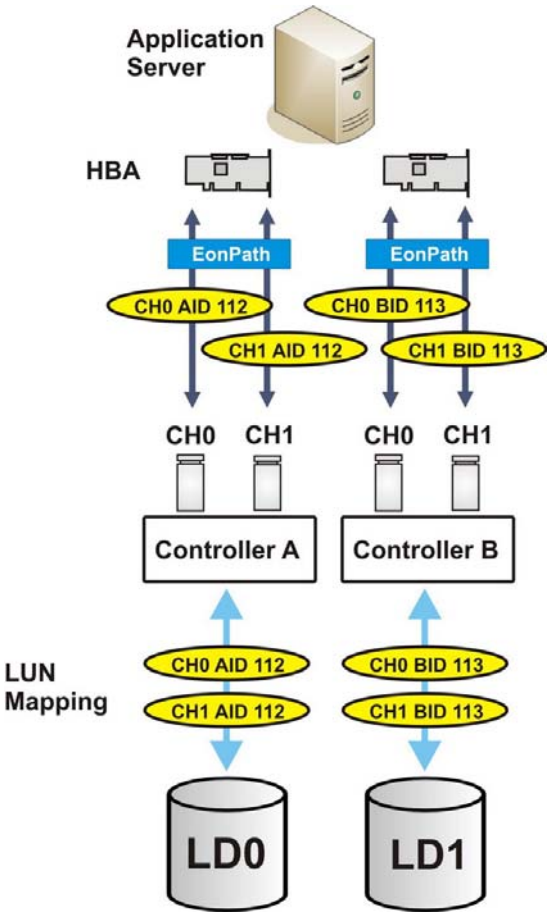

The LUN mapping process associate LDs with host channel IDs, and in this way LDs are presented through different host links.

The topology here only shows a basic, direct-attached configuration. Mapping multiple volumes in a SAN environment can be more complicated.

The complete LUN mapping steps are as follows:

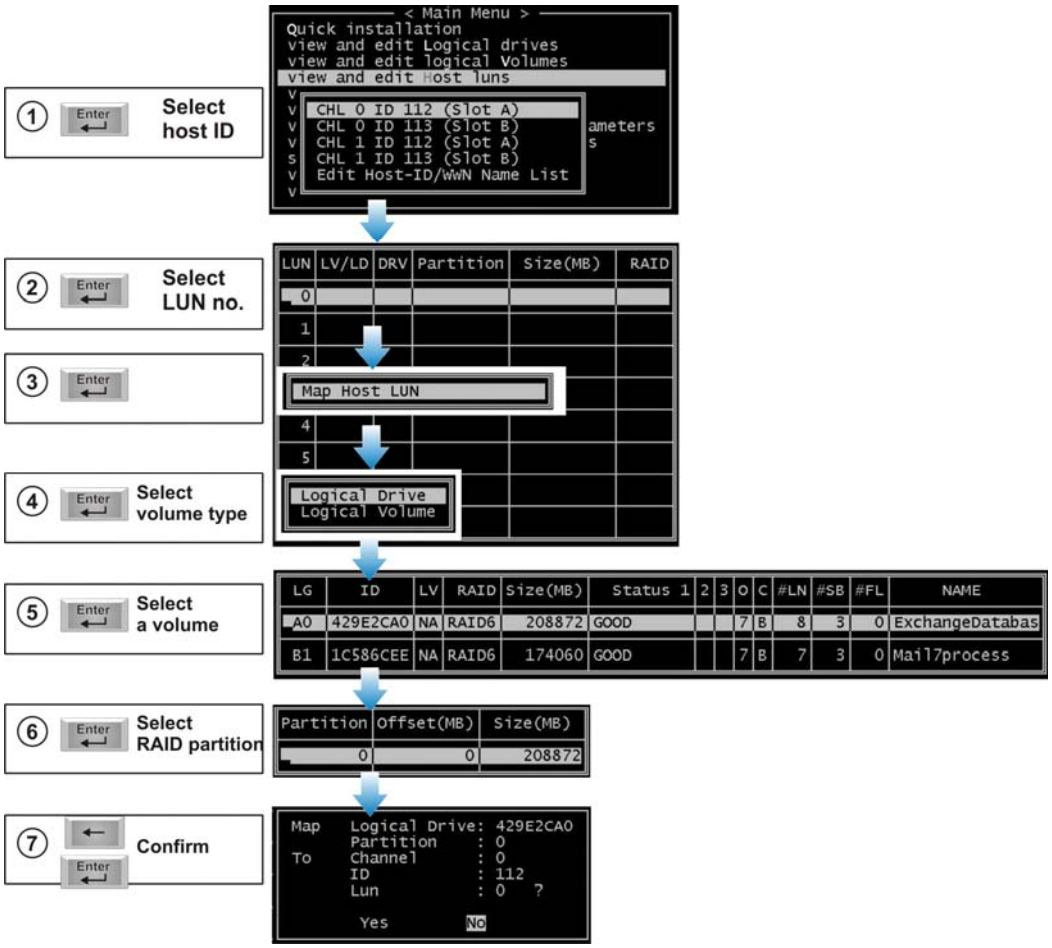

Step 16-1. Select a host channel ID. Note it is a Slot A or Slot B ID.

Step 16-2. Select an LUN number under this ID.

Step 16-3. Press Enter on seeing Map Host LUN.

Step 16-4. Select the volume type you are mapping to this host ID, Logical Drive or Logical Volume.

Step 16-5. Select a logical drive. Note the **LG** column. **A0** indicates the first LD, LD0, is assigned to controller A. The A0 LD is managed by controller A.

Step 16-6. Select a RAID partition within the LD. In this case, there is only one partition. Press Enter to proceed.

Step 16-7. Confirm your LUN mapping. It is recommended to check the details against your application plan and worksheet.

Step 17. Repeat the mapping process until you present all your LDs properly on the host busses according to your application plan.

Step 18. You should then see the volumes on your application server (using Windows Server 2003 as an example).

2 LDs on 4 data paths will appear 4 devices in the Disk drives menu of the Computer Management utility.

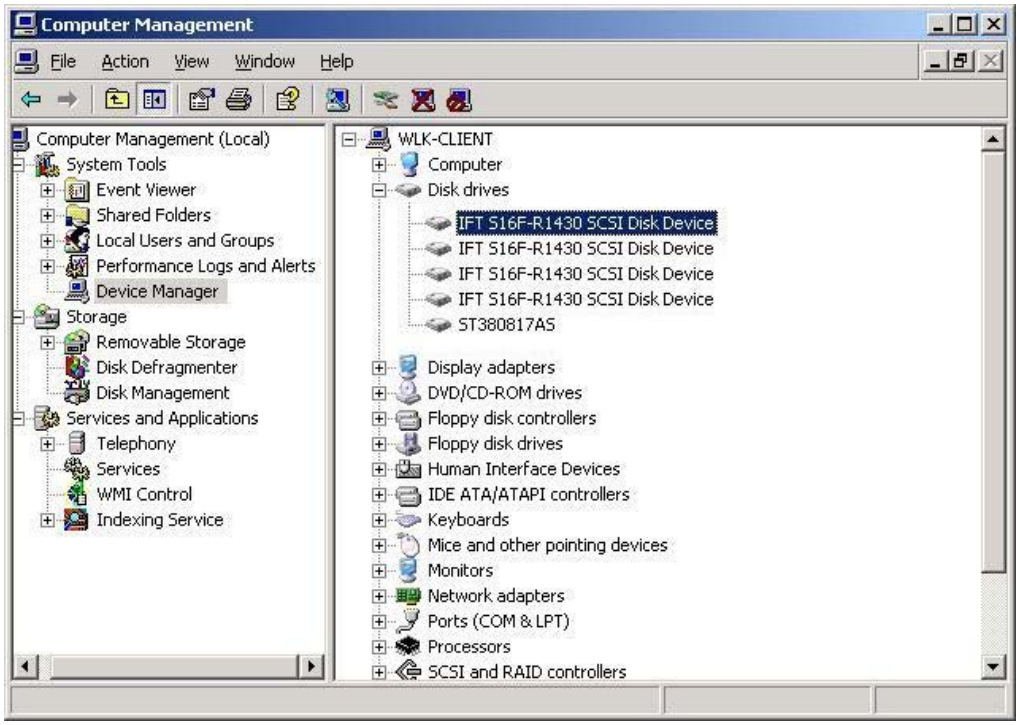

After installing the EonPath multi-pathing driver, the same LD appearing on 2 data paths will become a Multi-Path Disk Device. Installing EonPath requires your to reboot server. For details, please refer to EonPath's User Manual.

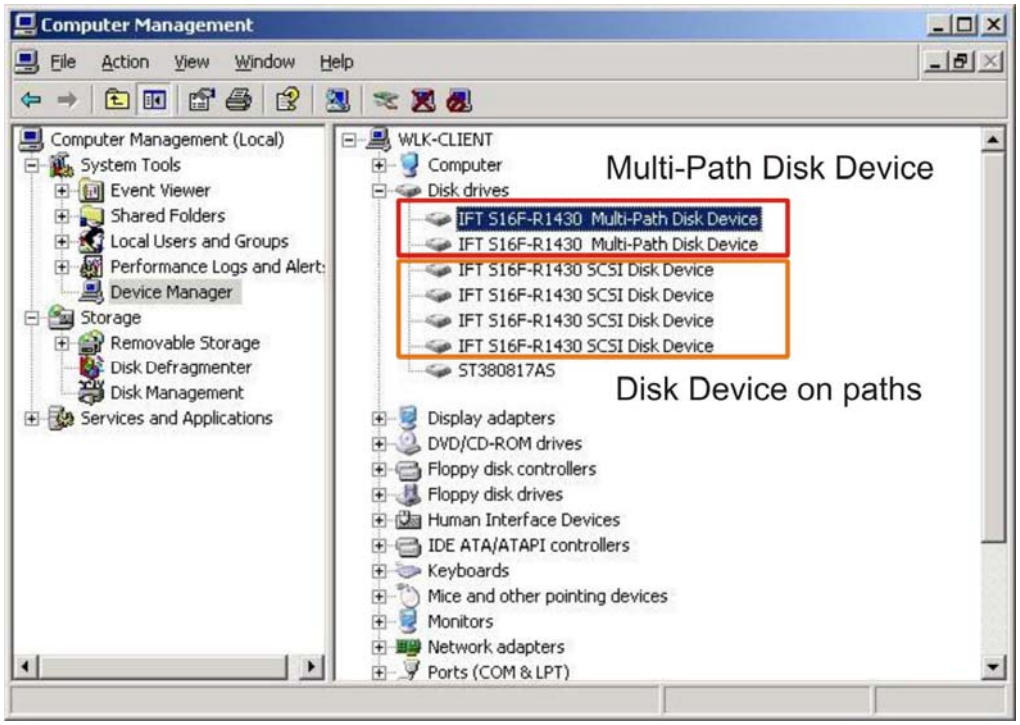

Configure and initialize the 2 LDs in the Disk Management window.

NOTE:

Make sure the firmware on your subsystem is EonPath compatible. Some earlier firmware revision, e.g., 3.42, may not work with EonPath.

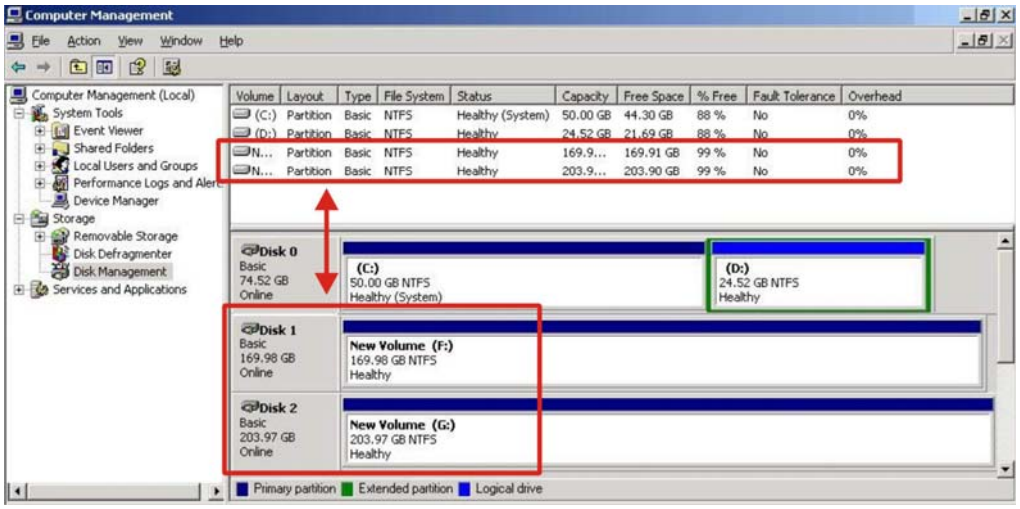

These 2 volumes are ready for use.

TIPS:

1. For the answers to some difficulties you might encounter during the initial configuration process, you can refer to Infortrend's website, the Support -> FAQ sections.

2. For specific, hardware-related details, such as the onboard hub or jumper settings, please refer to the Installation and Hardware Reference Manual that is included with your system package.

Appendix

1

Tunable Parameters

Fine-tune the subsystem and the array parameters for your host applications. Although the factory defaults guarantee the optimized operation, you may refer to the table below to facilitate tuning of your array. Some of the performance and fault-tolerance settings may also be changed later during the preparation process of your disk array.

Use this table as a checklist and make sure you have each item set to an appropriate value.

(1)	Parameters that should be configured at the initial stage of system configuration
(2)	Parameters that can be changed later
(3)	Non-critical

Controller Parameter Settings

User-Defined Parameters	Default	Alternate Settings
Fault Management:		
(1) Automatic Logical Drive Rebuild - Spare Drive	Enabled when Spare Drive is available	RAID 1 + Local Spare RAID 3 + Local Spare RAID 5 + Local Spare RAID 6 + Local Spare Global Spare Enclosure Spare (recommended in a multi-enclosure configuration)
(1) S.M.A.R.T.	Disabled	Detect Only Perpetual Clone Clone + Replace Fail Drive
(3) Clone Failing Drive	Manual function	Replace After Clone Perpetual Clone
(1) Rebuild Priority	Low (higher priority requires more system resources)	Low Normal Improved High
(1) Verification on Write	Disabled	On LD Initialization On LD Rebuild On Normal Drive Writes
(2) Periodic Cache Flush	Disabled	Continuous to 10 minutes
(2) Periodic Auto-Detect Failure Drive Swap Check Time	Disabled	Disabled, 5 to 60 seconds

(2) Periodic Drive Check Time	Disabled	Disabled, 0.5 to 30 seconds Note this option is not necessary in models using serial drive busses such as SAS or Fibre.
(2) Rebuild Priority	Normal	Low, normal, improved, high

Controller:		
(1) Channel Mode	*	Host, Drive, RCCOM, Drive + RCCOM (RCC options not configurable in the ASIC 400 models)
(1) Host and Drive Channel IDs	*	* preset
(1) Controller Unique Identifier	Preset on most models	Hex number from 0 to FFFFF (FW 3.25 and above)
(2) Data Rate	Auto	Depends on problems solving
(1) Date and Time	N/A	
(1) Time Zone	+ 8 hrs	

Optimization:			
(1) Write-back Cache	Enabled	Disabled	
(1) LD Stripe Size	Related to controller general setting & application I/O characteristics		32KB to 1024KB
(2) Adaptive Write Policy	Disabled	Enabled	
(2) LD Write Policy	LD-specific or dependent on system's general setting		W/B or W/T

Host- and Drive-side Parameters:		
(1) Data Transfer Rate	*	Host Side: Asynchronous to 4GHz Drive Side: Asynchronous to 3GHz
(1) Max Number of Tags Reserved for each Host-LUN Connection	32	1 to 1024
(1) Maximum Queued I/O Count	32	1 to 1024
(2) LUN's per ID	8	Up to 32
(1) Auto Rebuild on Drive Swap	Disabled	5 to 60 seconds
(1) Number of Concurrent Host-LUN Connection	4	1 to 1024
NOTE: LUN-per-ID x tags reserved= flag A Max. Number of Concurrent Host-LUN connection= flag B If A>B, Max=A; else, Max=B		
(1) Tags per Host-LUN Connection	32	1 to 256
(1) Wide Transfer	*	Enabled/Disabled
(1) Drive I/O Timeout	7	0.5 to 30 seconds
(3) Drive Spindown Idle Delay Period	Disabled	60, 300, 600 sec

Spin-Up Parameters: (available on SATA-based models)		
(1) Drive Motor Spin-Up	Disabled	Enabled
(1) Disk Access Delay Time	25 seconds; 30 seconds	No Delay, 5 to 75 seconds

	for specific SATA disk drives	

Data Integrity:		
(3) Task Scheduler	N/A	Execute on initialization Start time and date Execution period Media scan mode Media scan priority Select Logical drive

Fibre Channel Parameters:		
(1) Fibre Connection Options	*	Loop Only Point-to-Point Only
(1) Fibre Channel Dual-Loop	Enabled	Enabled by cabling connection
(1) Host ID/WWN Name List	*	User configurable
(1) RCC through Fibre Channel	*	Dedicated or sharing drive channel(s)

Array Configuration:		
(1) Disk Reserved Space	256MB	
(1) AV Optimization Mode	Disable	Fewer Streaming Multiple Streaming
(1) Max Drive Response Timeout	Disabled	160, 320, or 960ms
(2) Array Assignment	Primary controller	Secondary controller
(1) Array Partitioning	1	Up to 64
(1) Auto-assign Global Spare	disabled	enabled

Enclosure Monitoring:		
(2) Event Triggered Operation	N/A	Controller, fan, PSU, BBU, UPS, and elevated temperature Auto-shutdown: 2 mins~1 hour
(1) Thresholds for Voltage and Temperature Self-Monitoring	CPU temp: 0~90˚C Board temp: 0~80˚C 3.3V: 2.9~3.6V 5V: 4.5~5.5V 12V: 10.8~13.2V	User-defined; do not change parameters unless necessary

Others:		
(3) Password	N/A	User-Defined; Password Validation Timeout: 1 second to Always Check Configurable
(3) LCD Display Controller Name	N/A	User-defined
(1) UPS support	N/A	COM2 baud rate and related settings; event triggered operation
(1) Cylinder/Head/ Sector Mapping	Variable	Depends on host OS

Supported RAID Configurations on Both Sides of the 1GB Threshold

Feature	Default Value	
	< 1GB DIMM	>= 1GB DIMM
64-bit LBA support (>2TB)	Yes	Yes
Number of LDs	16 (max.)	32 (max.)
Number of LVs (Logical Volume)	8 (max.)	16 (max.)
Number of Partitions per LD	16 (max.)	64 (max.)
Number of LUNs per channel ID	8 (32 max.)	8 (32 max.)
Number of LUNs	128 (max.)	1024 (max.)
Caching Mode	Write-back	
Stripe size, RAID5	128KB	
Auto-assign Global Spare	Disabled	
Max LD capacity	64TB max.	
No. of Media Scan Task by scheduler	16 max.	
Max. No. of member drives per DIMM size, RAID5	128 HDD/512MB	
NOTE: A maximum of 128 members in a logical drive is a theoretical number. Rebuilding or scanning such a logical drive takes a long time.		

Appendix

2

Protection by Hot Spares

Infortrend's firmware provides the flexibility with three different kinds of hot spare drives:

- Local (dedicated) Spare
- Enclosure Spare
- Global Spare

When any drives fail in a RAID1, 3, 5, 6 logical drive, the hot spares automatically proceeds with online rebuild. This paper shows how these three types function and introduces related settings.

Rebuild Policies Using Hot Spares

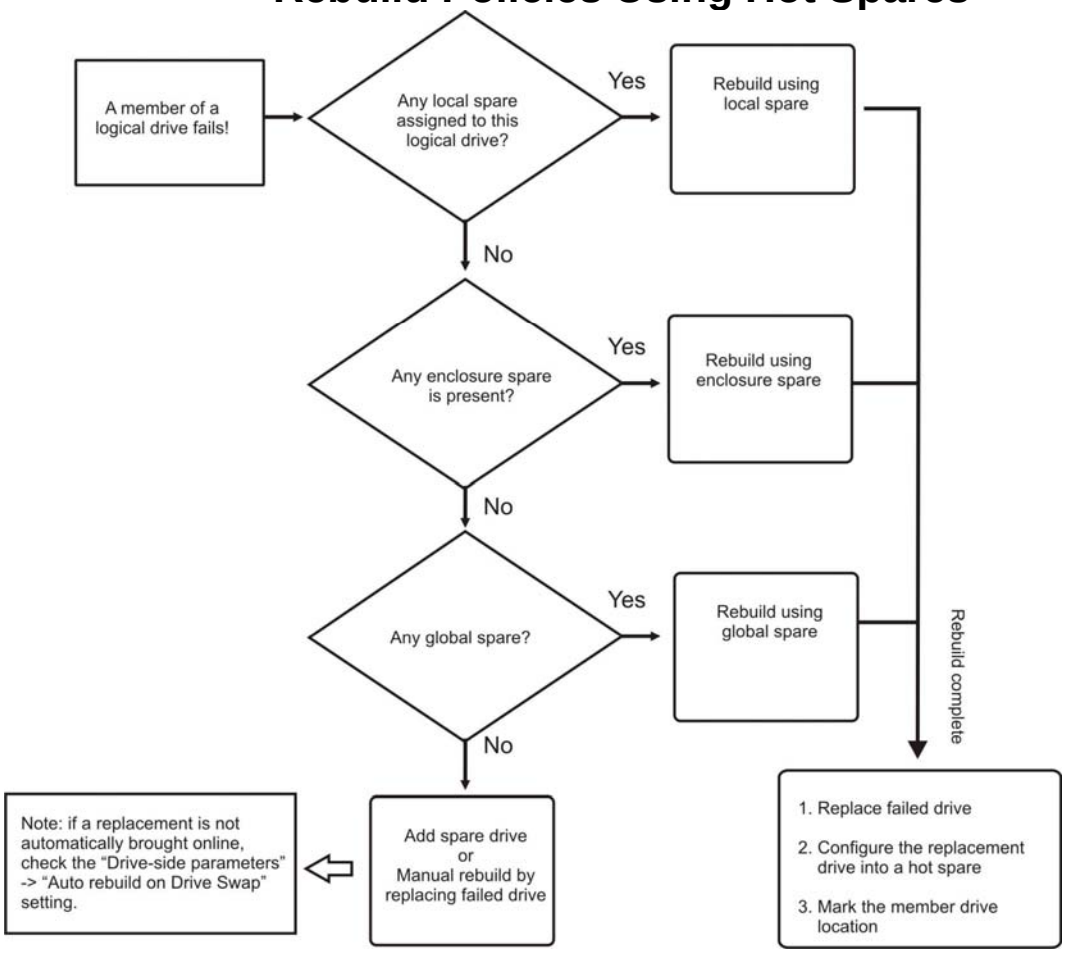

The mechanism above shows how the controller's embedded firmware determines whether to use Local, Enclosure, or Global Spares to rebuild a logical drive.

One important issue about rebuilding a logical drive is that users often forget to configure another hot spare after replacing a failed drive. Shown below is a standard procedure:

Typical Spare Rebuild Procedure

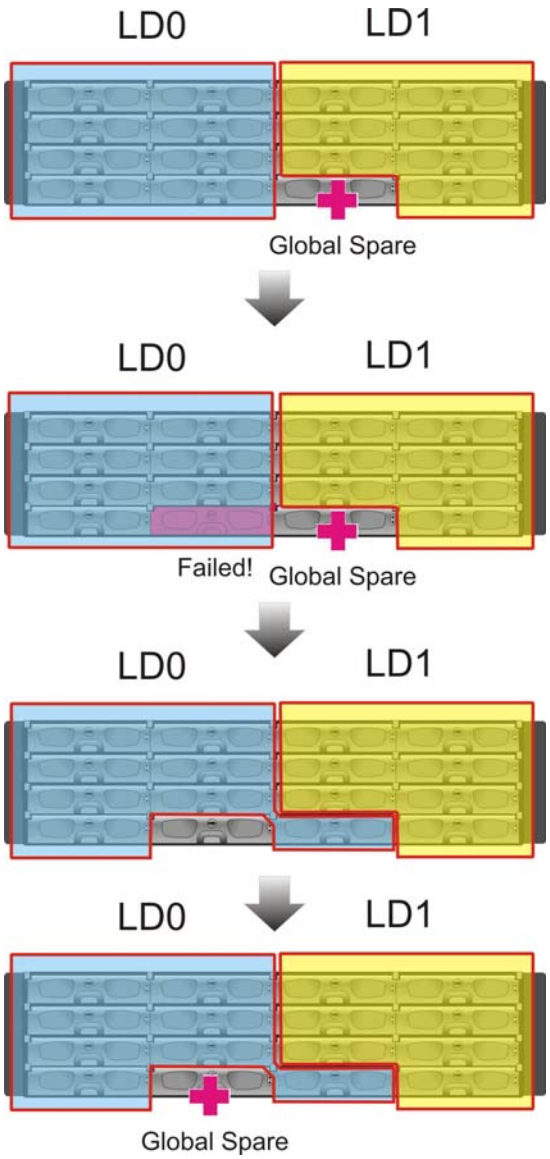

A typical 16-bay enclosure configuration is shown on the left.
Two logical drives: LD0 and LD1
Hot spare: Global

A member drive in LD0 fails.

The Global spare joins LD0 and automatically starts rebuild.
Note that the Global spare becomes the member of LD0.

The failed drive is replaced and configured as a Global spare. Doing so protects the array from another drive failure.

Auto-assign Global Spare

There is a way to prevent the situation when you forget to create a hot spare. Without hot spares, rebuild takes place only if a failed drive is replaced. Enable the function in firmware utility's "View and Edit Configuration Parameters" -> "Drive-side Parameters" -> "Auto-assign Global Spare."

Every disk drive that is not included in logical drives will be automatically configured into Global spares.

Strategies for Using Hot Spares

Users can assign specific disk drives as hot spares using RS232C terminal or SANWatch GUI.

Chapter 2 Local Spare Drive

If a particular logical drive is especially important to the user, it is recommended to assign a Local Spare Drive to this logical drive. A Local Spare drive only participates in the rebuild of the particular logical drive it is assigned to.

Enclosure Spare Drive

If a subsystem is connected to multiple expansion enclosures (JBOD or SBOD) or multiple subsystems are cascaded together, it is recommended to apply Enclosure Spare in each enclosure. An Enclosure Spare only participates in the rebuild of the logical drives within the same enclosure, and thus prevents the event of disorderly drive locations which happens when a Global Spare joins the rebuilding of a logical drive residing in another enclosure.

The below drawing shows the idea of a cross-enclosure rebuild:

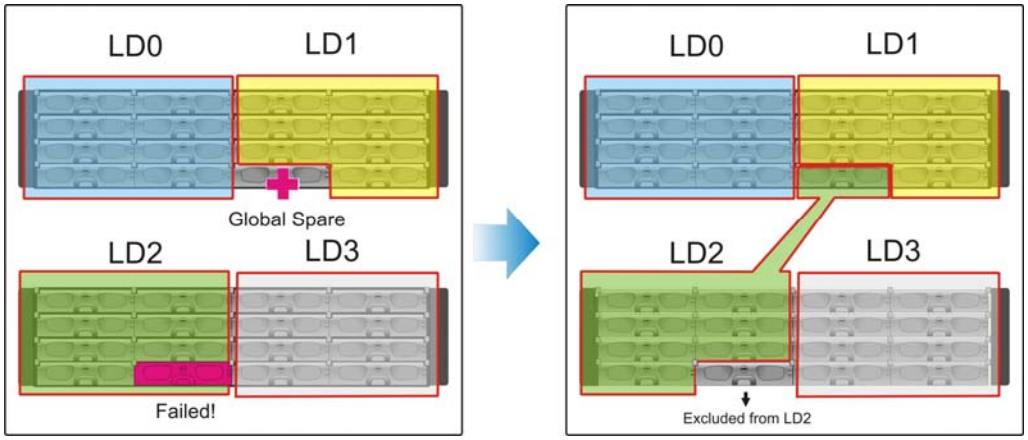

Having members across different enclosures may not bring ill effects on logical drive operation, however, it is easy to forget the locations of member drives and thus the chance of making mistakes will increase. For example, you may replace a wrong drive and destroy a logical drive when the logical drive is already in a degraded mode (having one failed member).

Global Spare Drive

A Global Spare Drive is a general hot spare which participates in the rebuild of all logical drives, even those in different enclosures. When Global spares are applied, make sure that the Global spare has a disk capacity equal or larger than all members in the array.

Spare Drive Limitation

Spare drives can only rebuild a logical drive with members of an equal or smaller capacity. Therefore, it is considered safer to tune down the Maximum Drive Capacity when creating logical drives. The Maximum Drive Capacity is the maximum of capacity used in each member drive to comprise a logical group. Some times disk drives labeled with the same capacity may actually come with different numbers of logical block units. With different block numbers, a slightly smaller spare may not be able to rebuild a logical drive composed of larger members.

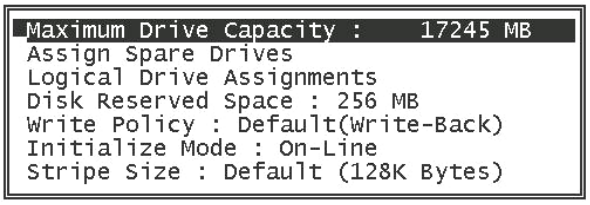

This page is intentionally left blank