MegaRAID® Express 000 Hardware Guide

MAN-762
10/2/99

American Megatrends, Inc.
6145F Northbelt Parkway
Norcross, GA 30071

Revision History

1/16/98	Initial release.
5/12/98	Corrected errors and deleted clustering information.
5/29/98	Added J1 description.
8/7/98	Removed references to 4 MB cache memory.
1/20/99	Added EDO DRAM support information.
10/2/99	Changed the product name to Express 000.

Table of Contents

Preface

Limited Warranty The buyer agrees if this product proves to be defective, that American Megatrends is obligated only to repair or replace this product at American Megatrends' discretion according to the terms and conditions of the warranty registration card that accompanies this product. American Megatrends shall not be liable in tort or contract for any loss or damage, direct, incidental or consequential resulting from the use of this product. Please see the Warranty Registration Card shipped with this product for full warranty details.

Limitations of Liability American Megatrends, Inc. shall in no event be held liable for any loss, expenses, or damages of any kind whatsoever, whether direct, indirect, incidental, or consequential (whether arising from the design or use of this product or the support materials provided with the product). No action or proceeding against American Megatrends may be commenced more than two years after the delivery of product to Licensee of Licensed Software.

Licensee agrees to defend and indemnify American Megatrends from any and all claims, suits, and liabilities (including attorney's fees) arising out of or resulting from any actual or alleged act or omission on the part of Licensee, its authorized third parties, employees, or agents, in connection with the distribution of Licensed Software to end-users, including, without limitation, claims, suits, and liability for bodily or other injuries to end-users resulting from use of Licensee's product not caused solely by faults in Licensed Software as provided by American Megatrends to Licensee.

Package Contents You should have received:

- a MegaRAID Express PCI SCSI Disk Array Controller,
- a CD with drivers, utilities, and documentation,
- a *MegaRAID Express 000 Hardware Guide* (on CD),
- a *MegaRAID Express 000 Configuration Software Guide* (on CD),
- a *MegaRAID Operating System Drivers Guide* (on CD),
- software license agreement (on CD), and
- a warranty registration card (on CD.)

Technical Support If you need help installing, configuring, or running the MegaRAID Express PCI SCSI Disk Array Controller, call American Megatrends technical support at 770-246-8600. Before you call, please complete the ***MegaRAID Express Problem Report*** form on the next screen.

Web Site We invite you to access the American Megatrends world wide web site at:

http://www.ami.com.

MegaRAID Express Problem Report Form

Customer Information	MegaRAID Express Information
Name	Today's Date
Company	Date of Purchase
Address	Invoice Number
City/State	Serial Number
Country	Number of Channels
email address	Cache Memory
Phone	Firmware Version
Fax	BIOS Version
System Information	
Motherboard:	BIOS manufacturer:
Operating System:	BIOS Date:
Op. Sys. Ver.:	Video Adapter:
MegaRAID Express Driver Ver.:	CPU Type/Speed:
Network Card:	System Memory:
Other disk controllers installed:	Other adapter cards installed:
Description of problem:	
Steps necessary to re-create problem: 1. 2. 3. 4.	

Logical Drive Configuration

Logical Drive	RAID Level	Stripe Size	Logical Drive Size	Cache Policy	Read Policy	Write Policy	# of Physical Drives
LD1							
LD2							
LD3							
LD4							
LD5							
LD6							
LD7							
LD8							

Physical Device Layout

	Channel 1	Channel 2
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		

	Channel 1	Channel 2
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		

Disclaimer

AMI certifies only that this product will work correctly when this product is used with the same jumper settings, the same system configuration, the same memory module parts, and the same peripherals that were tested by AMI with this product. The complete list of tested jumper settings, system configurations, peripheral devices, and memory modules are documented in the AMI Compatibility Report for this product. Call your AMI sales representative for a copy of the Compatibility Report for this product.

1 Overview

The MegaRAID Express PCI SCSI array controller adapter card works in a server that provides SCSI channels on the motherboard and also implements the American Megatrends IRQ routing logic.

MegaRAID Express is a RAID on Motherboard solution. RAID on the motherboard offers a cost-effective way to implement RAID in a server. MegaRAID Express includes MegaRAID features and performance, which offer the best way to differentiate a product in the world of Standard High Volume (SHV) servers.

Motherboard Requirements The motherboard in the system where MegaRAID Express will be installed must have an AMI interrupt routing module on the motherboard and Symbios Logic 53C875 SCSI controller chips on the motherboard.

Warning

The BIOS that controls the motherboard SCSI channels must be disabled to permit MegaRAID Express to control these SCSI channels correctly.

Warning

Normally, the BIOS the interrupt routing module is available for only one PCI expansion slot on a motherboard. This PCI expansion slot may be labeled RAID UPRGADE SLOT on the motherboard. See you computer owner's manual for additional information about the RAID Upgrade PCI expansion slot.

Cont'd

MegaRAID Express Features The MegaRAID Express features include:

- minimal cost because it does not use the American Megatrends advanced RAID ASIC,
- uses one reserved pin on the motherboard RAID expansion PCI slot,
- allows one or two motherboard SCSI channels to be controlled by MegaRAID Express,
- includes an Intel® i960RP/RD that performs RAID calculations and routing, and
- provides up to 128 MB of Fast Page Mode DRAM cache memory in a SIMM socket used for read and write-back caching and RAID 5 parity generation.

SCSI Channels MegaRAID Express does not have any SCSI channels. MegaRAID Express uses one or two SCSI channels on the motherboard.

> ***Warning***
>
> The BIOS that controls the motherboard SCSI channels must be disabled to permit MegaRAID Express to control these SCSI channels correctly.

NVRAM and Flash ROM An 8 KB x 8 NVRAM (upgradable to 32 KB x8) stores RAID system configuration information. The MegaRAID Express firmware is stored in flash ROM for easy upgrade.

SCSI Connectors MegaRAID Express does not have any SCSI connectors. MegaRAID Express is designed to work with the SCSI connectors on the motherboard.

Documentation

The MegaRAID Express documentation set includes:

- the *MegaRAID Express 000 Hardware Guide,*
- the MegaRAID *Express 000 Software Configuration Guide, and*
- the *MegaRAID Operating System Drivers Guide*.

Using MegaRAID Express Manuals This manual contains the RAID overview, RAID planning, and RAID system configuration information you will need first. Read the *MegaRAID Express 000 Hardware Guide* first.

MegaRAID Express 000 Configuration Software Guide This manual contains information about MegaRAID Express software utilities. The software configuration utilities configure and modify RAID systems. The software utilities include:

- MegaRAID Express Configuration utility,
- MegaRAID Express Manager, and
- Power Console.

MegaRAID Operating System Drivers Guide This manual provides detailed information about installing the MegaRAID Express operating system drivers.

MegaRAID Express Block Diagram

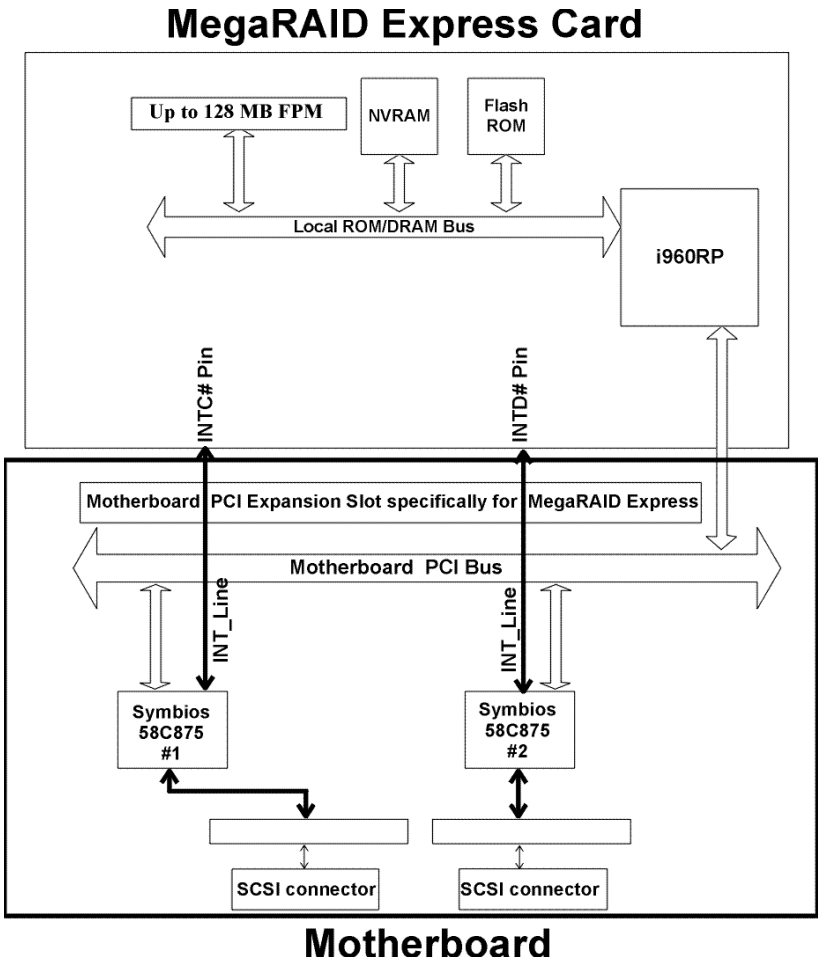

2 Introduction to RAID

RAID (Redundant Array of Independent Disks) is an array of multiple independent hard disk drives that provide high performance and fault tolerance. A RAID disk subsystem improves I/O performance over a computer using only a single drive. The RAID array appears to the host computer as a single storage unit or as multiple logical units. I/O is expedited because several disks can be accessed simultaneously. RAID systems improve data storage reliability and fault tolerance compared to single-drive computers. Data loss because of a disk drive failure can be prevented by reconstructing missing data from the remaining data and parity drives.

RAID Benefits

RAID has gained popularity because it: improves I/O performance, and increases storage subsystem reliability. RAID provides data security through fault tolerance and redundant data storage. The MegaRAID Express management software configures and monitors RAID disk arrays.

Improved I/O Although disk drive capabilities have improved drastically, actual performance has improved only three to four times in the last decade. Computing performance has improved over 50 times during the same time period.

Increased Reliability The electromechanical components of a disk subsystem operate more slowly, require more power, and generate more noise and vibration than electronic devices. These factors reduce the reliability of data stored on disks.

In This Chapter

The following topics are discussed:

Major Topic	Subtopic
Host-based solution	
RAID overview	
	Consistency check Fault tolerance Disk striping Disk spanning Disk mirroring Parity Hot spares Disk rebuilds Logical drives Hot swap SCSI drive states Logical drive states Disk array types

MegaRAID Express is a Host-Based RAID Solution

RAID products are either:

- host-based or
- SCSI-to-SCSI.

The MegaRAID Express controller is a host-based RAID solution. MegaRAID Express is a PCI adapter card that is installed in the RAID Upgrade PCI expansion slot in a host system.

Host-Based

A host-based RAID product puts all of the RAID intelligence on an adapter card that is installed in a network server. A host-based RAID product provides the best performance. MegaRAID Express is part of the file server, so it can transmit data directly across the computer's buses at data transfer speeds up to 132 MBs. The actual data transfer speed is determined by the number and type of SCSI channels and is usually between 20 and 40 MBs.

Host-based solutions must provide operating system-specific drivers.

SCSI-to-SCSI

A SCSI-to-SCSI RAID product puts the RAID intelligence inside the RAID chassis and uses a plain SCSI Host Adapter installed in the network server. The data transfer rate is limited to the bandwidth of the SCSI channels. A SCSI-to-SCSI RAID product that has two wide SCSI channels that operate at speeds up to 40 MBs must squeeze the data into a single wide SCSI (20 MBs) channel back to the Host computer.

In SCSI-to-SCSI RAID products, the hard drive subsystem uses only a single SCSI ID, which allows you to connect multiple drive subsystems to a single SCSI controller.

RAID Overview

RAID (Redundant Array of Independent Disks) is a collection of specifications that describe a system for ensuring the reliability and stability of data stored on large disk subsystems. A RAID system can be implemented in a number of different versions (or RAID Levels). The standard RAID levels are 0, 1, 3, and 5. MegaRAID Express supports all standard RAID levels and RAID levels 10, 30, and 50, special RAID versions supported by MegaRAID Express.

Consistency Check

In RAID, check consistency verifies the correctness of redundant data in an array. For example, in a system with dedicated parity, checking consistency means computing the parity of the data drives and comparing the results to the contents of the dedicated parity drive.

Fault Tolerance

Fault tolerance is achieved through cooling fans, power supplies, and the ability to hot swap drives. MegaRAID Express provides hot swapping through the hot spare feature. A hot spare drive is an unused online available drive that MegaRAID Express instantly plugs into the system when an active drive fails.

After the hot spare is automatically moved into the RAID subsystem, the failed drive is automatically rebuilt. The RAID disk array continues to handle request while the rebuild occurs.

Disk Striping

Disk striping writes data across multiple disk drives instead of just one disk drive. Disk striping involves partitioning each drive storage space into stripes that can vary in size from 2 KB to 128 KB. These stripes are interleaved in a repeated sequential manner. The combined storage space is composed of stripes from each drive. MegaRAID Express supports stripe sizes of 2 KB, 4 KB, 8 KB, 16 KB, 32 KB, 64 KB, or 128 KB.

For example, in a four-disk system using only disk striping (as in RAID level 0), segment 1 is written to disk 1, segment 2 is written to disk 2, and so on. Disk striping enhances performance because multiple drives are accessed simultaneously; but disk striping does not provide data redundancy.

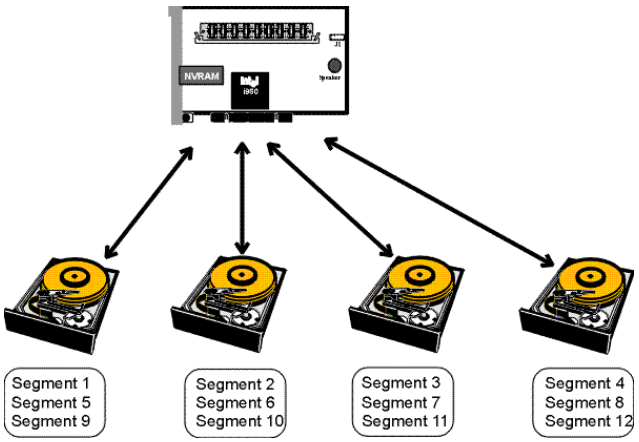

Stripe Width

Stripe width is the number of disks involved in an array where striping is implemented. For example, a four-disk array with disk striping has a stripe width of four.

Stripe Size

The stripe size is the length of the interleaved data segments that MegaRAID Express writes across multiple drives. MegaRAID Express supports stripe sizes of 2 KB, 4 KB, 8 KB, 16 KB, 32 KB, 64 KB, or 128 KB.

Disk Spanning

Disk spanning allows multiple disk drives to function like one big drive. Spanning overcomes lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources. For example, four 400 MB disk drives can be combined to appear to the operating system as one single 1600 MB drive.

Spanning alone does not provide reliability or performance enhancements. Spanned logical drives must have the same stripe size and must be contiguous. In the following graphic, RAID 1 array is turned into a RAID 10 array.

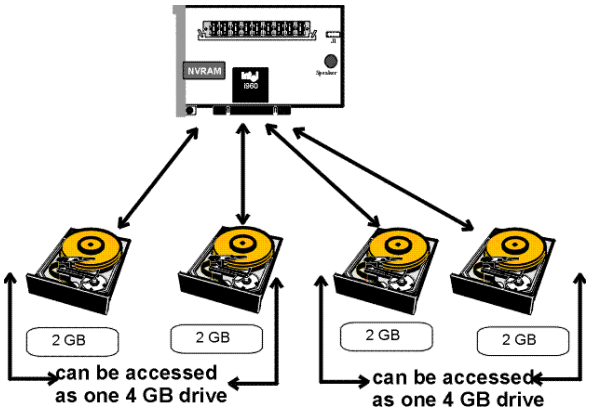

Spanning for RAID 10, RAID 30, or RAID 50

Level	Description
10	Configure RAID 10 by spanning two contiguous RAID 1 logical drives. The RAID 1 logical drives must have the same stripe size.
30	Configure RAID 30 by spanning two contiguous RAID 3 logical drives. The RAID 3 logical drives must have the same stripe size.
50	Configure RAID 50 by spanning two contiguous RAID 5 logical drives. The RAID 5 logical drives must have the same stripe size.
Note:	Spanning two contiguous RAID 0 logical drives does not produce a new RAID level or add fault tolerance. It does increase the size of the logical volume and improves performance by doubling the number of spindles.

Disk Mirroring

With mirroring (used in RAID 1), data written to one disk drive is simultaneously written to another disk drive. If one disk drive fails, the contents of the other disk drive can be used to run the system and reconstruct the failed drive. The primary advantage of disk mirroring is that it provides 100% data redundancy. Since the contents of the disk drive are completely written to a second drive, it does not matter if one of the drives fails. Both drives contain the same data at all times. Either drive can act as the operational drive.

Disk mirroring provides 100% redundancy, but is expensive because each drive in the system must be duplicated.

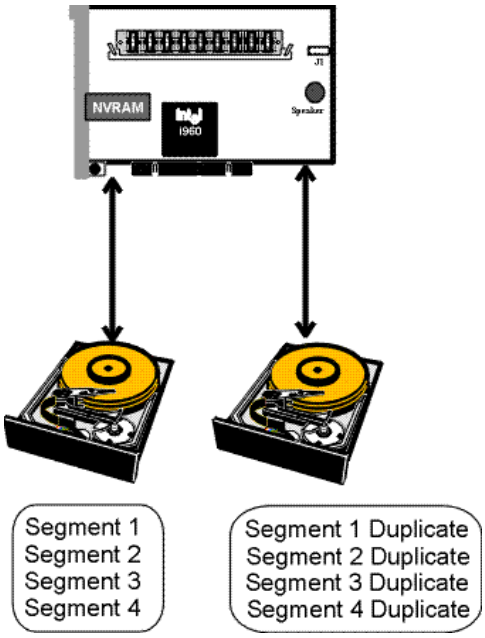

Parity

Parity generates a set of redundancy data from two or more parent data sets. The redundancy data can be used to reconstruct one of the parent data sets. Parity data does not fully duplicate the parent data sets. In RAID, this method is applied to entire drives or stripes across all disk drives in an array. The types of parity are:

Type	Description
Dedicated Parity	The parity of the data on two or more disk drives is stored on an additional disk.
Distributed Parity	The parity data is distributed across all drives in the system.

If a single disk drive fails, it can be rebuilt from the parity and the data on the remaining drives.

RAID level 3 combines dedicated parity with disk striping. The parity disk in RAID 3 is the last logical drive in a RAID set.

RAID level 5 combines distributed parity with disk striping. Parity provides redundancy for one drive failure without duplicating the contents of entire disk drives, but parity generation can slow the write process. A dedicated parity scheme during normal read/write operations is shown below:

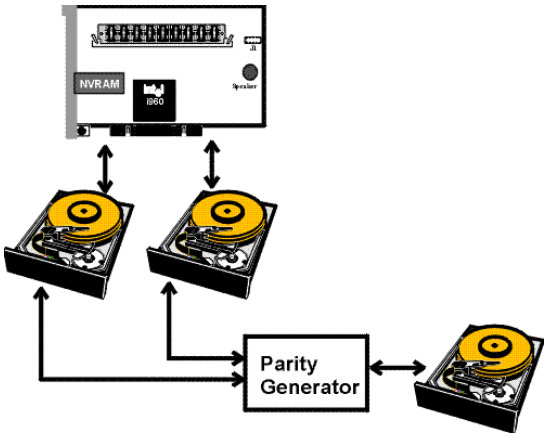

Hot Spares

A hot spare is an extra, unused disk drive that is part of the disk subsystem. It is usually in standby mode, ready for service if a drive fails. Hot spares permit you to replace failed drives without system shutdown or user intervention.

MegaRAID Express implements automatic and transparent rebuilds using hot spare drives, providing a high degree of fault tolerance and zero downtime. The MegaRAID Express RAID Management software allows you to specify physical drives as hot spares. When a hot spare is needed, the MegaRAID Express controller assigns the hot spare that has a capacity closest to and at least as great as that of the failed drive to take the place of the failed drive.

Important

Hot spares are only employed in arrays with redundancy, for example, RAID levels 1, 3, 5, 10, 30, and 50.

A hot spare connected to a specific MegaRAID Express controller can only be used to rebuild a drive that is connected to the same controller.

Disk Rebuild

You rebuild a disk drive by recreating the data that had been stored on the drive before the drive failed.

Rebuilding can be done only in arrays with data redundancy such as RAID level 1, 3, 5, 10, 30, and 50.

Standby (warm spare) rebuild is employed in a mirrored (RAID 1) system. If a disk drive fails, an identical drive is immediately available. The primary data source disk drive is the original disk drive.

A hot spare can be used to rebuild disk drives in RAID 1, 3, 5, 10, 30, or 50 systems. If a hot spare is not available, the failed disk drive must be replaced with a new disk drive so that the data on the failed drive can be rebuilt.

The MegaRAID Express controller automatically and transparently rebuilds failed drives with user-definable rebuild rates. If a hot spare is available, the rebuild starts automatically when a drive fails. MegaRAID Express automatically restarts the system and the rebuild if the system goes down during a rebuild.

Rebuild Rate

The rebuild rate is the fraction of the compute cycles dedicated to rebuilding failed drives. A rebuild rate of 100 percent means the system is totally dedicated to rebuilding the failed drive.

The MegaRAID Express rebuild rate can be configured between 0% and 100%. At 0%, the rebuild is only done if the system is not doing anything else. At 100%, the rebuild has a higher priority than any other system activity.

Physical Array

A RAID array is a collection of physical disk drives governed by the RAID management software. A RAID array appears to the host computer as one or more logical drives.

Logical Drive

A logical drive is a partition in a physical array of disks that is made up of contiguous data segments on the physical disks. A logical drive can consist of any of the following:

- an entire physical array,
- more than one entire physical array,
- a part of an array,
- parts of more than one array, or
- a combination of any two of the above conditions.

Hot Swap

A hot swap is the manual replacement of a defective physical disk unit while the computer is still running. When a new drive has been installed, you must issue a command to rebuild the drive.

SCSI Drive States

A SCSI disk drive can be in one of these states:

State	Description
Online (ONLIN)	The drive is functioning normally and is a part of a configured logical drive.
Ready (READY)	The drive is functioning normally but is not part of a configured logical drive and is not designated as a hot spare.
Hot Spare (HOTSP)	The drive is powered up and ready for use as a spare in case an online drive fails.
Fail (FAIL)	A fault has occurred in the drive placing it out of service.
Rebuild (REB)	The drive is being rebuilt with data from a failed drive.

Logical Drive States

State	Description
Optimal	The drive operating condition is good. All configured drives are online
Degraded	The drive operating condition is not optimal. One of the configured drives has failed or is offline.
Failed	The drive has failed.
Offline	The drive is not available to MegaRAID Express.

Disk Array Types

The RAID disk array types are:

Type	Description
Software-Based	The array is managed by software running in a host computer using the host CPU bandwidth. The disadvantages associated with this method are the load on the host CPU and the need for different software for each operating system.
SCSI to SCSI	The array controller resides outside of the host computer and communicates with the host through a SCSI adapter in the host. The array management software runs in the controller. It is transparent to the host and independent of the host operating system. The disadvantage is the limited data transfer rate of the SCSI channel between the SCSI adapter and the array controller.
Bus-Based	The array controller resides on the bus (for example, a PCI or EISA bus) in the host computer and has its own CPU to generate the parity and handle other RAID functions. A bus-based controller can transfer data at the speed of the host bus (PCI, ISA, EISA, VL-Bus) but is limited to the bus it is designed for. MegaRAID Express resides on a PCI bus, which can handle data transfer at up to 132 MBs. With MegaRAID Express, each channel can handle data transfer rates up to 40 MBs per SCSI channel.

3 RAID Levels

There are six official RAID levels (RAID 0 through RAID 5). MegaRAID Express supports RAID levels 0, 1, 3, and 5. American Megatrends has designed three additional RAID levels (10, 30, and 50) that provide additional benefits. The RAID levels that MegaRAID Express supports are:

RAID Level	Type
0	Standard
1	Standard
3	Standard
5	Standard
10	MegaRAID Express only
30	MegaRAID Express only
50	MegaRAID Express only

Select RAID Level To ensure the best performance, you should select the optimal RAID level when you create a system drive. The optimal RAID level for your disk array depends on a number of factors:

- the number of drives in the disk array,
- the capacity of the drives in the array,
- the need for data redundancy, and
- the disk performance requirements.

Selecting a RAID Level The factors you need to consider when selecting a RAID level are listed on the next screen

Selecting a RAID Level

Level	Description and Use	Pros	Cons	Max. Drives	Fault Tolerant
0	Data divided in blocks and distributed sequentially (pure striping). Use for non-critical data that requires high performance.	High data throughput for large files	No fault tolerance. All data lost if any drive fails.	One to 32	No
1	Data duplicated on another disk (mirroring). Use for read-intensive fault-tolerant systems	100% data redundancy	Doubles disk space. Reduced performance during rebuilds.	2, 4, 6, or 8	Yes
3	Disk striping with a dedicated parity drive. Use for non-interactive apps that process large files sequentially.	Achieves data redundancy at low cost	Performance not as good as RAID 1	Three to eight	Yes
5	Disk striping and parity data across all drives. Use for high read volume but low write volume, such as transaction processing.	Achieves data redundancy at low cost	Performance not as good as RAID 1	Three to eight	Yes
10	Data striping and mirrored drives.	High data transfers, complete redundancy	More complicated	4, 6, or 8	Yes
30	Disk striping with a dedicated parity drive.	High data transfers, redundancy	More complicated	Six to 32	Yes

Level	Description and Use	Pros	Cons	Max. Drives	Fault Tolerant
50	Disk striping and parity data across all drives.	High data transfers, redundancy	More complicated	Six to 32	Yes

RAID 0

RAID 0 provides disk striping across all drives in the RAID subsystem. RAID 0 does not provide any data redundancy, but does offer the best performance of any RAID level. RAID 0 breaks up data into smaller blocks and then writes a block to each drive in the array. The size of each block is determined by the stripe size parameter, set during the creation of the RAID set. RAID 0 offers high bandwidth. By breaking up a large file into smaller blocks, MegaRAID Express can use multiple SCSI channels and drives to read or write the file faster. RAID 0 involves no parity calculations to complicate the write operation. This makes RAID 0 ideal for applications that require high bandwidth but do not require fault tolerance.

Uses RAID 0 provides high data throughput, especially for large files. Any environment hat does not require fault tolerance.

Strong Points Provides increased data throughput for large files. No capacity loss penalty for parity.

Weak Points Does not provide fault tolerance. All data lost if any drive fails.

Drives One to 32

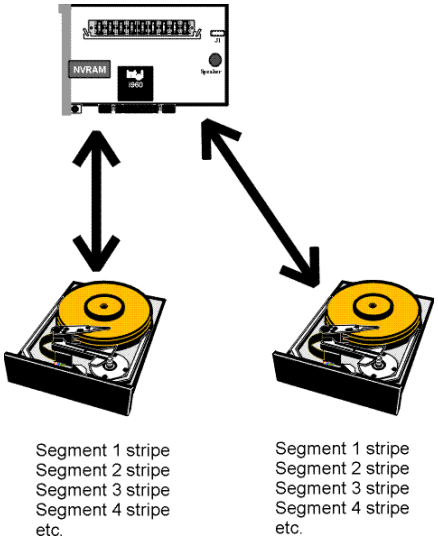

RAID 1

In RAID 1, MegaRAID Express duplicates all data from one drive to a second drive. RAID 1 provides complete data redundancy, but at the cost of doubling the required data storage capacity.

Uses Use RAID 1 for small databases or any other environment that requires fault tolerance but small capacity.

Strong Points RAID 1 provides complete data redundancy. RAID 1 is ideal for any application that requires fault tolerance and minimal capacity.

Weak Points RAID 1 requires twice as many disk drives. Performance is impaired during drive rebuilds.

Drives 2, 4, 6, or 8 drives.

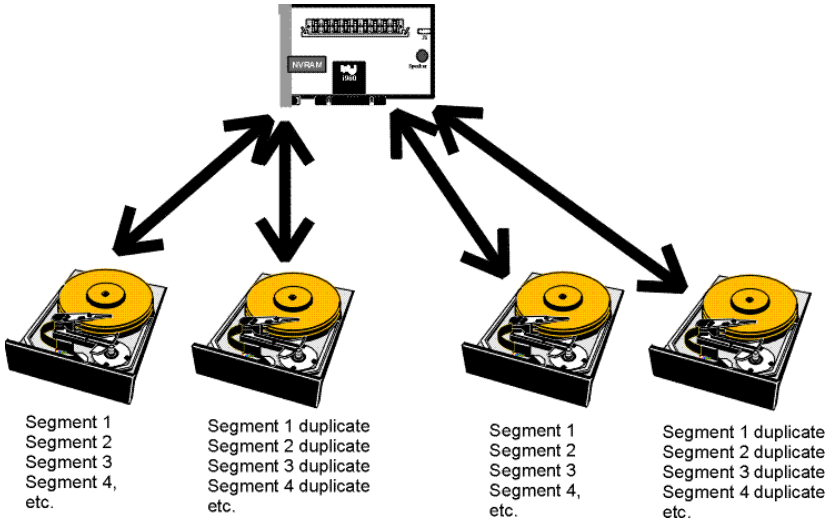

RAID 3

RAID 3 provides disk striping and complete data redundancy though a dedicated parity drive. The stripe size must be 64 KB if RAID 3 is used. RAID 3 handles data at the block level, not the byte level, so it is ideal for networks that often handle very large files, such as graphic images. RAID 3 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks, and then writes the blocks to all but one drive in the array. The parity data created during the exclusive-or is then written to the last drive in the array. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set. If a single drive fails, a RAID 3 array continues to operate in degraded mode. If the failed drive is a data drive, writes will continue as normal, except no data is written to the failed drive. Reads reconstruct the data on the failed drive by performing an exclusive-or operation on the remaining data in the stripe and the parity for that stripe. If the failed drive is a parity drive, writes will occur as normal, except no parity is written. Reads retrieve data from the disks.

Uses	Best suited for applications such as graphics, imaging, or video that call for reading and writing huge, sequential blocks of data.
Strong Points	Provides data redundancy and high data transfer rates.
Weak Points	The dedicated parity disk is a bottleneck with random I/O.
Drives	Three to eight

Cont’d

RAID 3, Continued

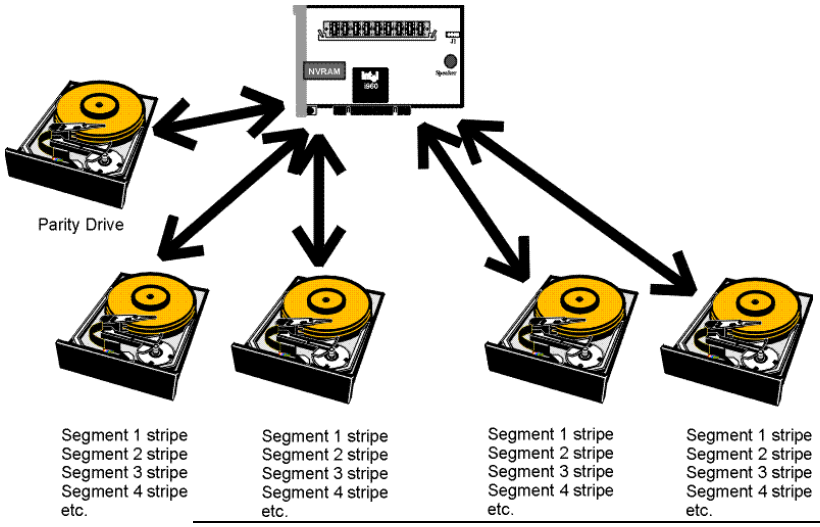

RAID 5 vs RAID 3 You may find that RAID 5 is preferable to RAID 3 even for applications characterized by sequential reads and writes, because MegaRAID Express has very robust caching algorithms and hardware based exclusive-or assist.

The benefits of RAID 3 disappear if there are many small I/O operations scattered randomly and widely across the disks in the logical drive. The RAID 3 fixed parity disk becomes a bottleneck in such applications. For example: The host attempts to make two small writes and the writes are widely scattered, involving two different stripes and different disk drives. Ideally both writes should take place at the same time. But this is not possible in RAID 3, since the writes must take turns accessing the fixed parity drive. For this reason, RAID 5 is the clear choice in this scenario.

RAID 5

RAID 5 includes disk striping at the byte level and parity. In RAID 5, the parity information is written to several drives. RAID 5 is best suited for networks that perform a lot of small I/O transactions simultaneously.

RAID 5 addresses the bottleneck issue for random I/O operations. Since each drive contains both data and parity numerous writes can take place concurrently. In addition, robust caching algorithms and hardware based exclusive-or assist make RAID 5 performance exceptional in many different environments.

Uses RAID 5 provides high data throughput, especially for large files. Use RAID 5 for transaction processing applications because each drive can read and write independently. If a drive fails, MegaRAID Express uses the parity drive to recreate all missing information. Use also for office automation and online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates.

Strong Points Provides data redundancy and good performance in most environments

Weak Points Disk drive performance will be reduced if a drive is being rebuilt. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.

Drives Three to eight

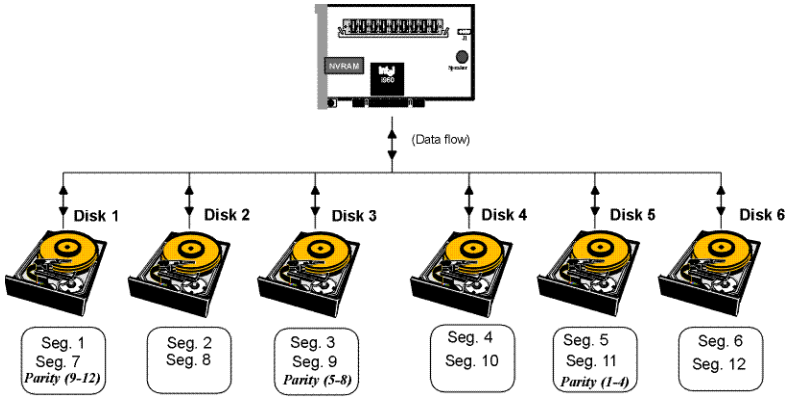

Parity is distributed across all drives in array.

RAID 10

RAID 10 is a combination of RAID 0 and RAID 1. RAID 10 has mirrored drives. RAID 10 breaks up data into smaller blocks, and then stripes the blocks of data to each RAID 1 raid set. Each RAID 1 raid set then duplicates its data to its other drive. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set. RAID 10 can sustain one to four drive failures while maintaining data integrity if each failed disk is in a different RAID 1 array.

Uses RAID 10 works best for data storage that must have 100% redundancy of mirrored arrays and that also needs the enhanced I/O performance of RAID 0 (striped arrays). RAID 10 works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate to medium capacity.

Strong Points RAID 10 provides both high data transfer rates and complete data redundancy.

Weak Points RAID 10 requires twice as many drives as all other RAID levels except RAID 1.

Drives 2n, where n is greater than 1.

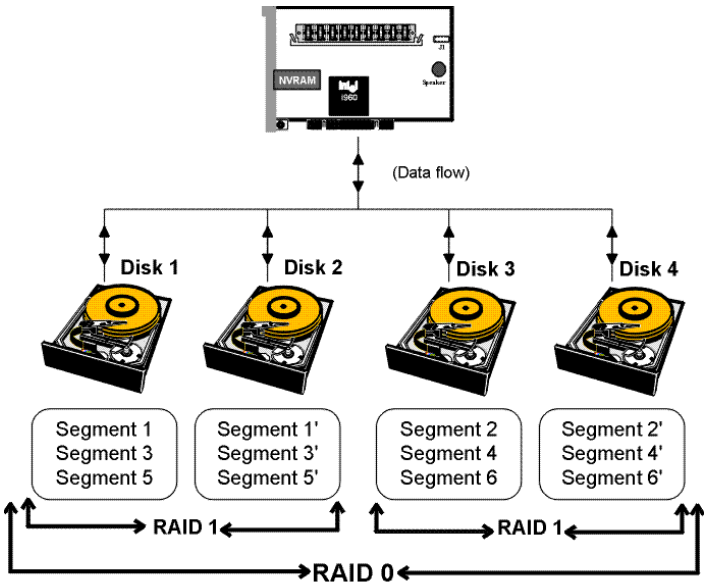

RAID 30

RAID 30 is a combination of RAID 0 and RAID 3. RAID 30 provides high data transfer speeds and high data reliability. RAID 30 is best implemented on two RAID 3 disk arrays with data striped across both disk arrays. RAID 30 breaks up data into smaller blocks, and then stripes the blocks of data to each RAID 3 raid set. RAID 3 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks, and then writes the blocks to all but one drive in the array. The parity data created during the exclusive-or is then written to the last drive in each RAID 3 array. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.

RAID 30 can sustain one to four drive failures while maintaining data integrity if each failed disk is in a different RAID 3 array.

Uses Use RAID 30 for sequentially written and read data, pre-press and video on demand that requires a higher degree of fault tolerance and medium to large capacity.

Strong Points Provides data reliability and high data transfer rates.

Weak Points Requires 2 – 4 times as many parity drives as RAID 3.

Drives Six to 32

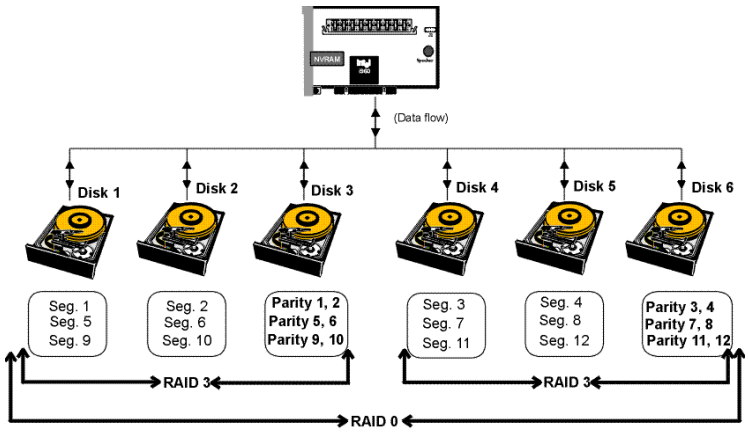

RAID 50

RAID 50 provides the features of both RAID 0 and RAID 5. RAID 50 includes both parity and disk striping across multiple drives. RAID 50 is best implemented on two RAID 5 disk arrays with data striped across both disk arrays. RAID 50 breaks up data into smaller blocks, and then stripes the blocks of data to each RAID 5 raid set. RAID 5 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks, and then writes the blocks of data and parity to each drive in the array. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.

RAID 50 can sustain one to four drive failures while maintaining data integrity if each failed disk is in a different RAID 5 array.

Uses RAID 50 works best when used with data that requires high reliability, high request rates, and high data transfer and medium to large capacity.

Strong Points RAID 50 provides high data throughput, data redundancy, and very good performance.

Weak Points Requires 2 to 4 times as many parity drives as RAID 5.

Drives Six to 32

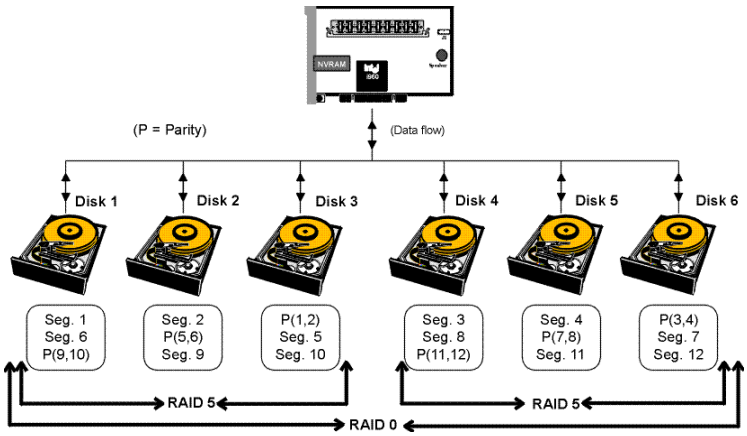

4 MegaRAID Express Features

MegaRAID Express is a family of high performance intelligent PCI-to-SCSI host adapters with RAID control capabilities. MegaRAID Express works with a motherboard that has one or two SCSI channels that support Ultra and Wide SCSI at data transfer rates of up to 40 MBs per SCSI channel. Each SCSI channel supports up to 15 Wide devices and up to seven non-Wide devices.

In This Chapter Topics described in this chapter include:

- new features,
- configuration features,
- hardware architecture features,
- array performance features,
- RAID management features,
- fault tolerance features,
- utility programs, and
- software drivers.

Hardware Requirements

MegaRAID Express can be installed in an IBM AT®-compatible or EISA computer with a motherboard that has a RAID Expansion PCI slot, PCI interrupt routing, onboard SCSI channels, and 5 volt PCI expansion slots. The computer must support PCI version 2.1. The computer should have an Intel Pentium, Pentium Pro, or more powerful CPU, a floppy drive, a color monitor and VGA adapter, a keyboard, and a mouse.

Configuration Features

Specification	Feature
RAID Levels	0, 1, 3, 5, 10, 30, and 50.
SCSI Channels	1 or 2 on motherboard
Maximum number of drives per channel	15
Array interface to host	PCI 2.1
Drive interface	Fast and Wide Ultra SE (SCSI 3)
Upgradable cache size	8 MB, 16 MB, 32 MB, 64 MB, or 128 MB
Cache Function	Write-through, write-back, ARA, NRA, RA
Multiple logical drives/arrays per controller	Up to 8 logical drives per controller
Maximum number of MegaRAID Express controllers per system	1
Online capacity expansion (with FlexRAID only)	Yes
Pool hot spare	Yes
Flashable firmware	Yes
Hot swap devices supported	Yes
Non-disk devices supported	Yes
Mixed capacity hard disk drives	Yes
Number of 16-bit internal connectors	None
Number of 16-bit external connectors	None
Support for hard disk drives with capacities of more than 8 GB.	Yes
Online RAID level migration (with FlexRAID only)	Yes
No reboot necessary after expansion	Yes
More than 200 Qtags per physical drive	Yes
User-specified rebuild rate	Yes

Hardware Architecture Features

The MegaRAID Express hardware architecture features include:

Specification	Feature
Processor	Intel i960RP 33 MHz
SCSI Controller	SymBIOS Logic 53C875 SCSI controllers on the motherboard
AMI ASIC	None
Size of Flash ROM	1 MB
Amount of NVRAM	32 KB
Direct I/O	Yes
SCSI bus termination	Active, single-ended
Double-sided SIMM	Yes
Direct I/O bandwidth	132 MBs

Array Performance Features

The MegaRAID Express array performance features include:

Specification	Feature
Host data transfer rate	132 MBs
Drive data transfer rate	40 MBs
Maximum Scatter/Gathers	26 elements
Maximum size of I/O requests	6.4 MB in 64 KB stripes
Maximum Queue Tags per drive	211
Stripe Sizes	2 KB, 4 KB, 8 KB, 16 KB, 32 KB, 64 KB, or 128 KB
Maximum number of concurrent commands	255
Support for multiple initiators	Yes

RAID Management Features

Specification	Feature
Support for SNMP	Yes
Performance Monitor provided	Yes
Remote control and monitoring	Yes
Event broadcast and event alert	Yes
Drive roaming	Yes
Support for concurrent multiple stripe sizes	Yes
Windows NT and NetWare server support via GUI client utility	Yes
OS/2, and UnixWare server support via GUI client utility	Yes
DMI support	Yes

Fault Tolerance Features

Specification	Feature
Drive failure detection	Automatic
Drive rebuild using hot spares	Automatic

Software Utilities

The MegaRAID Express software utility features include:

Specification	Feature
Graphical user interface	Yes
Management utility	Yes
Bootup configuration via MegaRAID Express Manager	Yes
Online Read, Write, and cache policy switching	Yes

Operating System Software Drivers

Operating System Drivers MegaRAID Express includes a DOS software configuration utility and drivers for most operating systems. See the *MegaRAID Operating System Drivers Guide* for additional information.

The DOS drivers for MegaRAID Express are contained in the firmware on MegaRAID Express except the DOS ASPI and CD-ROM drivers. Call American Megatrends technical support at 770-246-8600 for information about drivers for other operating systems or visit the AMI web site at www.ami.com.

MegaRAID Express Specifications

Parameter	Specification
Card Size	6" x 3.875"
Bus Type	PCI 2.1
Bus Data Transfer Rate	Up to 132 MBs
BIOS	MegaRAID Express BIOS
Cache Configuration	8, 16, 32, 64, or 128 MB using a 70 ns × 36 Fast Page Mode 72-pin SIMM
Operating Voltage	5.00 V ± 0.25 V
SCSI Data Transfer Rate	Up 40 MBs
SCSI Bus	Single-ended
SCSI Termination	Active
Termination Disable	Automatic through cable detection
Devices per SCSI Channel	Up to 15 wide or seven non-wide SCSI devices. Up to 6 non-disk SCSI drives per MegaRAID Express controller.
SCSI Device Types	Synchronous or Asynchronous. Disk and non-disk.
RAID Levels Supported	0, 1, 3, 5,10, 30, and 50

Specific Features

SMART Technology The MegaRAID Express Self Monitoring Analysis and Reporting Technology (SMART) detects up to 70% of all predictable drive failures. SMART monitors the internal performance of all motors, heads, and drive electronics. You can recover from drive failures through remapping and online physical drive migration.

Configuration on Disk **Configuration** on Disk (drive roaming) saves configuration information both in NVRAM on MegaRAID Express and on the disk drives connected to MegaRAID Express. If MegaRAID Express is replaced, the new MegaRAID Express controller can detect the actual RAID configuration, maintaining the integrity of the data on each drive, even if the drives have changed channel and/or target ID.

CPU The MegaRAID Express controller uses the 32-bit Intel i960RP running at 33 MHz. This processor directs all functions of the controller including command processing, PCI and SCSI bus transfers, RAID processing, drive rebuilding, cache management, and error recovery.

Cache Memory Cache memory resides in a memory bank requiring a 4 MB x 36, 16 MB x 36, or 32 MB x 36 72-pin 70 ns Fast Page Mode SIMM. Memory is interleaved. Possible configurations are 8, 16, 32, 64, or 128 MB. The MegaRAID Express controller supports write-through or write-back caching, selectable for each logical drive. To improve performance in sequential disk accesses, the MegaRAID Express controller uses read-ahead caching by default. You can disable read-ahead caching.

BIOS The BIOS resides on a 1 MB × 8 flash ROM for easy upgrade. The BIOS supports INT 13h calls to boot DOS without special software or device drivers. The BIOS provides an extensive setup utility that can be accessed by pressing <Ctrl> <M> at BIOS initialization. The MegaRAID Express Configuration utility is described in the *MegaRAID Express 000 Configuration Software Guide.*

Cont'd

Specific Features, Continued

Onboard Speaker MegaRAID Express has an onboard tone generator for audible warnings when system errors occur. Audible warnings can be generated through this speaker. The audible warnings are listed on page 85.

SCSI Bus MegaRAID Express uses up to two SCSI channels on the motherboard. The motherboard SCSI channels must use American Megatrends IRQ routing. Each channel supports up to 15 wide or seven non-wide SCSI devices at speeds up to 40 MBs per SCSI channel. MegaRAID Express supports up to six non-disk devices per controller.

SCSI Firmware The firmware handles all RAID and SCSI command processing and also supports:

Feature	Description
Disconnect/ Reconnect	Optimizes SCSI Bus seek.
Tagged Command Queuing	Multiple tags to improve random access
Scatter/Gather	Multiple address/count pairs
Multi-threading	Up to 255 simultaneous commands with elevator sorting and concatenation of requests per SCSI channel
Stripe Size	Variable for all logical drives: 2 KB, 4 KB, 8 KB, 16 KB, 32 KB, 64 KB, or 128 KB.
Rebuild	Multiple rebuilds and consistency checks with user-definable priority.

RAID Management These utilities manage and configure the RAID system and MegaRAID Express, create and manage multiple disk arrays, control and monitor multiple RAID servers, provide error statistics logging, and provide online maintenance facilities. These utilities include:

- MegaRAID Express Configuration utility,
- Power Console Plus, and
- MegaRAID Express Manager

Cont'd

Specific Features, Continued

MegaRAID Configuration Utility The MegaRAID Express Configuration utility configures and maintains RAID arrays, formats disk drives, and manages the RAID system. It is independent of any operating system. See the *MegaRAID Express 000 Configuration Software Guide* for additional information.

Power Console Plus Power Console Plus runs in Windows NT. It configures, monitors, and maintains multiple RAID servers from any network node or a remote location.

MegaRAID Express Manager MegaRAID Express Manager is a character-based utility that works in DOS, SCO UnixWare, OS/2 2.x and OS/2 Warp, Novell NetWare.

Fault-Tolerance Features The fault-tolerance features are:

- automatic failed drive detection,
- automatic failed drive rebuild with no user intervention required,
- hot swap manual replacement without bringing the system down, and
- cache memory.

Detect Failed Drive The MegaRAID Express firmware automatically detects and rebuilds failed drives. This can be done transparently with hot spares.

Hot Swap MegaRAID Express supports the manual replacement of a disk unit in the RAID subsystem without system shutdown.

Compatibility

MegaRAID Express compatibility issues include:

- server management,
- SCSI device compatibility, and
- software compatibility

Server Management As an SNMP agent, MegaRAID Express supports all SNMP managers.

SCSI Device Compatibility MegaRAID Express supports SCSI hard disk drives, CD-ROM drives, tape drives, optical drives, DAT drives and other SCSI peripheral devices.

Software All SCSI backup and utility software should work with MegaRAID Express. Software that has been tested and approved for use with MegaRAID Express includes Cheyenne®, CorelSCSI®, Arcserve®, and Novaback®. This software is not provided with MegaRAID Express.

Summary

MegaRAID Express features were discussed in this chapter.

Hardware installation is discussed in Chapter 6.

5 Configuring MegaRAID Express

Configuring SCSI Physical Drives

SCSI Channels Physical SCSI drives must be organized into logical drives. The arrays and logical drives that you construct must be able to support the RAID level that you select.

> ***Warning***
>
> The BIOS that controls the motherboard SCSI channels must be disabled to permit MegaRAID Express to control these SCSI channels correctly.

MegaRAID Express adapter supports one or two SCSI channels that are on the motherboard..

Distributing Drives If your MegaRAID Express adapter supports more than one SCSI channel, distribute the disk drives among available channels for optimal performance. It is best to stripe across channels instead of down channels. Performance is most affected for sequential reads and writes. MegaRAID Express supports SCSI CD-ROM drives, SCSI tape drives, and other SCSI devices as well as SCSI hard disk drives. For optimal performance, all non-disk SCSI devices should be attached to one SCSI channel.

Basic Configuration Rules Observe the following guidelines when connecting and configuring SCSI devices in a RAID array:

- attach non-disk SCSI devices to a single SCSI channel that does not have any disk drives,
- distribute the SCSI hard disk drives equally among all available SCSI channels except any SCSI channel that is being reserved for non-disk drives,
- you can place up to eight physical disk drives in an array,
- an array can contain SCSI devices that reside on an array on any channel,
- include all drives that have the same capacity to the same array,
- make sure any hot spare has a capacity that is at least as large as the largest drive that may be replaced by the hot spare, and
- when replacing a failed drive, make sure that the replacement drive has a capacity that is at least as large as the drive being replaced.

Current Configuration

SCSI ID	Device Description	Termination?
SCSI Channel 1		
0		
1		
2		
3		
4		
5		
6		
8		
9		
10		
11		
12		
13		
14		
15		
SCSI Channel 2		
0		
1		
2		
3		
4		
5		
6		
8		
9		
10		
11		
12		
13		
14		
15		

Logical Drive Configuration

Logical Drive	RAID Level	Stripe Size	Logical Drive Size	Cache Policy	Read Policy	Write Policy	# of Physical Drives
LD1							
LD2							
LD3							
LD4							
LD5							
LD6							
LD7							
LD8							

Cont’d

Physical Device Layout

	Channel 1	Channel 2
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		

	Channel 1	Channel 2
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		
Target ID		
Device Type		

	Channel 1	Channel 2
Logical Drive Number/ Drive Number		
Manufacturer/Model Number		
Firmware level		

Configuring Arrays

Organize the physical disk drives in arrays after the drives are connected to MegaRAID Express, formatted, and initialized. Each array can consist of one to eight physical disk drives.

MegaRAID Express supports up to eight arrays. The number of drives in an array determines the RAID levels that can be supported.

Arranging Arrays You must arrange the arrays to provide additional organization for the drive array. You must arrange arrays so that you can create system drives that can function as boot devices.

You can sequentially arrange arrays with an identical number of drives so that the drives in the group are spanned. Spanned drives can be treated as one large drive. Data can be striped across multiple arrays as one logical drive.

You can create spanned drives by using the MegaRAID Express Configuration utility or the MegaRAID Express Manager.

Creating Hot Spares Any drive that is present, formatted, and initialized but is not included in a array or logical drive is automatically designated as a hot spare. You can also designate drives as hot spares via MegaRAID Express Configuration utility , the MegaRAID Express Manager, or Power Console. See the *MegaRAID Express 000 Configuration Software Guide* for additional information.

Creating Logical Drives Logical drives are arrays or spanned arrays that are presented to the operating system. You must create one or more logical drives.

The logical drive capacity can include all or any portion of a array. The logical drive capacity can also be larger than an array by using spanning. MegaRAID Express supports up to 8 logical drives.

Configuration Strategies

The most important factors in RAID array configuration are: drive capacity, drive availability (fault tolerance), and drive performance. You cannot configure a logical drive that optimizes all three factors, but it is easy to choose a logical drive configuration that maximizes one factor at the expense of the other two factors, although needs are seldom that simple.

Maximize Capacity RAID 0 achieves maximum drive capacity, but does not provide data redundancy. Maximum drive capacity for each RAID level is shown below. OEM level firmware that can span up to 4 logical drives is assumed.

RAID Level	Description	Drives Required	Capacity
0	Striping without parity	1 – 32	(Number of disks) X capacity of smallest disk
1	Mirroring	2	(Capacity of smallest disk) X (1)
3	Striping with fixed parity drive	3 – 8	(Number of disks) X (capacity of smallest disk) - (capacity of 1 disk)
5	Striping with floating parity drive	3 – 8	(Number of disks) X (capacity of smallest disk) - (capacity of 1 disk)
10	Mirroring and Striping	4 – 8 (Must be a multiple of 2)	(Number of disks) X (capacity of smallest disk) / (2)
30	RAID 3 and Striping	6 – 32 (Must be a multiple of arrays)	(Number of disks) X (capacity of smallest disk) – (capacity of 1 disk X number of Arrays)
50	RAID 5 and Striping	6 – 32 (Must be a multiple of arrays)	(Number of disks) X (capacity of smallest disk) – (capacity of 1 disk X number of Arrays)

Cont’d

Maximizing Drive Availability You can maximize the availability of data on the physical disk drive in the logical array by maximizing the level of fault tolerance. The levels of fault tolerance provided by the RAID levels are:

RAID Level	Fault Tolerance Protection
0	No fault tolerance.
1	Disk mirroring, which provides 100% data redundancy.
3	100% protection through a dedicated parity drive.
5	100% protection through striping and parity. The data is striped and parity data is written across a number of physical disk drives.
10	100% protection through data mirroring.
30	100% protection through data striping. All data is striped across all drives in two or more arrays.
50	100% protection through data striping and parity. All data is striped and parity data is written across all drives in two or more arrays.

Maximizing Drive Performance You can configure an array for optimal performance. But optimal drive configuration for one type of application will probably not be optimal for any other application. A basic guideline of the performance for RAID drive arrays at each RAID level is:

RAID Level	Performance Characteristics
0	Excellent for all types of I/O activity, but provides no data security.
1	Provides data redundancy and good performance.
3	Provides data redundancy.
5	Provides data redundancy and good performance in most environments.
10	Provides data redundancy and excellent performance.
30	Provides data redundancy and good performance in most environments.
50	Provides data redundancy and very good performance.

Assigning RAID Levels

Only one RAID level can be assigned to each logical drive. The drives required per RAID level is:

RAID Level	Minimum Number of Physical Drives	Maximum Number of Physical Drives
0	One	32
1	Two	Two
3	Three	Eight
5	Three	Eight
10	four	Eight
30	Six	32
50	Six	32

Configuring Logical Drives

After you have installed the MegaRAID Express controller in the server and have attached all physical disk drives, perform the following actions to prepare a RAID disk array:

Step	Action
1	Optimize the MegaRAID Express controller options for your system. See Chapter 6 for additional information.
2	Perform a low-level format the SCSI drives that will be included in the array and the drives to be used for hot spares.
3	Press <Ctrl> <M> to run the MegaRAID Express Manager.
4	Define and configure one or more logical drives. Select Easy Configuration in MegaRAID Express Manager or select New Configuration to customize the RAID array.
5	Create and configure one or more system drives (logical drives). Select the RAID level, cache policy, read policy, and write policy.
6	Save the configuration.
7	Initialize the system drives. After initialization, you can install the operating system.

Optimizing Data Storage

Data Access Requirements Each type of data stored in the disk subsystem has a different frequency of read and write activity. If you know the data access requirements, you can more successfully determine a strategy for optimizing the disk subsystem capacity, availability, and performance.

Servers that support Video on Demand typically read the data often, but write data infrequently. Both the read and write operations tend to be long. Data stored on a general-purpose file server involves relatively short read and write operations with relatively small files.

Array Functions You must first define the major purpose of the disk array. Will this disk array increase the system storage capacity for general-purpose file and print servers? Does this disk array support any software system that must be available 24 hours per day? Will the information stored in this disk array contains large audio or video files that must be available on demand? Will this disk array contain data from an imaging system?

You must identify the purpose of the data to be stored in the disk subsystem before you can confidently choose a RAID level and a RAID configuration.

Planning the Array Configuration

Answer the following questions about this array:

Question	Answer
Number of MegaRAID Express SCSI channels	
Number of physical disk drives in the array	
Purpose of this array. Rank the following factors:	
Maximize drive capacity	
Maximize the safety of the data (fault tolerance)	
Maximize hard drive performance and throughput	
How many hot spares?	
Amount of cache memory installed on the MegaRAID Express	
Are all of the disk drives and the server that MegaRAID Express is installed in protected by a UPS?	

Using the Array Configuration Planner The following table lists the possible RAID levels, fault tolerance, and effective capacity for all possible drive configurations for an array consisting of one to eight drives.

The following table does not take into account any hot spare (standby) drives. You should always have a hot spare drive in case of drive failure.

RAID 1 and RAID 10 require 2, 4, 6, or 8 drives. RAID 30 and RAID 50 require at least 6 drives.

Array Configuration Planner

Number of Drives	Possible RAID Levels	Relative Performance	Fault Tolerance	Effective Capacity
1	None	Excellent	No	100%
1	RAID 0	Excellent	No	100%
2	None	Excellent	No	100%
2	RAID 0	Excellent	No	100%
2	RAID 1	Good	Yes	50%
3	None	Excellent	No	100%
3	RAID 0	Excellent	No	100%
3	RAID 3	Good	Yes	67%
3	RAID 5	Good	Yes	67%
4	None	Excellent	No	100%
4	RAID 0	Excellent	No	100%
4	RAID 1	Good	Yes	50%
4	RAID 3	Good	Yes	75%
4	RAID 5	Good	Yes	75%
4	RAID 10	Good	Yes	50%
5	None	Excellent	No	100%
5	RAID 0	Excellent	No	100%
5	RAID 3	Good	Yes	80%
5	RAID 5	Good	Yes	80%
6	None	Excellent	No	100%
6	RAID 0	Excellent	No	100%
6	RAID 1	Good	Yes	50%
6	RAID 3	Good	Yes	83%
6	RAID 5	Good	Yes	83%
6	RAID 10	Good	Yes	50%
6	RAID 30	Good	Yes	67%
6	RAID 50	Good	Yes	67%
7	None	Excellent	No	100%
7	RAID 0	Excellent	No	100%
7	RAID 3	Good	Yes	86%
7	RAID 5	Good	Yes	86%

6 Hardware Installation

Requirements You must have the following items before installing the MegaRAID Express controller in a server:

- a MegaRAID Express card,
- a motherboard that includes SCSI channels, SymBIOS SCSI controller chips, and supports American Megatrends IRQ routing,
- a host computer with an available PCI expansion slot,
- the MegaRAID Express Installation CD,
- the necessary SCSI cables and terminators (depends on the number and type of SCSI devices to be attached),
- an Uninterruptible Power Supply (UPS) for the entire system, and
- Fast SCSI 2 or Wide SCSI hard disk drives and other SCSI devices, as desired.

Optional Equipment You may also want to install SCSI cables that interconnect MegaRAID Express to external SCSI devices.

MegaRAID Express Card Layout

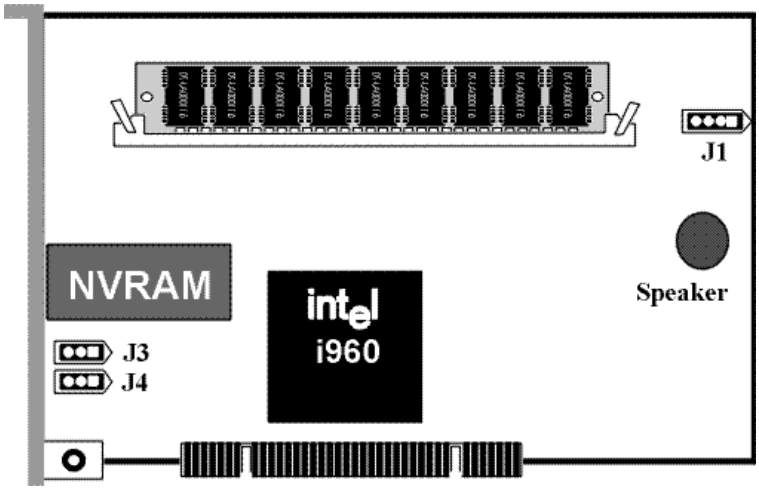

Checklist

Warning

The BIOS that controls the motherboard SCSI channels must be disabled to permit MegaRAID Express to control these SCSI channels correctly.

Perform the steps in the installation checklist:

Check	Step	Action
❑	1	Turn all power off to the server and all hard disk drives, enclosures, and system, components.
❑	2	Prepare the host system. See the host system technical documentation.
❑	3	Determine the SCSI ID and SCSI termination requirements.
❑	4	Install the cache memory on the MegaRAID Express card.
❑	5	Install the MegaRAID Express card in a PCI expansion slot that supports American Megatrends IRQ routing in the server and attach the SCSI cables and terminators as needed to the SCSI connectors in the server. MegaRAID Express has no SCSI connectors. Make sure Pin 1 on the cable matches Pin 1 on the connector. Make sure that the SCSI cables you use conform to all SCSI specifications.
❑	6	Perform a safety check. Make sure all cables are properly attached. Make sure the MegaRAID Express card is properly installed. Turn power on after completing the safety check.
❑	7	Install and configure the MegaRAID Express software utilities and drivers.
❑	8	Format the hard disk drives as needed.
❑	9	Configure system drives (logical drives).
❑	10	Initialize the logical drives.
❑	11	Install the network operating system drivers as needed.

Installation Steps

Step	Action
1	Unpack the MegaRAID Express controller and inspect for damage. Make sure all items are in the package. If damaged, call American Megatrends technical support at 770-246-8600.
2	Turn the computer off and remove the cover.
3	Make sure the motherboard jumper settings are correct.
4	Install cache memory on the MegaRAID Express card. 8 MB minimum cache memory is required.
5	Install the MegaRAID Express card in a PCI expansion slot that supports American Megatrends IRQ routing.
6	Set the target IDs for the SCSI devices.
7	Replace the computer cover and turn the power on. Be sure the SCSI devices are powered up before or at the same time as the host computer.
8	Run MegaRAID Express Configuration utility.
9	Install software drivers for the desired operating systems.

Step 1 Unpack

Unpack and install the hardware in a static-free environment. MegaRAID Express is packed inside an anti-static bag between two sponge sheets. Remove the controller card and inspect it for damage. If the card appears damaged or if any of items listed below are missing, contact Technical Support at 770-246-8600. MegaRAID Express is shipped with the following items that are on CD:

- the *MegaRAID Express 000 Configuration Software Guide,*
- the *MegaRAID Operating System Drivers Guide,*
- the *MegaRAID Express 000 Hardware Guide,*
- the software license agreement,
- the MegaRAID Express Configuration Utilities for DOS, and
- the warranty registration card.

Step 2 Power Down

Turn off the computer and remove the cover. Make sure the computer is turned off and disconnected from any networks before installing the controller card.

Step 3 Configure MegaRAID Express

Make sure the motherboard is configured correctly for MegaRAID Express. MegaRAID Express is a SCSI Controller. Each MegaRAID Express card you install will require an available PCI IRQ; make sure an IRQ is available for each controller you install.

J1 Termination Enable J1 is a 3-pin header that specifies hardware or software control of SCSI termination.

Type of SCSI Termination	J1 Setting
Software control of SCSI termination via drive detection.	Short Pins 1-2
Permanently disable all onboard SCSI termination.	Short Pins 2-3
Permanently enable all onboard SCSI termination.	OPEN

J3, J4 Motherboard SCSI Channel Select J3 and J4 are 3-pin jumpers that control the number of motherboard SCSI channels used by MegaRAID Express. SCSI channels not used by MegaRAID Express can be used as standalone SCSI channels. To allow MegaRAID Express to use both motherboard SCSI channels, short Pins 1-2 of J3 and J4.

Motherboard SCSI Channel Configuration	J3	J4
Both motherboard SCSI channels are used by MegaRAID Express. The MegaRAID Express PCI slot generates an IRQ on the INTA# pin on the MegaRAID Express PCI slot.	Short Pins 1-2 (Factory setting)	Short Pins 1-2 (Factory setting)
Motherboard SCSI channel 1 is used by MegaRAID Express. The MegaRAID Express PCI slot generates an IRQ on the INTA# pin on the MegaRAID Express PCI slot and passes the IRQ of the second SCSI channel back to the motherboard on the INTB# pin. If Pins 2-3 of J3 and J4 are shorted, SCSI channel 2 can be used by loading the standard Symbios Logic SCSI drivers. In this case, the MegaRAID Express drivers will not see any SCSI devices connected to SCSI channel. 2.	Short Pins 2-3 This setting is Reserved. Not currently supported by any motherboard	Short Pins 2-3 This setting is Reserved. Not currently supported by any motherboard

SCSI channel 1 is the channel whose interrupt is invoked to the INTC# pin on the MegaRAID Express PCI slot. SCSI channel 2 is the channel whose interrupt is routed to the INTD# pin on the MegaRAID Express PCI slot.

Step 4 Install Cache Memory

> ***Important***
>
> *A minimum of 8 MB of cache memory is recommended. The cache memory must be installed before MegaRAID Express is operational*. The SIMM slot *must be populated with one cache memory SIMM. You can use Fast Page Mode or EDO DRAM SIMMs.* ***You can only use EDO DRAM if the firmware version is version 2.05 or higher.***

Fast Page Mode Memory Use a +5V Fast Page Mode x 36 SIMM with 60 ns RAS Access time. The memory technology in the SIMM must support equal row and column size.

Volt	Speed	Parity	Type	BBU Support	Bank I	Total Memory
5 V	60ns	Yes	Single-sided	No	1M x 36	4 MB
5 V	60ns	Yes	Single-sided	No	4M x 36	16 MB
5 V	60ns	Yes	Single-sided	No	16M x 36	64 MB
5 V	60ns	Yes	Double-sided	No	2M x 36	8 MB
5 V	60ns	Yes	Double-sided	No	8M x 36	32 MB
5 V	60ns	Yes	Double-sided	No	32M x 36	128 MB

EDO Memory Use a +5V EDO x 36 SIMM with 50 ns RAS Access time.

Volt	Speed	Parity	Type	BBU Support	Bank I	Total Memory
5 V	50ns	Yes	Single-sided	No	1M x 36	4 MB
5 V	50ns	Yes	Single-sided	No	4M x 36	16 MB
5 V	50ns	Yes	Single-sided	No	16M x 36	64 MB
5 V	50ns	Yes	Double-sided	No	2M x 36	8 MB
5 V	50ns	Yes	Double-sided	No	8M x 36	32 MB
5 V	50ns	Yes	Double-sided	No	32M x 36	128 MB

Recommended SIMMs Call AMI technical support at 770-246-8600 for the latest list of approved memory vendors.

Cont'd

Install the SIMM on the MegaRAID Express controller card. Lay the controller card component-side up on a clean static-free surface and install the SIMM. The SIMM clicks into place, indicating proper seating in the socket, as shown below.

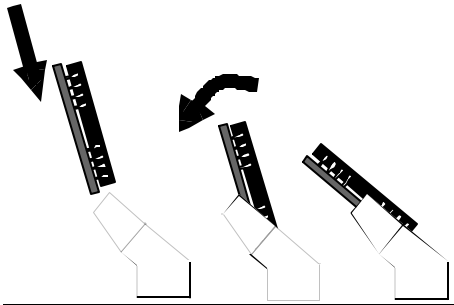

Step 5 Install MegaRAID Express

Install the MegaRAID Express card in a PCI expansion slot that supports American Megatrends IRQ routing, as shown below:

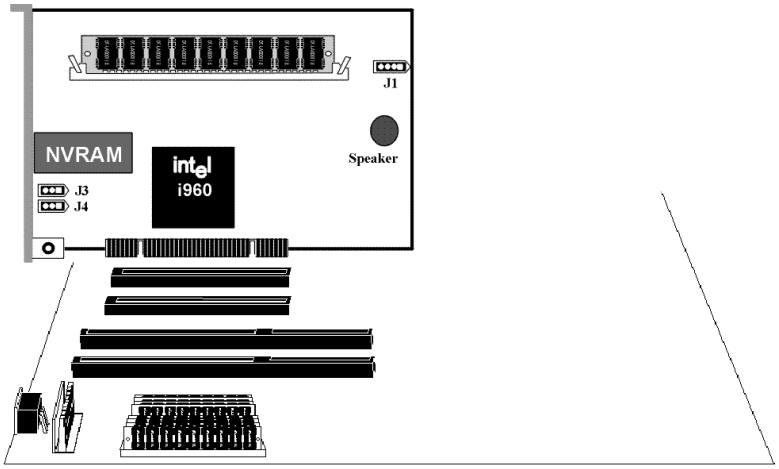

Step 6 Set Target IDs

Set target identifiers (TIDs) on the SCSI devices. Each device in a specific SCSI channel must have a unique TID in that channel. Non-disk devices (CD-ROM or tapes) should have unique SCSI IDs *regardless of the channel where they are connected*. See the documentation for each SCSI device to set the TIDs. MegaRAID Express automatically occupies TID 7 in each SCSI channel. Eight-bit SCSI devices can only use the TIDs from 0 to 6. 16-bit devices can use the TIDs from 0 to 15. The arbitration priority for a SCSI device depends on its TID.

Priority	Highest											Lowest
TID	7	6	5	…	2	1	0	15	14	…	9	8

Important

Non-disk devices (CD-ROM or tapes) should have unique SCSI IDs regardless of the channel they are connected to.

Device Identification

All Logical Drives on each SCSI bus are identified to the host as ID 0. Differentiation is made possible by utilizing Logical Unit Identifiers (LUNs). ID 0 cannot be used for non-disk devices because they are limited to IDs 1 through 6. The MegaRAID Express is limited to eight logical drives because LUNs are used to present logical drives. The SCSI-2 ANSI specification has a limit of 8 LUNs per ID. The SCSI-3 specification has increased the number of LUNs to 16. An example of the MegaRAID Express ID mapping method is shown below.

Example of MegaRAID Express ID Mapping

ID	Channel 1	Channel 2
0	A1-1	A1-2
1	A2-1	Scanner
2	CD	A2-3
3	A2-5	A2-6
4	CD	A3-1
5	A4-1	Tape
6	Optical	A5-1
7	Reserved	Reserved
8	A5-2	A5-3
9	A5-6	A5-7
10	A6-1	A6-2
11	A6-4	A6-5
12	A6-7	A6-8
13	A7-2	A7-3
14	A7-5	A7-6
15	A7-8	A8-1

As Presented to the Operating System

ID	LUN	Device	ID	LUN	Device
0	0	Disk (A1-X)	1	0	Scanner
0	1	Disk (A2-X)	2	0	CD
0	2	Disk (A3-X)	3	0	Tape
0	3	Disk (A4-X)	4	0	CD
0	4	Disk (A5-X)	5	0	Tape
0	5	Disk (A6-X)	6	0	Optical
0	6	Disk (A7-X)			
0	7	Disk (A8-X)			

Step 7 Power Up

Replace the computer cover and reconnect the AC power cords. Turn power on to the host computer. Set up the power supplies so that the SCSI devices are powered up at the same time as or before the host computer. If the computer is powered up before a SCSI device, the device might not be recognized.

During boot, the MegaRAID Express BIOS message appears:

```
MegaRAID Express Disk Array Adapter BIOS Version x.xx date
Copyright (c) American Megatrends, Inc.
Firmware Initializing... [ Scanning SCSI Device …(etc.)… ]
```

The firmware takes several seconds to initialize. During this time the adapter will scan each SCSI channel. When it is ready, the following lines appear:

```
Host Adapter-1 Firmware Version x.xx DRAM Size 4 MB
0 Logical Drives found on the Host Adapter
0 Logical Drives handled by BIOS
Press <Ctrl><M> to run MegaRAID Express BIOS Configuration Utility
```

The <Ctrl> <M> prompt times out after several seconds.

The MegaRAID Express host adapter (controller) number, firmware version, and cache DRAM size are displayed in the second portion of the BIOS message. The numbering of the controllers follows the PCI slot scanning order used by the host motherboard.

Step 8 Run MegaRAID Express Configuration utility

Press <Ctrl> <M> to run the MegaRAID Express Configuration utility utility. See the *MegaRAID Express 000 Configuration Software Guide* for information about running MegaRAID Express Configuration utility.

Step 9 Install the Operating System Driver

> ***Important***
>
> *When booting the system from a drive connected to a MegaRAID controller and using EMM386.EXE, MEGASPI.SYS must be loaded in CONFIG.SYS before EMM386.EXE is loaded. If you do not do this, you cannot access the boot drive after EMM386 is loaded.*

DOS ASPI Driver The MegaRAID Express ASPI driver can be used under DOS, Windows 3.*x*, and Windows 95. The DOS ASPI driver supports:

- up to six non-disk SCSI devices (each SCSI device must use a unique SCSI ID regardless of the SCSI channel it resides on. SCSI IDs 1 through 6 are valid,
- up to six MegaRAID Express adapters (although you should only configure one MegaRAID adapter per system if possible).

ASPI Driver The ASPI driver is MEGASPI.SYS. It supports disk drives, tape drives, CD-ROM drives, etc. You can use it to run CorelSCSI, Novaback, PC Tools, and other software that requires an ASPI driver. *CorelSCSI, Novaback, and PC Tools are not provided with MegaRAID Express.* Copy MEGASPI.SYS to your hard disk drive. Add the following line to CONFIG.SYS. *MEGASPI.SYS must be loaded in CONFIG.SYS before EMM386.EXE is loaded.*

```
device=<path>\MEGASPI.SYS /v
```

Parameters The MEGASPI.SYS parameters are:

Parameter	Description
/h	INT 13h support is not provided.
/v	Verbose mode. All message are displayed on the screen.
/a	Physical drive access mode. Direct access to physical drives is permitted.
/q	Quiet mode. All message except error message are suppressed.

Cont'd

Step 9 Install Operating System Driver, Continued

CD-ROM Driver A device driver is provided with MegaRAID Express for CD-ROM drives operating under DOS, Windows 3.*x*, and Windows 95. The driver filename is AMICDROM.SYS. The MEGASPI.SYS ASPI manager must be added to the CONFIG.SYS file before you can install the CD-ROM device driver. See the instructions on the previous screen for adding the MEGASPI.SYS driver. Copy AMICDROM.SYS to the root directory of the C: drive. Add the following line to CONFIG.SYS, making sure it is preceded by the line for MEGASPI.SYS:

```
DEVICE=C:\AMICDROM.SYS
```

Add the following to AUTOEXEC.BAT. Make sure it precedes the SMARTDRV.EXE line.

```
MSCDEX /D:MSCD001
```

MSCDEX is the CD-ROM drive extension file that is supplied with MS-DOS® and PC-DOS® Version 5.0 or later. See your DOS manual for the command line parameters for MSCDEX.

Summary

This chapter discussed hardware installation. See the *MegaRAID Express 000 Configuration Software Guide* for information about MegaRAID Express software utilities. You must configure the RAID system via software configuration utilities. The utility programs for configuring MegaRAID Express are:

Configuration Utility	Operating System
MegaRAID Express Configuration utility	independent of the operating system
MegaRAID Express Manager	DOS Novell NetWare 3.x, 4.x UnixWare
Power Console	Microsoft Windows NT

7 Troubleshooting

Problem	Suggested Solution
The system hangs during the boot process after installation.	Make sure the SCSI BIOS on the motherboard has been disabled.
The system hangs during the boot process after installation.	Make sure the MegaRAID Express adapter card is installed in the proper PCI expansion slot. It must be installed in the RAID Upgrade PCI slot..
Some operating systems do not load in a computer with a MegaRAID Express adapter.	Check the system BIOS configuration for PCI interrupt assignments. Make sure some Interrupts are assigned for PCI. Initialize the logical drive before installing the operating system.
One of the hard drive in the array fails often	Check the drive error counts using Power Console. Format the drive. Rebuild the drive If the drive continues to fail, replace the drive with another drive with the same capacity.
Pressed <Ctrl> <M>. Ran Megaconf.exe and tried to make a new configuration. The system hangs when scanning devices.	Check the drives IDs on each channel to make sure each device has a different ID. Check the termination. The device at the end of the channel must be terminated. Replace the drive cable.
Multiple drives connected to MegaRAID Express using the same power supply. There is a problem spinning the drives all at once.	Set the drives to spin on command. This will allow MegaRAID Express to spin two devices simultaneously.
Pressing <Ctrl> <M> or running megaconf.exe does not display the Management Menu.	These utilities require a color monitor.
At system power-up with the MegaRAID Express installed, the screen display is garbled or does not appear at all.	For proper cache memory operation, you must install at least 8 MB of cache memory in MegaRAID Express.
Cannot flash or update the EEPROM.	You may need a new EEPROM.

Problem	Suggested Solution
`Firmware Initializing...` appears and remains on the screen.	Make sure that TERMPWR is being properly provided to each peripheral device populated channel. Make sure that each end of the channel chain is properly terminated using the recommended terminator type for the peripheral device. The channel is automatically terminated at the MegaRAID Express card if only one cable is connected to a channel. Make sure (on a channel basis) only two type of cables are connected at any one time. If using an FPM SIMM, make sure the RAS Access Time parameter is 70 ns. Make sure that the MegaRAID Express controller is properly seated in the PCI slot.
What SCSI IDs can a non-hard disk device have and what is maximum number allowed per adapter?	Non-hard disk devices can only accommodate SCSI IDs 1, 2, 3, 4, 5 or 6, regardless of the channel used. A maximum of six non-hard disk devices are supported per MegaRAID Express adapter.
Why does a failed logical array still get a drive assignment?	To maintain the DOS Path statement integrity.

BIOS Boot Error Messages

Message	Problem	Suggested Solution
Adapter BIOS Disabled. No Logical Drives Handled by BIOS	The MegaRAID Express BIOS is disabled. Sometimes the BIOS is disabled to prevent booting from the BIOS.	Enable the BIOS via the MegaRAID Express Configuration utility utility.
Host Adapter at Baseport xxxx Not Responding	The BIOS cannot communicate with the adapter firmware.	Make sure MegaRAID Express is properly installed.
No MegaRAID Express Adapter	The BIOS cannot communicate with the adapter firmware.	Make sure MegaRAID Express is properly installed.
Configuration of NVRAM and drives mismatch. Run View/Add Configuration option of Configuration Utility. Press any key to run the Configuration Utility.	The configuration stored in the MegaRAID Express adapter does not match the configuration stored in the drives.	Press a key to run MegaRAID Express Manager. Choose View/Add Configuration from the Configure menu. Use View/Add Configuration to examine both the configuration in NVRAM and the configuration stored on the disk drives. Resolve the problem by selecting one of the configurations.
1 Logical Drive Failed	A logical drive failed to sign on.	Make sure all physical drives are properly connected and are powered on. Run MegaRAID Express Manager to find out if any physical drives are not responding. Reconnect, replace, or rebuild any drive that is not responding.

Message	Problem	Suggested Solution
X Logical Drives Degraded	x number of logical drives signed on in a degraded state.	Make sure all physical drives are properly connected and are powered on. Run MegaRAID Express Manager to find out if any physical drives are not responding. Reconnect, replace, or rebuild any drive that is not responding.
1 Logical Drive Degraded	A logical drive signed on in a degraded state.	Make sure all physical drives are properly connected and are powered on. Run MegaRAID Express Manager to find out if any physical drives are not responding. Reconnect, replace, or rebuild any drive that is not responding.
Insufficient memory to run BIOS. Press any key to continue…	Not enough MegaRAID Express memory to run MegaRAID Express BIOS.	Make sure MegaRAID Express memory has been properly installed.
Insufficient Memory	Not enough memory on the MegaRAID Express adapter to support the current configuration.	Make sure MegaRAID Express memory has been properly installed.
The following SCSI IDs are not responding: Channel x:a.b.c	The physical drives with SCSIO IDs a, b, and c are not responding on SCSI channel x.	Make sure the physical drives are properly connected and are powered on.

Other BIOS Error Messages

Message	Problem	Suggested Solution
Following SCSI disk not found and no empty slot available for mapping it	The physical disk roaming feature did not find the physical disk with the displayed SCSI ID. No slot is available to map the physical drive. MegaRAID Express cannot resolve the physical drives into the current configuration.	Reconfigure the array.
Following SCSI IDs have the same data y, z Channel x: a, b, c	The physical drive roaming feature found the same data on two or more physical drive on channel x with SCSI IDs a, b, and c. MegaRAID Express cannot determine the drive that has the duplicate information.	Remove the drive or drives that should not be used.
Unresolved configuration mismatch between disks and NVRAM on the adapter	The configuration stored in the MegaRAID Express NVRAM does not match the configuration stored on the drives.	Press a key to run MegaRAID Express Manager. Choose View/Add Configuration from the Configure menu. Use View/Add Configuration to examine both the configuration in NVRAM and the configuration stored on the disk drives. Resolve the problem by selecting one of the configurations.

DOS ASPI Driver Error Messages

Message	Corrective Action
American Megatrends Inc. ASPI Manager has NOT been loaded.	The ASPI manager is not loaded. One of the failure codes listed below is displayed next.
Controller setup FAILED error code=[0xab]	Correct the condition that caused the failure. The failure codes are: 0x40 No MegaRAID Express adapters found 0x80 Timed out waiting for interrupt to be posted 0x81 Timed out waiting for MegaRAID Express Response command. 0x82 Invalid command completion count. 0x83 Invalid completion status received. 0x84 Invalid command ID received. 0x85 No MegaRAID Express adapters found or no PCI BIOS support. 0x90 Unknown Setup completion error
No non-disk devices were located	The driver did not find any non-hard drive devices during scanning. A SCSI device that is not a hard disk drive, such as a tape drive or CD-ROM drive, must be attached to this SCSI channel. The SCSI ID must be unique for each adapter and cannot be SCSI ID 0. The supported SCSI IDs are 1, 2, 3, 4, 5, and 6.
'ERROR: VDS support is *INACTIVE* for MegaRAID Express logical drives	The /h option is appended to driver in CONFIG.SYS or this driver is used with a BIOS that is earlier than v1.10, or no logical drives are configured.

Other Potential Problems

Topic	Information
DOS ASPI	MEGASPI.SYS, the MegaRAID Express DOS ASPI manager, uses 6 KB of system memory once it is loaded.
CD-ROM drives under DOS	At this time, copied CDs are not accessible from DOS even after loading MEGASPI.SYS and AMICDROM.SYS.
Physical Drive Errors	To display the MegaRAID Express Manager Media Error and Other Error options, press <F2> after selecting a physical drive under the Physical Drive menu, selected from the Objects menu. A Media Error is an error that occurred while actually transferring data. An Other Error is an error that occurs at the hardware level because of a device failure, poor cabling, bad termination, signal loss, etc.
Virtual Sizing	The Virtual Sizing option enables RAID expansion. Virtual Sizing must be enabled to increase the size of a logical drive or add a physical drive to an existing logical drive. Run MegaRAID Express Manager by pressing <Ctrl> <M> to enable Virtual Sizing. Select the Objects menu, then select the Logical Drive menu. Select View/Update Parameters. Set Virtual Sizing to Enabled.
BSD Unix	We do not provide a driver for BSDI Unix. MegaRAID Express does not support BSDI Unix.
Multiple LUNs	MegaRAID Express supports one LUN per each target ID. No multiple LUN devices are supported.
MegaRAID Express Power Requirements	The Maximum MegaRAID Express power requirements are 15 watts at 5V and 3 Amps.

Topic	Information
SCSI Bus Requirements	The ANSI specification dictates the following: The maximum signal path length between terminators is 3 meters when using up to 4 maximum capacitance (25 pF) devices and 1.5 meters when using more than 4 devices. SCSI devices should be uniformly spaced between terminators, with the end devices located as close as possible to the terminators. The characteristic impedance of the cable should be 90 +/- 6 ohms for the /REQ and /ACK signals and 90 +/- 10 ohms for all other signals. The stub length(the distance from the controller's external connector to the mainline SCSI bus) shall not exceed.1m (approximately 4 inches). The spacing of devices on the mainline SCSI bus should be at least three times the stub length. All signal lines shall be terminated once at both ends of the bus powered by the TERMPWR line.

Topic	Information
Windows NT Installation	When Windows NT is installed via a bootable CD, the devices on the MegaRAID Express will not be recognized until after the initial reboot. The Microsoft documented workaround is in SETUP.TXT: SETUP.TXT is on the CD To install drivers when Setup recognizes one of the supported SCSI host adapters without making the devices attached to it available for use: 1 Restart Windows NT Setup. 2 When Windows NT Setup displays Setup is inspecting your computer's hardware configuration..., press <F6> to prevents Windows NT Setup from performing disk controller detection. This allows you to install the driver from the Drivers disk you created. All SCSI adapters must be installed manually. 3 When Windows NT Setup displays Setup could not determine the type of one or more mass storage devices installed in your system, or you have chosen to manually specify an adapter, press S to display a list of supported SCSI host adapters. 4 Select Other from the bottom of the list. 5 Insert the Drivers Disk you made when prompted to do so and select MegaRAID Express from this list. In some cases, Windows NT Setup repeatedly prompts to swap disks. Windows NT will now recognize any devices attached to this adapter. Repeat this step for each host adapter not already recognized by Windows NT Setup.

A SCSI Cables and Connectors

Cable Considerations

MegaRAID Express Single-Ended Ultra SCSI Understanding the cable requirements, termination and stub lengths is key to the successful implementation of a Ultra-SCSI subsystem.

SCSI Cables - Up to Four Devices The total external SCSI cable length for single-ended when using up to 4 Ultra-SCSI devices (maximum. capacitance of device = 25pf) should be less than or equal to:

```
(3 meter-(SCSI signal length on the MegaRAID Express)-(SCSI length in storage box)

= (3 meter - 0.305 meter - SCSI length in storage box)

= 2.695 - SCSI length in Storage box
```

SCSI Cables - More than Four Devices The total external SCSI cable length for single-ended when using from five to eight Ultra-SCSI devices (max. cap of device = 25pf) should be less than or equal to:

```
(1.5 meter-(SCSI signal length on MegaRAID Express)-(SCSI length in storage box)

= (1.5 meter - 0.305 meter - SCSI length in storage box)

= 1.195 - SCSI length in Storage box
```

Spacing Devices The SCSI devices should be uniformly spaced between terminators with the end devices located as close as possible to the terminators.

SCSI Signal Path The SCSI signal path is a controlled impedance environment with the following characteristic impedance:

```
90 ohms +/- 6 ohms for the REQ and ACK signals

90 ohms +/- 10 ohms for all other signals
```

Termination and Stub Placement for MegaRAID Express Adapters

Termination placement	For a well behaved SCSI bus/cabling system, there should be termination enabled at both ends of the SCSI bus. Avoid adding terminators in the middle of the bus. The end devices must be located as close as possible to the terminators. Active terminators must be used. Terminators employing a 220 ohm resistor to 5 Volts and a 330 ohm resistor to ground on each signal shall not be used.
Stub length	The stub length shall not exceed 0.1 meter. The spacing of devices on the SCSI bus should be at least three times the stub length to avoid stub clustering.
Cables	Teflon flat ribbon cables give the best performance in the Ultra-SCSI environment. These cables should be used for all internal cabling. To minimize discontinuities and signal reflections, the use of cables with different impedance's on the same bus should be minimized.

SCSI Connectors

MegaRAID Express has no SCSI connectors. The SCSI connectors must be on the motherboard.

68-Pin High Density SCSI Internal Connectors

The motherboard should provide 68-pin high density 0.050” pitch unshielded SCSI connectors. These SCSI connectors provide all signals needed to connect MegaRAID Express to wide SCSI devices. The connector pinouts are for a single-ended primary bus (P-CABLE) as specified in SCSI-3 Parallel Interface X3T9.2, Project 885-D, revision 12b, date July 2, 1993.

The cable assemblies that interface with this 68-pin connector are:

- flat ribbon or twisted pair cable for connecting internal wide SCSI devices,
- flat ribbon or twisted pair cable for connecting internal and external wide SCSI devices,
- cable assembly for converting from internal wide SCSI connectors to internal non-wide (Type 2) connectors,
- cable assembly for converting from internal wide to internal non-wide SCSI connectors (Type 30), and
- cable assembly for converting from internal wide to internal non-wide SCSI connectors.

Cont’d

68-Pin High Density Connectors, Continued

Cable Assembly for Internal Wide SCSI Devices The cable assembly for connecting internal wide SCSI devices is shown below:

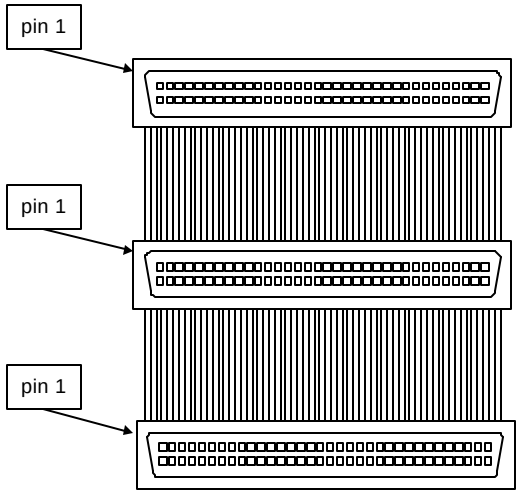

Connectors: 68 position plug (male)
AMP - 786090-7
Cable: Flat Ribbon or Twisted-Pair Flat Cable
68 Conductor 0.025 Centerline
30 AWG

Cont'd

68-Pin High Density Connectors, Continued

Connecting Internal and External Wide Devices The cable assembly for connecting internal wide and external wide SCSI devices is shown below:

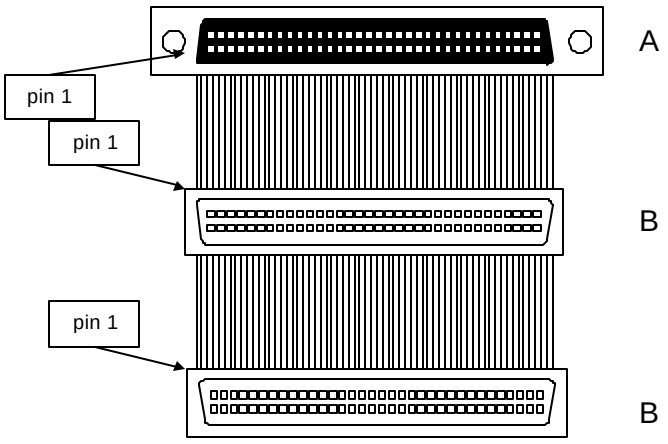

Connector A: 68 position panel mount receptacle
with 4-40 holes (female)
AMP - 786096-7

NOTE: To convert to 2-56 holes, use screwlock kit 749087-1, 749087-2, or 750644-1 from AMP

Connector B: 68 position plug (male)
AMP - 786090-7

Cable: Flat Ribbon or Twisted-Pair Flat Cable
68 Conductor 0.025 Centerline
30 AWG

Cont'd

Converting Internal Wide to Internal Non-Wide (Type 2) The cable assembly for converting internal wide SCSI connectors to internal non-wide SCSI connectors is shown below:

68 POSITION CONNECTOR CONTACT NUMBER		50 POSITION CONNECTOR CONTACT NUMBER
6		1
40		2
7		3
41		4
*		
*		
*		
49		20
16		21
50		22
17	OPEN	23
51	OPEN	24
18	OPEN	25
52		26
19		27
*		
*		
*		
29		47
63		48
30		49
64		50

TABLE 1: CONNECTOR CONTACT CONNECTION FOR WIDE TO NON-WIDE CONVERSION

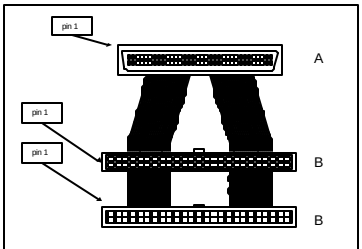

Connector A68 position plug (male)
AMP - 749925-5
Connector B:50 position IDC receptacle (female)
AMP - 499252-4, 1-746285-0, 1-746288-
Wire: Twisted-Pair Flat Cable or
Laminated Discrete Wire Cable
25 pair 0.050 centerline
28 AWG

Cont'd

68-Pin High Density Connectors, Continued

Converting Internal Wide to Internal Non-Wide (Type 30) The cable assembly for connecting internal wide SCSI devices to internal non-wide SCSI devices is shown below:

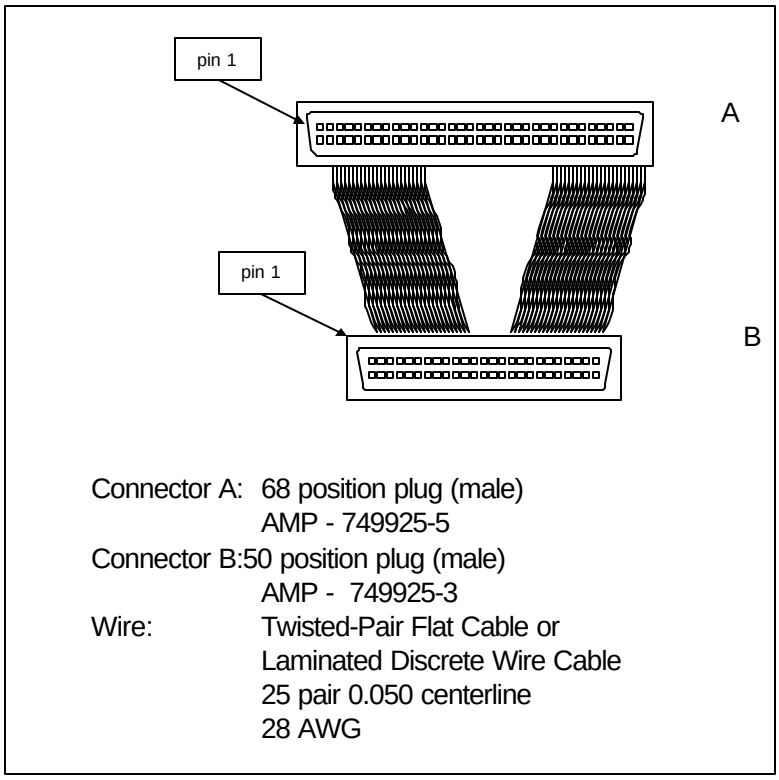

Connector A: 68 position plug (male)
AMP - 749925-5
Connector B:50 position plug (male)
AMP - 749925-3
Wire: Twisted-Pair Flat Cable or
Laminated Discrete Wire Cable
25 pair 0.050 centerline
28 AWG

Cont'd

Converting from Internal Wide to Internal Non-Wide (Type 3) The cable assembly for connecting internal wide SCSI devices to internal non-wide (Type 3) SCSI devices is shown below:

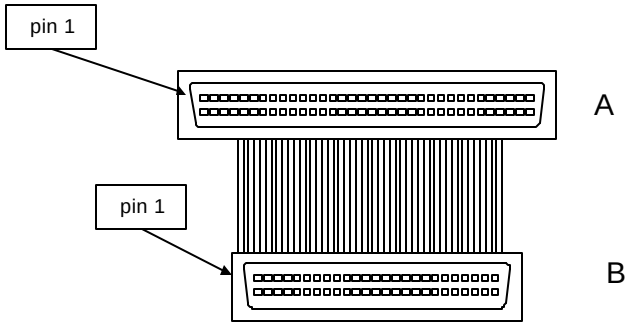

Connector A: 68 position plug (male)
AMP - 786090-7
Connector B:50 position plug (male)
AMP - 786090-7
Wire: Flat ribbon or twisted-pair flat cable
50 conductor 0.025 centerline
30 AWG

SCSI Cable Vendors

Manufacturer	Telephone Number
Cables To Go	Voice: 800-826-7904 Fax: 800-331-2841
System Connection	Voice: 800-877-1985
Technical Cable Concepts	Voice: 714-835-1081
GWC	Voice: 818-579-0888

SCSI Connector Vendors

Manufacturer	Connector Part Number	Back Shell Part Number
AMP	749111-4	749193-1
Fujitsu	FCN-237R050-G/F	FCN-230C050-D/E
Honda	PCS-XE50MA	PCS-E50LA

High-Density 68-Pin SCSI Connector Pinout

Signal	Connector Pin	Cable Pin	Cable Pin	Connector Pin	Signal
Ground	1	1	2	35	-DB(12)
Ground	2	3	4	36	-DB(13)
Ground	3	5	6	37	-DB(14)
Ground	4	7	8	38	-DB(15)
Ground	5	9	10	39	-DB(P1)
Ground	6	11	12	40	-DB(0)
Ground	7	13	14	41	-DB(1)
Ground	8	15	16	42	-DB(2)
Ground	9	17	18	43	-DB(3)
Ground	10	19	20	44	-DB(4)
Ground	11	21	22	45	-DB(5)
Ground	12	23	24	46	-DB(6)
Ground	13	25	26	47	-DB(7)
Ground	14	27	28	48	-DB(P)
Ground	15	29	30	49	SWAP L
Ground	16	31	32	50	SHELF_OK
TERMPWR	17	33	34	51	TERMPWR
TERMPWR	18	35	36	52	TERMPWR
Reserved	19	37	38	53	Reserved
Ground	20	39	40	54	FAULT_CLK H
Ground	21	41	42	55	-ATN
Ground	22	43	44	56	FAULT_DATA H
Ground	23	45	46	57	-BSY
Ground	24	47	48	58	-ACK
Ground	25	49	50	59	-RST
Ground	26	51	52	60	-MSG
Ground	27	53	54	61	-SEL
Ground	28	55	56	62	-C/D
Ground	29	57	58	63	-REQ
Ground	30	59	60	64	-I/O
Ground	31	61	62	65	-DB(8)
Ground	32	63	64	66	-DB(9)
Ground	33	65	66	67	-DB(10)
Ground	34	67	68	68	-DB(11)

Cont’d

68-Pin SCSI Connector Pinout, Continued

High-Density Connector The following applies to the high-density SCSI connector table on the previous screen:

- A hyphen before a signal name indicates that signal is active low.
- The connector pin refers to the conductor position when using 0.025 inch centerline flat ribbon cable with a high-density connector (AMPLIMITE.050 Series connectors).
- Eight-bit devices connected to the P-Cable must leave the following signals open: -DB (8), -DB (9), -DB (10), -DB (11), -DB(12), -DB (13), -DB (14), -DB 15), and -DB (P1).
- All other signals should be connected as defined.

Caution

Lines labeled RESERVED should be connected to Ground in the bus terminator assemblies or in the end devices on the SCSI cable.

RESERVED lines should be open in the other SCSI devices, but can be connected to Ground.

B Audible Warnings

MegaRAID Express has an onboard tone generator that indicates events and errors.

Tone Pattern	Meaning	Examples
Three seconds on and one second off	A logical drive is offline.	One or more drives in a RAID 0 configuration failed. Two or more drives in a RAID 1, 3, or 5 configuration failed.
One second on and one second off	A logical drive is running in degraded mode.	One drive in a RAID 3 or 5 configuration failed.

Glossary

Array A grouping or array of disk drives combines the storage space on the disk drives into a single segment of contiguous storage space. MegaRAID can group disk drives on one or more SCSI channels into an array. A hot spare drive does not participate in an array.

Array Management Software Software that provides common control and management for a disk array. Array Management Software most often executes in a disk controller or intelligent host bus adapter, but can also execute in a host computer. When it executes in a disk controller or adapter, Array Management Software is often called firmware.

Array Spanning Array spanning by a logical drive combines storage space in two arrays of disk drives into a single, contiguous storage space in a logical drive. MegaRAID logical drives can span consecutively numbered arrays that each consist of the same number of disk drives. Array spanning promotes RAID levels 1, 3, and 5 to RAID levels 10, 30, and 50, respectively. See also *Disk Spanning*.

Asynchronous Operations Operations that bear no relationship to each other in time and can overlap. The concept of asynchronous I/O operations is central to independent access arrays in throughput-intensive applications.

Cache I/O A small amount of fast memory that holds recently accessed data. Caching speeds subsequent access to the same data. It is most often applied to processor-memory access, but can also be used to store a copy of data accessible over a network. When data is read from or written to main memory, a copy is also saved in cache memory with the associated main memory address. The cache memory software monitors the addresses of subsequent reads to see if the required data is already stored in cache memory. If it is already in cache memory (a cache hit), it is read from cache memory immediately and the main memory read is aborted (or not started.) If the data is not cached (a cache miss), it is fetched from main memory and saved in cache memory.

Channel An electrical path for the transfer of data and control information between a disk and a disk controller.

Cont'd

Consistency Check An examination of the disk system to determine whether all conditions are valid for the specified configuration (such as parity.)

Cold Swap A cold swap requires that you turn the power off before replacing a defective hard drive in a disk subsystem.

Data Transfer Capacity The amount of data per unit time moved through a channel. For disk I/O, bandwidth is expressed in megabytes per second (MBs).

Degraded A drive that has become non-functional or has decreased in performance.

Disk A non-volatile, randomly addressable, rewritable mass storage device, including both rotating magnetic and optical disks and solid-state disks, or non-volatile electronic storage elements. It does not include specialized devices such as write-once-read-many (WORM) optical disks, nor does it include so-called RAM disks implemented using software to control a dedicated portion of a host computer volatile random access memory.

Disk Array A collection of disks from one or more disk subsystems combined with array management software. It controls the disks and presents them to the array operating environment as one or more virtual disks.

Disk Duplexing A variation on disk mirroring where a second disk adapter or host adapter and redundant disk drives are present.

Disk Mirroring Writing duplicate data to more than one (usually two) hard disks to protect against data loss in the event of device failure. It is a common feature of RAID systems.

Disk Spanning Disk spanning allows multiple disk drives to function like one big drive. Spanning overcomes lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources. For example, four 400 MB disk drives can be combined to appear to the operating system as one single 1600 MB drive. See also *Array Spanning* and *Spanning*.

Cont'd

Disk Striping A type of disk array mapping. Consecutive stripes of data are mapped round-robin to consecutive array members. A striped array (RAID Level 0) provides high I/O performance at low cost, but provides lowers data reliability than any of its member disks.

Disk Subsystem A collection of disks and the hardware that connects them to one or more host computers. The hardware can include an intelligent controller or the disks can attach directly to a host computer I/O a bus adapter.

Double Buffering A technique that achieves maximum data transfer bandwidth by constantly keeping two I/O requests for adjacent data outstanding. A software component begins a double-buffered I/O stream by issuing two requests in rapid sequence. Thereafter, each time an I/O request completes, another is immediately issued. If the disk subsystem is capable of processing requests fast enough, double buffering allows data to be transferred at the full-volume transfer rate.

Failed Drive A drive that has ceased to function or consistently functions improperly.

Fast SCSI A variant on the SCSI-2 bus. It uses the same 8-bit bus as the original SCSI-1, but runs at up to 10MB (double the speed of SCSI-1.)

Firmware Software stored in read-only memory (ROM) or Programmable ROM (PROM). Firmware is often responsible for the behavior of a system when it is first turned on. A typical example would be a monitor program in a computer that loads the full operating system from disk or from a network and then passes control to the operating system.

FlexRAID Power Fail Option The FlexRAID Power Fail option allows a reconstruction to restart if a power failure occurs. This is the advantage of this option. The disadvantage is, once the reconstruction is active, the performance is slower because an additional activity is added.

Cont'd

Format	The process of writing zeros to all data fields in a physical drive (hard drive) to map out unreadable or bad sectors. Because most hard drives are factory formatted, formatting is usually only done if a hard disk generates many media errors.
GB	Shorthand for 1,000,000,000 (10 to the ninth power) bytes. It is the same as 1,000 MB (megabytes).
Host-based Array	A disk array with an Array Management Software in its host computer rather than in a disk subsystem.
Host Computer	Any computer that disks are directly attached to. Mainframes, servers, workstations, and personal computers can all be considered host computers.
Hot Spare	A stand-by drive ready for use if another drive fails. It does not contain any user data. Up to eight disk drives can be assigned as hot spares for an adapter. A hot spare can be dedicated to a single redundant array or it can be part of the global hot-spare pool for all arrays controlled by the adapter.
Hot Swap	The substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is running (performing its normal functions). Hot swaps are manual.
I/O Driver	A host computer software component (usually part of the operating system) that controls the operation of peripheral controllers or adapters attached to the host computer. I/O drivers communicate between applications and I/O devices, and in some cases participates in data transfer.
Initialization	The process of writing zeros to the data fields of a logical drive and generating the corresponding parity to put the logical drive in a Ready state. Initializing erases previous data and generates parity so that the logical drive will pass a consistency check. Arrays can work without initializing, but they can fail a consistency check because the parity fields have not been generated.

Cont'd

Logical Disk A set of contiguous chunks on a physical disk. Logical disks are used in array implementations as constituents of logical volumes or partitions. Logical disks are normally transparent to the host environment, except when the array containing them is being configured.

Logical Drive A virtual drive within an array that can consist of more than one physical drive. Logical drives divide the contiguous storage space of an array of disk drives or a spanned group of arrays of drives. The storage space in a logical drive is spread across all the physical drives in the array or spanned arrays. Each MegaRAID adapter can be configured with up to eight logical drives in any combination of sizes. Configure at least one logical drive for each array.

Mapping The conversion between multiple data addressing schemes, especially conversions between member disk block addresses and block addresses of the virtual disks presented to the operating environment by Array Management Software.

MB (Megabyte) An abbreviation for 1,000,000 (10 to the sixth power) bytes. It is the same as 1,000 KB (kilobytes).

Multi-threaded Having multiple concurrent or pseudo-concurrent execution sequences. Used to describe processes in computer systems. Multi-threaded processes allow throughput-intensive applications to efficiently use a disk array to increase I/O performance.

Operating Environment The operating environment includes the host computer where the array is attached, any I/O buses and adapters, the host operating system, and any additional software required to operate the array. For host-based arrays, the operating environment includes I/O driver software for the member disks, but does not include Array Management Software, which is regarded as part of the array itself.

Cont’d

Parity Parity is an extra bit added to a byte or word to reveal errors in storage (in RAM or disk) or transmission. Parity is used to generate a set of redundancy data from two or more parent data sets. The redundancy data can be used to reconstruct one of the parent data sets. However, parity data does not fully duplicate the parent data sets. In RAID, this method is applied to entire drives or stripes across all disk drives in an array. Parity consists of dedicated parity, in which the parity of the data on two or more drives is stored on an additional drive, and distributed parity, in which the parity data are distributed among all the drives in the system. If a single drive fails, it can be rebuilt from the parity of the respective data on the remaining drives.

Partition An array virtual disk made up of logical disks rather than physical ones. Also known as logical volume.

Physical Disk A hard disk drive that stores data. A hard disk drive consists of one or more rigid magnetic discs rotating about a central axle with associated read/write heads and electronics.

Physical Disk Roaming The ability of some adapters to detect when hard drives have been moved to a different slots in the computer, for example, after a hot swap.

Protocol A set of formal rules describing how to transmit data, especially across a network. Low level protocols define the electrical and physical standards to be observed, bit- and byte- ordering, and the transmission and error detection and correction of the bit stream. High level protocols deal with the data formatting, including the message syntax, the terminal-to-computer dialogue, character sets, and sequencing of messages.

RAID Redundant Array of Independent Disks is an array of multiple small, independent hard disk drives that yields performance exceeding that of a Single Large Expensive Disk (SLED). A RAID disk subsystem improves I/O performance on a server using only a single drive. The RAID array appears to the host server as a single storage unit. I/O is expedited because several disks can be accessed simultaneously.

Cont'd

RAID Levels	A style of redundancy applied to a logical drive. It can increase the performance of the logical drive and can decrease usable capacity. Each logical drive must have a RAID level assigned to it. The RAID level drive requirements are: RAID 0 requires one or more physical drives, RAID 1 requires exactly two physical drives, RAID 3 requires at least three physical drives, RAID 5 requires at least three physical drives. RAID levels 10, 30, and 50 result when logical drives span arrays. RAID 10 results when a RAID 1 logical drive spans arrays. RAID 30 results when a RAID 3 logical drive spans arrays. RAID 50 results when a RAID 5 logical drive spans arrays.
RAID Migration	RAID migration is used to move between optimal RAID levels or to change from a degraded redundant logical drive to an optimal RAID 0. In Novell, the utility used for RAID migration is MEGAMGR and in Windows NT its Power Console. If a RAID 1 is being converted to a RAID 0, instead of performing RAID migration, one drive can be removed and the other reconfigured on the controller as a RAID 0. This is due to the same data being written to each drive.
Read-Ahead	A memory caching capability in some adapters that allows them to read sequentially ahead of requested data and store the additional data in cache memory, anticipating that the additional data will be needed soon. Read-Ahead supplies sequential data faster, but is not as effective when accessing random data.
Ready State	A condition in which a workable hard drive is neither online nor a hot spare and is available to add to an array or to designate as a hot spare.
Rebuild	The regeneration of all data from a failed disk in a RAID level 1, 3, 4, 5, or 6 array to a replacement disk. A disk rebuild normally occurs without interruption of application access to data stored on the array virtual disk.
Rebuild Rate	The percentage of CPU resources devoted to rebuilding.

Cont'd

Reconstruct The act of remaking a logical drive after changing RAID levels or adding a physical drive to an existing array.

Redundancy The provision of multiple interchangeable components to perform a single function to cope with failures or errors. Redundancy normally applies to hardware; a common form of hardware redundancy is disk mirroring.

Replacement Disk A disk available to replace a failed member disk in a RAID array.

Replacement Unit A component or collection of components in a disk subsystem that are always replaced as a unit when any part of the collection fails. Typical replacement units in a disk subsystem includes disks, controller logic boards, power supplies, and cables. Also called a hot spare.

SAF-TE SCSI Accessed Fault-Tolerant Enclosure. An industry protocol for managing RAID enclosures and reporting enclosure environmental information.

SCSI (Small Computer System Interface) A processor-independent standard for system-level interfacing between a computer and intelligent devices, including hard disks, floppy disks, CD-ROM, printers, scanners, etc. SCSI can connect up to 7 devices to a single adapter (or host adapter) on the computer's bus. SCSI transfers eight or 16 bits in parallel and can operate in either asynchronous or synchronous modes. The synchronous transfer rate is up to 40 MBs. SCSI connections normally use single ended drivers, as opposed to differential drivers. The original standard is now called SCSI-1 to distinguish it from SCSI-2 and SCSI-3, which include specifications of Wide SCSI (a 16-bit bus) and Fast SCSI (10 MBs transfer).

SCSI Channel MegaRAID controls the disk drives via SCSI-2 buses (channels) over which the system transfers data in either Fast and Wide or Ultra SCSI mode. Each adapter can control up to three SCSI channels. Internal and external disk drives can be mixed on channels 0 and 1, but not on channel 2.

Cont’d

SCSI Drive states A SCSI disk drive (physical drive) can be in one of these four states:

- Online — Powered-on and operational.
- Hot Spare — Powered-on stand-by disk drive, ready for use if an online disk fails.
- Rebuild — A disk drive to which one or more logical drives is restoring data.
- Not Responding — The disk drive is not present, is not powered-on, or has failed.

SMARTer Self-Monitoring, Analysis, and Reporting Technology with Error Recovery. An industry standard protocol for reporting server system information. Self-Monitoring, Analysis and Reporting Technology for disk drives is a specification designed to offer an early warning for some disk drive failures. These failures are predicted based upon actual performance degradation of drive components that are then reported back to a user through a graphical interface.

SNMP Simple Network Management Protocol, the most widely used protocol for communication management information between the managed elements of a network and a network manager. SNMP focuses primarily on the network backbone. The Internet standard protocol developed to manage nodes on an Internet Protocol (IP) network.

Spanning Array spanning by a logical drive combines storage space in two arrays of disk drives into a single, contiguous storage space in a logical drive. MegaRAID logical drives can span consecutively numbered arrays that each consist of the same number of disk drives. Array spanning promotes RAID levels 1, 3, and 5 to RAID levels 10, 30, and 50, respectively. See also *Disk Spanning* and *Spanning*.

Cont'd

Spare	A hard drive available to back up the data of other drives.
Stripe Size	The amount of data contiguously written to each disk. You can specify stripe sizes of 4 KB, 8 KB, 16 KB, 32 KB, 64 KB, and 128 KB for each logical drive. For best performance, choose a stripe size equal to or smaller than the block size used by the host computer.
Stripe Width	The number of disk drives across which the data are striped.
Striping	Segmentation of logically sequential data, such as a single file, so that segments can be written to multiple physical devices in a round-robin fashion. This technique is useful if the processor can read or write data faster than a single disk can supply or accept it. While data is being transferred from the first disk, the second disk can locate the next segment. Data striping is used in some modern databases and in certain RAID devices.
Terminator	A resistor connected to a signal wire in a bus or network for impedance matching to prevent reflections, e.g., a 50 ohm resistor connected across the end of an Ethernet cable. SCSI chains and some LocalTalk wiring schemes also require terminators.
Ultra-SCSI	An extension of SCSI-2 that doubles the transfer speed of Fast-SCSI, providing 20MBs on an 8-bit connection and 40MBs on a 16-bit connection.
Virtual Sizing	FlexRAID Virtual Sizing is used to create a logical drive up to 80 GB. A maximum of eight logical drives can be configured on a RAID controller and RAID migration is possible for all logical drives except the eighth. Because it is not possible to do migration on the last logical drive, the maximum space available for RAID migration is 560 GB.
Wide SCSI	A variant on the SCSI-2 interface. Wide SCSI uses a 16-bit bus, double the width of the original SCSI-1. Wide SCSI devices cannot be connected to a SCSI-1 bus. Wide SCSI supports transfer rates up to 20 MBs, like Fast SCSI.

Cont'd

Write-Through/Write-Back When the processor writes to main memory, the data is first written to cache memory, assuming that the processor will probably read this data again soon. In write-through cache, data is written to main memory at the same time it is written to cache memory. In write-back cache, data is written only to main memory when it is forced out of cache memory. Write-through caching is simpler than write-back because an entry to cache memory that must be replaced can be overwritten in cache memory because it will already have been copied to main memory. Write-back requires cache memory to initiate a main memory write of the flushed entry followed (for a processor read) by a main memory read. However, write-back is more efficient because an entry can be written many times to cache memory without a main memory access.

Index

G

H

I

J

L

M

O

P

R

S

T

U

V

W