Intel® RAID Controller SRCMR

Technical Product Specification

November 2001

Order Number A78493-001

The Intel® RAID Controller SRCMR may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are documented in the Intel RAID Controller SRCMR Specification Update.

Revision History

Revision	Revision History	Date
-001	Final release of the Intel® RAID Controller SRCMR Technical Product Specification.	November 2001

Contents

Figures

Tables

1 Introduction

1.1 Acronyms

The following acronyms are used throughout this document.

Table 1. Common Acronyms

Acronym	Description
API	Application Programmer Interface
DLL	Dynamic Linked Library
DOS	Generic term to reference either MS-DOS† or ROM-DOS
DMI	Desktop Management Interface – a system management specification
ECC	Error Correction Code (also “error correcting code” and “error correcting circuits”)
FUU	Flash Update Utility
FW	Firmware
HBA	Host Bus Adapter
IOP	I/O Processor, the Intel® 80302 I/O processor
Kb	Kilobit
KB	Kilobyte
LVD	Low Voltage Differential SCSI
Mb	Megabit
MB	Megabyte
MROMB	Modular RAID On Motherboard
PCB	Printed Circuit Board
PCI	Peripheral Component Interconnect
RAID	Redundant Array of Independent Disks
ROMB	RAID On Motherboard
SAF-TE	SCSI Accessed Fault Tolerant Enclosure
SES	SCSI Enclosure Services
SE	Single Ended SCSI device; SCSI device type
SCA	Single Connector Attachment – 80-pin SCSI connector on hot-swappable SE and LVD hard disks.
SCSI	Small Computer Systems Interface
SNMP	Simple Network Management Protocol
SRCMR	Product code for this Intel® RAID Controller. S = Server, R = RAID, C = Controller, MR = Modular ROMB (RAID On Motherboard)
StorCon	Storage Console – a character-based, menu-driven tool used for setting up, monitoring, and maintaining mass storage device subsystems based on SRCMR controllers
StorCon+	Storage Console Plus – A GUI-based, menu-driven tool used for setting up, monitoring, and maintaining mass storage device subsystems based on SRCMR controllers; runs only on Microsoft Windows†-based systems
XROM	PCI Expansion ROM – BIOS utility accessed at system POST

1.2 Documentation Conventions

The terms “RAID controller,” “RAID adapter,” “the controller,” and “the adapter” are used interchangeably throughout this document. Each term represents the physical PCB that integrates all the components of the SRCMR RAID PCI add-in card.

1.3 Product Overview

The Intel RAID Controller SRCMR is a low profile, 64-bit/66 MHz PCI, Modular RAID on Motherboard (MROMB) add-in adapter. It is a cost-effective “Intelligent RAID” solution. It was specifically developed to provide system design engineers with a flexible storage solution that offers a full-featured intelligent RAID I/O subsystem for their entry-level server platforms. The MROMB design reduces cost by utilizing the SCSI controller and connectors that are integrated into the system’s motherboard.

NOTE

The initial release of this product is specifically designed to work in only the Intel® server board SCB2 platform distributed by Intel’s Enterprise Platforms and Services Division (EPSD). Future server products will also be supported. For the latest list of supported servers, please visit, support.intel.com

The major components of the SRCMR RAID controller are:

- The Intel 80302 I/O Processor
- 32 MB (megabytes) of integrated SDRAM memory
- A 16 Mb (2 MB) Flash memory unit

Additional components required on the system motherboard are:

- The Adaptec AIC-7899W, dual-channel, Ultra160 (U160 or U3) SCSI controller
- One PCI slot with RAIDIOS (RAID I/O Steering) logic. See Section 2.8 for more details.
- One LVD SCSI connector for each channel of the SCSI controller (two total). Also proper SCSI cabling is required to connect the RAID controller to the SCSI controller’s channels

NOTE

These additional component requirements of the system motherboard can directly affect the RAID controller’s overall performance and functionality. Review the SRCMR User’s Guide for more detailed information concerning hardware installation, cabling, and termination.

1.4 Operating System Support

The following operating systems (OS) are fully validated and supported:

- Microsoft Windows 2000 Advanced Server (service pack 2a or higher)
- Microsoft Windows NT† 4.0 Server Enterprise Edition (service pack 6a or higher)

- Novell NetWare† 5.1 (support pack 2a or higher)
- Caldera UnixWare† 7.1.1
- Red Hat† Linux† 7.1 (2.4 kernel)
- Red Hat Linux 6.2 SBE2

The following operating systems are supported with limited compatibility validation:

- Windows 2000 Server and Professional
- Windows NT 4.0 Server, Terminal Server, and Workstation

1.5 List of Features

- Supports RAID levels 0, 1, 4, 5, and 10
- Supports up to two-channel, Adaptec Ultra160 SCSI controller (160 MB/second per channel)
- Online RAID level migration and capacity expansion without reboot
- RAID array roaming
- Instant availability and background initialization
- Automatic rebuild with private (dedicated) or pooled (global) hot fix (spare) drives
- Variable data strip size configurable per array
- Non hard disk drive SCSI device support (tape, CD-ROM)
- 32 MB of embedded ECC SDRAM
- Read/write controller and disk drive caching
- SAF-TE and SES intelligent enclosure support
- Hot-plug drive auto detection configurable for non-intelligent enclosures
- Hot-plug drive support

2 Hardware

2.1 Physical Layout

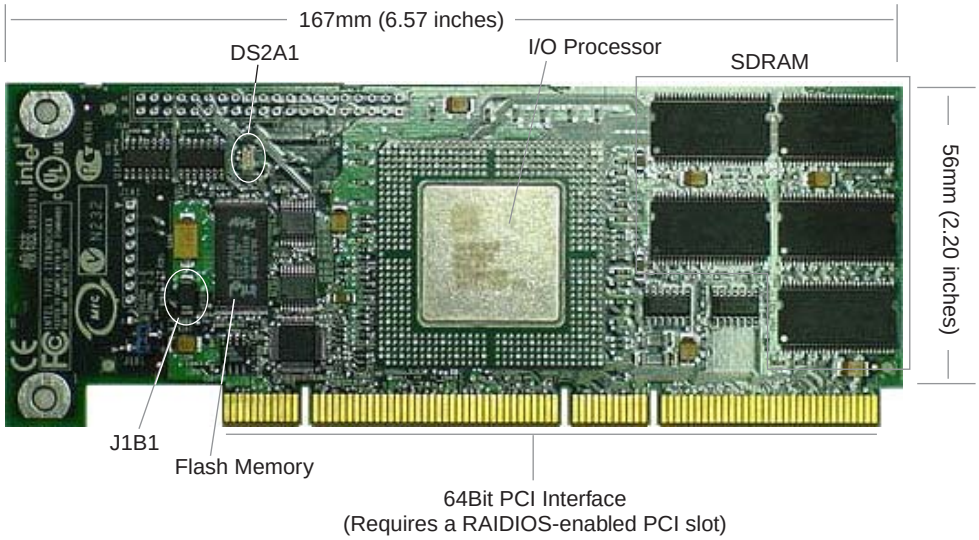

Figure 1. RAID Adapter Physical Layout

2.2 Major Components

2.2.1 Intel® 80302 I/O Microprocessor

The adapter features the Intel 80302 I/O processor (IOP). The major components of the 80302 IOP are: the processor core, PCI-to-PCI Bridge, Address Translation Units (ATU), Application Accelerator Unit (AAU), and Memory Controller Unit (MCU). The IOP's architecture is particularly useful in RAID applications. The 80302 and the 80303 IOPs are identical with the exception that the internal bus and memory interface of the 80302 run at 66 MHz. For more detailed information on the Intel 80302 IOP, visit the 80303 IOP's website at: *http://developer.intel.com/design/iio/index.htm*

2.2.1.1 i960® Core

The 80302 IOP uses the 100 MHz Intel® 80960JT-100 core. The 80302 IOP uses a 64-bit/66 MHz internal bus that can pump 528 MB/sec of data to and from the internal IOP peripherals. Among other features, it contains a 128-bit register bus, 16 KB two-way instruction cache, 4 KB direct-mapped data cache, 1 KB zero wait state data RAM, and single clock execution of most instructions.

2.2.1.2 PCI-to-PCI Bridge Unit

NOTE

The SRCMR controller does not utilize this feature.

The primary and secondary PCI buses of the PCI-to-PCI Bridge Unit support 64-bit/66 MHz PCI. Data throughput is 528 MB/sec. The upstream delayed read completion queue is 256 bytes, optimizing the 66 MHz PCI bus' performance.

2.2.1.3 Memory Controller Unit

The Memory Controller has been designed to use the latest, most cost effective, and highest performance SDRAM technology available. The 80302 IOP supports up to 512 MB of 64-bit 100 MHz ECC SDRAM operating at 66 MHz. It is backward compatible to support 64, 128, and 256 MB SDRAM technologies (the SRCMR only utilizes 32 MB of embedded SDRAM).

2.2.1.4 Address Translation Unit

The inbound R/W queues of the ATU are 256 bytes, twice that of the previous generation IOPs. The 80302 IOP can support up to 360 MB/sec for ATU inbound reads (initiated by a PCI master) and up to 520 MB/sec for ATU inbound writes and DMA transfers.

2.2.1.5 Application Acceleration Unit

The AAU, which is primarily used in RAID applications, is user-programmable for a 512 bytes or 1 KB queue.

2.2.2 Intel® Smart 3 FlashFile Flash Memory

This 3.3v, 16 Mb (2 MB) flash memory chip is used to store the RAID firmware. This non-volatile storage can be accessed for firmware updates and recovery. For firmware recovery, set the IOP mode select jumper (Figure 1: J2B1) to reset; place jumper on pins 1 and 2. For normal firmware updates, place the jumper on pins 2 and 3 (or remove totally).

2.2.3 Cache Memory Modules

The adapter provides 32 MB of 3.3 volt PC-100 ECC unbuffered CAS 2 latency SDRAM. The memory is embedded into the adapter's PCB (printed circuit board) and is not upgradeable. It is connected directly to the memory controller interface bus of the IOP and serves as storage for the executable code transferred from the flash memory. It also serves as the controller cache memory. The IOP memory controller provides single-bit ECC error correction.

2.2.4 PCI Interface

The adapter has a 64-bit/66 MHz PCI interface. Although designed for this interface specification, it is compatible with both 32-bit/33 MHz and 64-bit/33 MHz PCI interfaces. The PCI interface is universally keyed for 3.3v and 5v slots and is PCI 2.2 specification compliant. The RAID controller requires 3.3 volts. (Note: the adapter requires a RAIDIOS enabled PCI slot on the system baseboard.)

2.2.5 Adapter Jumpers and Switches

Table 2. Jumper Settings

Jumper Block	Jumper Position	Definition
J2B1	Pins [1-2]	IOP Reset / FW Recovery Enabled
	Pins [2-3]	IOP Normal-Run / FW Update Enabled

J2B1 – IOP Mode Select jumper block: This jumper is used to place the IOP in reset, which enables the flash chip to be programmed to recover resident firmware (FW). This is only necessary if the content of the flash part is corrupted and needs to be erased and reprogrammed. The jumper is installed on pins 1 and 2 to enable this function. During normal operation and during normal firmware updates, the jumper is installed on pins 2 and 3 (or removed totally).

2.2.6 Diagnostic Features

The adapter provides an LED (Figure 1: DS2A1) to indicate SCSI bus I/O activity.

2.3 Architecture Features (HW)

Table 3. Hardware Architecture

Component	Features
I/O Microprocessor	The 80302 uses the 100 MHz Intel 80960JT-100 core with a 64-bit / 66 MHz internal bus, which can pump 528 MB/sec of data to and from the internal IOP peripherals.
Cache Memory	The 80302 IOP supports up to 512 MB of 64-bit 100 MHz ECC SDRAM operating at 66 MHz (the SRCMR utilizes 32 MB of embedded memory).
Flash Memory	3.3 V, 16 Mb (2 MB) flash memory chip is used to store the RAID firmware
I/O Interface (PCI)	PCI 2.2 compliant, universally keyed for 3.3 or 5 volt PCI slots
PCI Transfer Rate	528 MB/sec (Burst)
PCI Signaling	3.3 or 5 volt

2.4 Electrical Characteristics

Table 4. Electrical Specifications

Attribute	Measurements
Voltage Requirements	+5, 3.3 volts (all +/- 5% tolerance)
Power Consumption	+5 V @ 0.25 Amps = 1.25 W +3.3 V @ 2.5 Amps = 8.25 W

2.5 Environmental Specifications

Table 5. Environmental Specifications

Environmental Stress Test	Required Conditions
Thermal: Non-Operating	-40 °C (Celsius) to +70 °C
Humidity: Non-Operating	50% to 92% Relative Humidity, non-condensing at 25 °C to 50 °C for 120 hours
Operating Temperature & Voltage Shmoo	-5 °C to 60 °C, and +/- 6% V_{cc} Shmoo
Form Factor (physical dimensions)	Height: 56 mm (2.20 inches) Length: 167 mm (6.57 inches)

2.6 Supported Hard Drive Technology

The RAID adapter supports up to 15 SCSI devices per SCSI channel. It supports up to 15 hard disks drives (or 14 hard disks drives if one of the SCSI ID's is occupied by a SAF-TE processor) per channel of the SCSI controller (30 disk drives total for the SRCMR; assuming a dual-channel SCSI controller on the motherboard).

The adapter supports both Single-ended (SE) and Low Voltage Differential (LVD) hard disk drives but it is recommended that you use only one type of drive technology (SE or LVD) on any one channel at a time. Each is explained below. The RAID adapter is designed to optimally utilize an Ultra160 SCSI controller implementation on the motherboard; yet, it is backward compatible with older SCSI hard disk drive specifications.

NOTE

If both SE and LVD devices are attached to the same channel/bus, the entire bus must operate at the single ended mode speed of the slower device. See Table 6 for the maximum cable length distances that apply to each mode.

Table 6. Supported SCSI Hard Disk Drive Standards

SCSI Drive Standard	Speed (MB/Sec)	Bus Width (Bits)	Maximum Cable Length (meters)(1)		Adapter Maximum Devices per SCSI Channel
			Single-Ended	LVD	
SCSI-1 (2)	5	8	6	(3)	8
Fast SCSI (2)	10	8	3	(3)	8
Fast Wide SCSI	20	16	3	(3)	16
Ultra SCSI (2)	20	8	1.5	(3)	8
Ultra SCSI (2)	20	8	3	NA	4
Wide Ultra SCSI	40	16	NA	(3)	16
Wide Ultra SCSI	40	16	1.5	NA	8
Wide Ultra SCSI	40	16	3	NA	4
Ultra2 SCSI (2, 4)	40	8	(4)	12	8
Wide Ultra2 (4)	80	16	(4)	12	16
Ultra160 (Ultra3)	160	16	(4)	12	16

NOTES:

1. May be exceeded in Point-to-Point and engineered specific applications.
2. Use of the word “Narrow,” preceding SCSI, Ultra SCSI, or Ultra2 SCSI (for example, Narrow SCSI) is optional.
3. LVD was not defined in the original SCSI standards for this speed. If all devices on the bus support LVD, then 12-meters operation is possible at this speed. However, if any device on the bus is single-ended only, then the entire bus switches to single-ended mode and the distances in the single-ended column apply.
4. Single-ended is not defined for speeds beyond Ultra.
5. After Ultra2, all new speeds are wide only.

The preceding information is based on the “STA-Endorsed Terms & Terminology for SCSI Parallel Interface Technology” specification published by the SCSI Trade Association. It can be viewed at: *http://www.scsita.org/terms/scsiterms.html*

2.7 Support for Non-Hard-Disk-Drive SCSI Devices (Non-Direct-Access Devices)

The RAID controller will pass through to the host operating system direct access to non-direct-access SCSI devices which are connected to a SCSI bus (channel) of the RAID controller. The RAID controller passes through all control of these devices to the host operating system.

NOTE

For MBOB adapters, the SCSI connectors are integrated into the system motherboard.

Types of supported non-Direct-Access SCSI devices (this doesn't cover specific vendors and models):

- SAF-TE Processors
- Tape Backups
- CD-ROMs
- Scanners

2.8 Other Hardware Design Considerations

2.8.1 Motherboard and System Design Guidelines

MROMB (often called Zero-channel RAID) is a unique RAID solution since it requires specific design considerations by the design engineer of the targeted computer system and motherboard. It is a hybrid between a true ROMB solution and a RAID add-in card, which makes it a platform-dependent solution. The SRCMR implementation of MROMB, due to its platform dependency and use of a regular PCI compliant slot, requires that the motherboard and system design engineers make certain considerations during their designing of the targeted platform.

3 Software

3.1 Software Architecture Overview

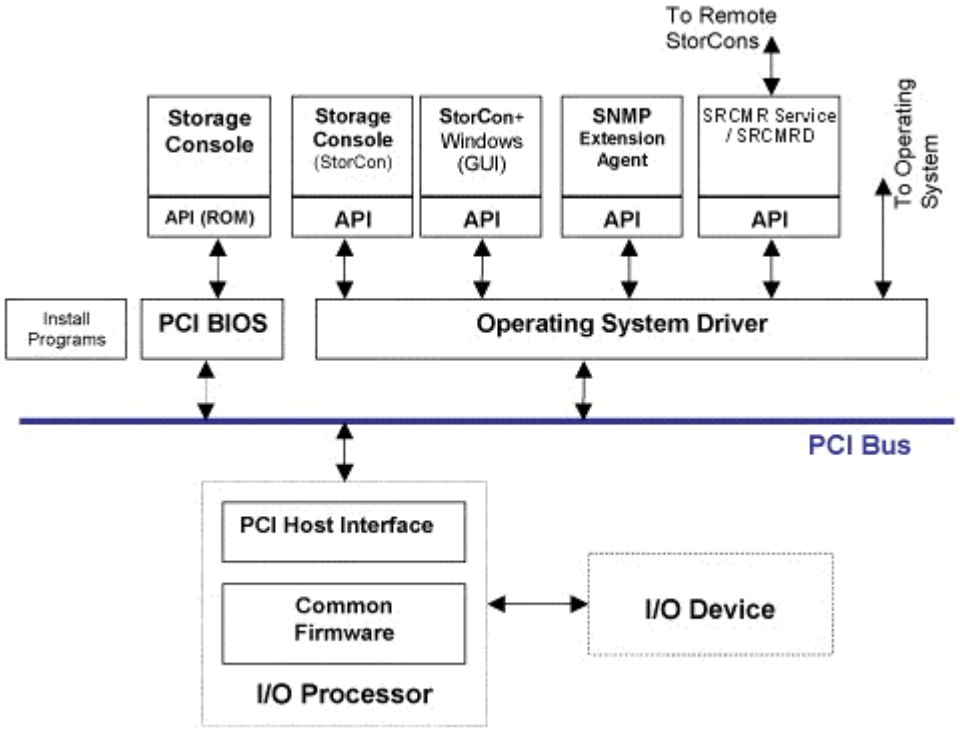

Figure 2. RAID Software Stack Architecture Block Diagram

The RAID software stack is composed of two major component groupings: the RAID firmware embedded in the Flash memory and a set of host resident drivers and utilities installed on the host system. All host-based software contains an OS-dependent portion and an OS-independent portion. This allows for a consistent "look and feel" across operating system platforms. A simple, custom messaging protocol is used to communicate between the host driver or utility and the embedded RAID firmware. The firmware is independent of the OS, I/O processor, and I/O bus through the use of abstraction layers. This layered RAID software executes on a custom, multi-tasking, real-time software executive and relies on the reuse of internal software communication "building blocks" to incorporate new technology and provide for new products.

NOTE

The architecture block diagram in Figure 2 is a generalization. Its goal is to cover all OS implementations. Certain blocks may or may not be relevant to each specific OS.

3.1.1 User Interface

3.1.1.1 Storage Console

The Storage Console is a text-based user interface. It is a full-featured monitoring and configuration utility for managing all aspects of the RAID subsystem as well as many features of the RAID adapter. It can be accessed via two methods. The first is during system boot time when entering the adapter's BIOS by depressing the <Ctrl> + <G> keys when prompted. The utility is accessing the RAID subsystem via the PCI BIOS.

The other method of accessing the Storage Console is via launching the application from within the host operating system.

Storage Console communicates with the firmware via a common API both during system POST and from within the host OS. This unique feature allows for a common UI between both OS and pre-OS environments.

3.1.1.2 Storage Console+

Storage Console+ is a GUI implementation of the Storage Console. However, it runs only on Microsoft Windows-based systems and has the same features and functionality with only a few exceptions.

3.1.2 System Management

3.1.2.1 SNMP Extension Agent

The SNMP Extension Agent is the interface between SNMP and the SRCMR Management Information Base (MIB). It interacts with the host resident MIB as well as the SRCMR HBA (via the SRCMR API) to respond to SNMP requests.

3.1.3 Common Layers

3.1.3.1 PCI BIOS

The PCI BIOS is the Expansion ROM software as defined in the PCI specification. It performs SRCMR initialization from host system memory during POST.

3.1.3.2 SRCMR API

The Intel® Integrated RAID (SRCMR) API is a C++ class library consisting of OS-independent classes and methods. This layer encapsulates sequences of lower level C library API functions and builds OS-independent data structures used for communicating with the HBA. This API relies on an OS-dependent layer that communicates with the local driver (if one is present) or a remote driver via the TCP/IP or IPX/SPX network protocol.

3.1.3.3 Operating System Driver

The SRCMR device driver is the OS specific driver that communicates between the host resident application and the SRCMR HBA using the SRCMR communications protocol.

3.1.3.4 RAID Configuration Service

The SRCMR Service / SRCMRD allows remote access to the SRCMR HBA. It runs as a service on Windows-based systems and as a daemon on Unix-based systems. In order to use the remote access capability, an administrator has to create user accounts. The passwords for these accounts are encrypted for security.

3.1.4 RAID Firmware

The SRCMR firmware is composed of multiple software layers allowing for maximum flexibility, re-use and maintainability. At the highest level is a host interface abstraction layer in the common firmware which is also composed of multiple internal layers of service and library modules. The firmware is common among all SRCMR products both internal and external allowing for a common disk data format across the entire SRCMR product line.

The I/O device pictured next to the firmware (Figure 2) also communicates over an abstraction layer allowing for elegant integration of multiple I/O device types from different vendors.

4 RAID Functionality and Features

4.1 Hierarchy

A fundamental purpose of a RAID system is to present a usable (with some level of redundancy) data storage medium (or drive) to a host operating system. In accomplishing this, the Intel RAID firmware is based on a four level hierarchal model. Each level has its "own drives" associated with it. The basic rule is: to build drives on a given level in the hierarchy, the "drives" of the next lower level are used as components. So, in order to construct and create a data drive (host drive/RAID volume) and present it to the host operating system, the following steps are typically followed by the RAID firmware to accomplish this:

1. One or more physical drives are selected and initialized.
2. A logical drive is created for each physical device.
3. The logical drives are grouped and an array drive is created.
4. The RAID firmware designates the array drive as a host drive and presents it to the host OS.

Level 1

"Physical drives" are located on the lowest level of the hierarchy. This includes hard disk drives, removable hard disks, and some MO drives. They are the basic components of all "drive constructions." However, before they can be used by the firmware, these hard drives must be "prepared" by a process called initialization. During initialization each hard disk has configuration information written to its physical medium in non-user accessible redundant areas. This information allows a univocal identification even if the SCSI-ID or the controller is changed. For reasons of data coherency, this information is extremely important for any drive construction consisting of two or more physical drives.

Level 2

"Logical drives" are constructed to obtain full independence of the physical coordinates of a physical device. This is important because it allows one to rearrange the SCSI-IDs and channel location of the physical drives of a disk array without compromising the integrity of the RAID array disk. It also is what allows one to interchange disk arrays between compatible controllers. Logical drives always consist of physical drives.

The user can create a logical drive manually by using one or more available physical drives. A logical drive created manually in this way is presented directly to the host OS as a Host drive (see host drives in Level 4). However, it is not associated with any array drive (see array drives in Level 3).

Level 3

"Array drives" are located at this level in the hierarchy. Array drives always consist of logical drives and consist of the following drive types:

- RAID 0 drives
- RAID 1 drives
- RAID 4 drives
- RAID 5 drives
- RAID 10 drives

The user can manually create an array drive by using two or more logical drives that were manually created at level 2. This action combines the original host drives associated with each of the logical drives into a single host drive at a chosen RAID level.

Level 4

"Host drives" (RAID volumes) are created at the highest level of the hierarchy by the RAID firmware. This is done automatically upon the creation of an array drive. It is also done automatically upon the manual creation of a logical drive. Each host drive is assigned a drive number that matches the drive number of its array drive or logical drive. Host drives are the only hierarchal level drives that are detectable by the host operating system. The three lowest level hierarchal drives are transparent to the host operating system.

After a capacity expansion of a given array drive, the added capacity appears to the host OS as a new host drive on this level. It can then be configured as a separate host drive or, using the "Merge" feature in the "Host Drive" menu, be combined with the initial host drive of the array drive into one single host drive.

WARNING

Any data located on the original host drive will be lost using the "Merge" feature. Do not use this feature if the initial host drive already contains data that you do not wish to lose.

Within the Storage Console utility, each level of hierarchy has its own special menu:

- Level 1 → Menu: Configure Physical Devices
- Level 2 → Menu: Configure Logical Drives
- Level 3 → Menu: Configure Array Drives
- Level 4 → Menu: Configure Host Drives

4.1.1 RAID Host/Array Drive Status

Table 7 lists the available states of RAID Host and Array Drives.

Table 7. Array Drive Status

Drive Status	Attribute	Drive Type	Redundant	Description
Idle	RW	RAID 4 / 5 / 10	No	Newly defined array prior to build process starting
Build	RW	RAID 1 / 4 / 5 / 10	No	The initial process of configuring redundancy information upon creation of a drive
Ready	RW	RAID 1 / 4 / 5 / 10	Yes	The array drive is fully operational; normal state
Rebuild	RW	RAID 1 / 4 / 5 / 10	No	Array drive assumes this status after the automatic activation of Host Fix or after manual replacement (hot-plug)
Expand	RW	RAID 4 / 5	Yes	This status indicates that the RAID level and or capacity is (are) being migrated
Fail	RW	RAID 1 / 4 / 5 / 10	No	This status indicates that a “logical drive” failed
Error	RO	RAID 4 / 5 / 10	No	If a component of a disk array is missing during startup and “fail mode” is not activated, the array is set to read only
OK	RW	RAID 1	Yes	The RAID 1 array drive is fully operational; normal state
Patch	RW	RAID 4 / 5	Yes	This status indicates that the array drive has gone through a significant procedure or has been patched from the “error” status to “fail” status
RW= Read Write RO= Read Only			The drive “attribute” indicates the access level that the host OS has to the drive	

4.1.2 Logical Drive Status

Table 8. Host Drive Status

Drive Status	Attribute	Description
OK	RW	The drive is operational and functioning normally
Missing	-	The drive is missing or no longer detected by the RAID controller
Fault	RW	The drive is no longer operating with expected parameters
RW= Read Write RO= Read Only		The drive “attribute” indicates the access level that the host OS has to the drive

4.1.3 RAID Controller Drive Limitations (Host, Array, Logical, and Physical)

The following are limitations assuming the following:

- U160 speed
- Dual-channel SCSI controller
- Cabling that meets U160 specifications

Physical drives are limited by the number of SCSI channels being controlled by the RAID controller. The firmware/software supports a maximum of 15 hard disk drives per channel (or 14 if one SCSI ID is being occupied by an intelligent enclosure processor).

The maximum number of array drives is limited to 35 by the RAID firmware. The actual maximum limit of the SRCMR RAID controller is 15 (see Table 9). The firmware supports channel spanning where an array can consist of physical drives that are attached to either one or to both channels of the RAID controller. An array drive requires a minimum of two hard disk drives (or logical drives). Therefore the maximum array limitation for each RAID controller is the physical drive limit of that RAID controller divided by two. An array drive can contain (or have reside on it) up to a maximum of 2 host drives.

Host drives are limited by the RAID firmware to a maximum of 30 host drives (RAID volumes) per controller. There is a minimum of at least one host drive per array drive. A host drive can only be associated with (or reside on) a single array drive. Like array drives, host drives can reside on hard disk drives that are isolated to a single channel or span both of the channels of the RAID controller.

Table 9. RAID Controller Drive Maximum Limitations (SRCMR)

Drive Type	Per Channel	Per Controller	Per Array Drive	Per Host Drive
Physical Disk Drives	15	30	30 (2 min)	30
RAID Array Drives	7	15	-	1
RAID Host Drives	-	30	2	-

4.2 Utilities and Tools

Table 10. Utilities and Tools

Management/ Monitoring	Description	Actions
Storage Console (StorCon)	This is a text-based UI that allows full management and monitoring of the RAID controller and its subsystem; consistent look and feel across all supported operating systems	Can be launched during boot up, <Ctrl> + <G>, or within host OS Can run locally or remotely using TCP/IP or SPX/IPX network protocols
Storage Console Plus (StorCon+)	GUI based version of the Storage Console with exception to a few functions; it only runs locally on Windows servers and clients	Can be launched in normal or admin modes Can run locally or remotely using TCP/IP or SPX/IPX network protocols
Management/ Reporting	**Description**	**Actions**
View Statistics	Allow viewing of I/O activity of host, logical, and physical drives and cache activity	Accessed via the Storage Console menu: Express or Advanced Setup → View Events. User can adjust synch rate and enable and disable logging
View Events	Allows viewing and saving of all events regarding the adapter since its last boot up	Accessed via the Storage Console menu: Express or Advanced Setup → View Events

continued

Table 10. Utilities and Tools

Management/ Monitoring	Description	Actions
View Hard Disk Information	Lists all hard drives connected to the RAID adapter with information about each hard drive	Accessed via the Storage Console tools. User can select each drive individually to bring up a detailed list of information that includes detected defects since the drive was first detected by the RAID adapter
Diagnostics	**Description**	**Actions**
Save Information	Gives the administrator the ability to save the configuration information of the RAID adapter in an ASCII file for viewing	Accessed via the Storage Console. User can save file to be used for troubleshooting or documentation.
Memory Test	Non-destructive tests are written to ECC memory and verified. The different modes determine type of pattern and quantity. Tests are not non-destructive for non-ECC memory.	Accessed via Storage Console menu: Advanced Setup → Configure Controller → Controller Settings. User selected options are: "No test," "Standard," "Double Scan," and "Intensive."
Parity Verify	Selecting this feature causes the adapter to verify the parity on RAID 4 and 5 arrays	Accessed via Storage Console menu: Advanced Setup → Configure Array Drives → Select an Array
Check Surface: Physical Disks	A surface-check of the selected hard disk drive done. The controller writes and reads certain data patterns and checks them for correctedness.	Accessed via Storage Console menu: Advanced Setup → Configure Physical Devices → Select Physical Drive → Check Surface Note: This action destroys all data on the selected drive.
Firmware	**Description**	**Actions**
Firmware Update	Utility that allows the updating of the adapter's firmware using: 1. DOS utility for major and minor updates (for example, major: 2.32.xx to 2.33.xx) 2. StorCon for minor updates (for example, minor: 2.32.xy to 2.32.xz)	1. Accessed via DOS utility. 2. Accessed via StorCon during boot up (<Ctrl> + <G>) or from with the host OS menu: Advanced Setup → Configure Controller → Firmware Update.
Firmware Recovery	This utility allows the recovery of corrupted firmware or reprogramming of entire FLASH code	Requires the user set the adapter's J2B1 jumper to "Reset/Firmware Recovery" mode and the use of the DOS firmware recovery utility

4.3 RAID Features

4.3.1 RAID Level Support

Table 11. Supported RAID Levels

RAID Level	Description	Configuration Drives Requirements
0	Data striping	2 min, 30 max (2 channel limit)
1	Drive mirroring	2 min, 30 max (2 channel limit)
4	Data striping with dedicated parity drive	3 min, 30 max (2 channel limit)
5	Data striping with distributed parity	3 min, 30 max (2 channel limit)
10	Combination RAID 0 and 1, striped mirrors	4 min, 30 max (2 channel limit; drives must be added in pairs)
Disk Pass-Through to Host		
Single Disk	Presented to host OS as a host drive	1
Chaining	Presented to host OS as a host drive	2 min, 30 max (2 channel limit)

4.3.2 Caching

There are two levels or modes of caching related to the adapter - each independent of the other. Caching can be enabled on the controller, which sets caching on all the RAID array/host drives configured on the adapter. This mode of caching utilizes the 32 MB of caching memory embedded on the adapter. The other caching mode is enabling the caching feature of the hard disk drives. In this method, the cache memory that is utilized is located on the disk drive and does not use the adapter's memory.

NOTE

In the event of power loss to the computer system, data located in the disk cache is not protected by the battery backup unit (of those controllers that support optional battery backup). The battery backup only protects data that is located in the controller cache.

Controller cache selections:

- Caching on/off selectable (Read Ahead)
- Delayed Write on/off selectable (Write Back)

Physical Disk cache selections:

- Read cache on/off selectable
- Write cache on/off selectable

Table 12. Supported Cache Settings

Disk Cache	Controller Cache		Cache Configuration
Write _ Read _	Cache _	Delayed Write •_	No cache
Write _ Read •	Cache _	Delayed Write •_	Disk Read
Write • Read _	Cache _	Delayed Write •_	Disk Write
Write • Read •	Cache _	Delayed Write •_	Disk Read Write
Write _ Read _	Cache •	Delayed Write _	Controller Read
Write _ Read •	Cache •	Delayed Write _	Disk Read, Controller Read
Write • Read _	Cache •	Delayed Write _	Disk Write, Controller Read
Write • Read •	Cache •	Delayed Write _	Disk Read Write, Controller Read
Write _ Read _	Cache •	Delayed Write •	Controller Read Write
Write _ Read •	Cache •	Delayed Write •	Disk Read, Controller Read Write
Write • Read _	Cache •	Delayed Write •	Disk Write, Controller Read Write
Write • Read •	Cache •	Delayed Write •	Disk Read Write, Controller Read Write
Legend	• enabled / on _ disabled / off •_ either		

4.3.3 Hot Fix (Spare) Disk Drives

There are two types of hot fix drives:

- Private (dedicated): This type of hot fix drive is assigned to a specific RAID 1, 4, 5, or 10 array drive. It cannot be used by any other RAID array drive configured on the controller.
- Pooled (global): This type of hot fix drive is available for any RAID 1, 4, 5, or 10 array drive that has be configured with “Pooled Hot Fix Access” enabled.

NOTE

When adding a Pool Hot Fix to a specific array, the access for this array will be automatically enabled. For other arrays this feature has to be manually activated with the Pool Hot Fix Access switch.

The capacities of Hot fix drives are required to be the same size or larger than the capacity of the smallest physical disk drive in the RAID array drives that they are protecting. Therefore, if you have two RAID arrays where “Array_1” has all 9 GB drives and “Array_2” has all 18 GB drives then “Array_1” would require, at a minimum, a 9 GB “hot fix drive” and “Array_2” would require, at a minimum, an 18 GB “hot fix drive.”

Also, when using a single “pooled hot fix drive” to protect several RAID array drives, the “pooled hot fix drive” must meet the proper capacity requirement to protect all of the arrays. To select the proper sized “pooled hot fix drive,” first determine the capacity of the smallest physical disk drive in each RAID array to be protected. Next, select a “pooled hot fix drive” that is equal to or larger than the capacity of the largest of these disk drives. For example, in Table 13, if the pooled hot fix drive for the four RAID arrays is 9 GB, then it would only protect against 1 and 2. Therefore, the proper hot fix drive selection to protect all four of the array drives would have to be of a minimum capacity of 36 GB (see note below Table 13).

Table 13. Example of a 9GB Pooled Hot Fix Drive Selected to Protect Four RAID Array Drives

Array (smallest disk)	Array_1 (4GB)	Array_2 (9GB)	Array_3 (18GB)	Array_4 (36GB)
Protected by Pooled Hot Fix?	Yes	Yes	No	No

NOTE

This is only an example of how the pooled hot fix drive feature works. It is not meant to represent any particular or practical configuration. Obviously it would not be practical to use only one pooled hot fix drive to protect all four of the arrays in this example unless hardware configuration limitations only allowed for one extra drive as a pooled hot fix (that is to say that adding the pooled hot fix drive brings the total number of hard disk drives to the maximum supported by the controller).

4.3.4 Hot-Plug Disk Drive Support

This feature allows the ability to remove and replace SCA (single connect adapter) drives while I/O activity is taking place on the same SCSI bus, provided that both the hard disk drive and backplane fully support hot-swap, without interruption of operations on any other drives. It supports the hot-plug of new drives in both intelligent (SAF-TE and SES) and non-intelligent enclosures.

4.3.5 Non-Intelligent Drive Enclosure Auto-detection of Host Plug Disk Drives

This is a special feature that allows the use of non-intelligent disk enclosures (requires truly hot-plug disk drives and backplane connectors) as though they were intelligent enclosures. This feature is configurable and allows the user to set up non-intelligent enclosures to detect the insertion or removal of hot-plug disk drives and report the event to the RAID firmware. The RAID configuration is automatically updated to the new configuration. Access this feature through the "Advanced Setup" menu of Storage Console.

4.3.6 Auto-Declare Hot Fix (Spare) Drive

If the RAID controller has a RAID array drive that is in "failed (degraded) state," and you connect to the controller a new hard disk drive that is the same size or larger than the smallest disk drive in that "failed (degraded)" RAID array then the RAID firmware will automatically mark this new disk drive as a hot fix (spare) drive for the "failed (degraded) RAID array drive." "Rebuild" will then automatically commence. If the new hard disk is smaller than the smallest hard disk drive in the "failed (degraded) RAID array drive," the new disk drive will not be marked as a spare and the "failed (degraded) RAID array drive" will remain "failed (degraded)." With an intelligent SAF-TE enclosure (or a non-intelligent enclosure that is configured to auto-detect insertion of hot-plug disk drives), a bus scan occurs automatically when the hard disk drive is inserted. No manual intervention is required. When a non-intelligent enclosure that has not been configured for auto-detection of hot-plug disk drives is used, a bus scan or reboot is required for the "Auto Declare Hot Fix" feature to commence an auto-rebuild.

A new hard disk drive may have the same or a different SCSI ID as the failed hard disk drive that it replaces.

4.3.7 RAID Array Drive Roaming

Array Roaming allows the user the ability to move a complete RAID array from one computer system to another computer system and preserve the RAID configuration information and user data on that RAID array. Compatible RAID controllers (another Intel RAID Controller SRCMR) must control the RAID subsystems of the two computer systems. The transferred RAID array may be brought online while the target server continues to run if the hard disk drives and disk enclosure support hot-plug capabilities. The hard disk drives are not required to have the same SCSI ID in the target system that they did in the original system that they are removed from. The RAID array drive that is being roamed must not be of type "*Privat*e." This includes all non-private "host," "array," and "logical drives."

WARNING

Do not attempt RAID Array Drive Roaming between RAID controllers that are not compatible with the SRCMR. Unpredictable behavior may include, but is not limited to, data loss or corruption.

4.3.8 Online RAID Array Configurations

4.3.8.1 Capacity Expansion Without Reboot

Online capacity expansion refers to the ability of the RAID controller to present new storage space to the host OS without requiring that the computer system be taken off-line and rebooted (for those operating systems that support this feature). The host OS is able to detect the new capacity and format and partition it for immediate use. The RAID controller offers several ways of creating additional capacity while online.

Creating new RAID array drives from an available physical hard disk drive:

This is the simplest method. The user just creates a new "host drive" from available physical disk drives that are connected to the RAID controller (the assumption here is that there were disk drives that were already connected but unused or there were new disk drives inserted into open slots in a hot-plug drive enclosure already connected to the RAID controller).

Creating new capacity to existing RAID array components (Expand Array Drive):

Using this method you are allowed to do one of the following:

- Convert unused free space on the existing logical drives of the RAID array drive into a separate host drive, or
- Add additional physical disk drives to the existing RAID array drive.

4.3.8.2 RAID Level Migration

RAID level migration is accomplished using the "Expand Array Drive" feature. To expand a RAID 0 drive to a RAID 4 or 5, requires that you add at least one additional drive to the array. Migrating from RAID 4 or 5 to RAID 0 frees one disk. The RAID controller allows the following RAID level migrations of a given array drive:

1. RAID 0 → RAID 4: add new disk, parity is calculated and written to new disk
2. RAID 0 → RAID 5: add new disk, parity calculated and written to new disk, parity is distributed over all disk.
3. RAID 5 → RAID 0: parity written to one disk, parity disk freed and removed
4. RAID 4 → RAID 0: parity disk freed and removed
5. RAID 4 → RAID 5: parity distributed over all disk
6. RAID 5 → RAID 4: parity written to one disk

Basic migration process: RAID 0 ←→ RAID 4 ←→ RAID 5

To initiate a migration the RAID array must have a "ready" status.

4.3.8.3 Data Strip Size Configurable per RAID Array

The strip size for each RAID array can be configured at the time of creation of the RAID array. This is a one-time configuration and cannot be changed or migrated once the array has been created. For RAID 0, 4, 5, and 10 arrays the following strip sizes are possible:

- 16 KB
- 32 KB
- 64 KB
- 128 KB

4.3.9 Background Initialization and Instant Availability

The initialization of RAID array drives is done in the background. Array drives have a status of build during this process and are immediately accessible to the host OS if the host OS supports online capacity expansion.

Three are two build modes available when creating RAID arrays, "Standard" and "Destructive." Destructive is much faster than the standard build mode. When in destructive build mode, the firmware writes a pattern of zeros across all disks. If this build process is interrupted by rebooting the computer, the build process will continue in the much slower "standard non-destructive" build mode. The destructive build mode is only available from within the BIOS version (<Ctrl> + <G>) of the Storage Console. For the SRCMR controller, background initialization does not continue during POST until after the OS driver is loaded.

4.3.10 Configurable SCSI Parameters

When a SCSI hard disk drive is initialized the first time by the RAID controller, its SCSI parameters are automatically set to their optimal settings. Manual configuration is not required. However, the RAID controller allows for the custom configuration of several SCSI parameters on a hard_disk_drive-by-hard_disk_drive basis. There are several settings that can be configured by using the Storage Console menu Advanced Setup → Configure Physical Devices → Select Physical

Drive → SCSI Parameter/Initialize (for SRCMR, most settings are set automatically and can not be configured manually).

Table 17. Disk Drive SCSI Parameters

Parameter	Setting/Value	Description
Synch. Transfer	Not Configurable	This setting, when enabled, allows the controller to operate in synchronous transfer mode.
Synch. Transfer Rate	Not Configurable	Allows for setting of the speed for the SCSI hard disks (160 MB/sec for U160 drives). No matter the setting, the SCSI bus will negotiate the fastest speed up to this setting. Lowering the setting will force the disk drive to transfer at the lower speed.
Disconnect	Not Configurable	Enabling this setting allows for the disk drive to disconnect from the SCSI bus when it's not participating in a transfer. This allows for optimal bus utilization by all devices on the bus.
Tagged Queues	Not Configurable	When enabled, this feature allows the SCSI disk drive to execute more than one command at a time.
Disk Read and Write Cache	On / Off	For performance reasons, the "Read Ahead" and "Write" cache of the hard disk drives should always be "On."
Domain Validation	Not Configurable	Using the <F4> key while in this menu accesses this parameter. When this is set to "On," Domain Validation allows for a cyclical check of the correct data transfer at a given rate.

5 Certifications and Supported Technologies

5.1 OS Certifications

The product will be validated with the latest vendor OS certification test suites. Pre-submission tests will be passed and the certifications listed in Table 15 will be submitted to the proper submission process as required per OS. The pre-submission test will be run on the final gold production release candidate of the RAID software suite OS drivers and RAID firmware. The product will not be held up from shipping while awaiting final passing notification from the OS vendors (and in the case of Windows 2000, the digitally signed versions of the OS driver).

Table 15. OS Certification Requirements

OS Vendor	Details	Test Suite Version
Microsoft	The product shall be WHQL Certified as a RAID Adapter for the following: • Windows 2000 Advanced Server • Windows NT 4.0 Server Enterprise Edition	HCT ver. 9.502 or latest available
Novell	The product shall be certified to receive “Yes Tested and Approved” with the Storage Access Tests for the following: • NetWare 5.1 with clustering • NetWare 6.0 with clustering (pending release)	Test kit ver. 3.1 or latest available
Caldera	The product shall be certified to receive “Works with Caldera” certification for the following: • UnixWare 7.1.1 • OpenServer† 5.0.6	• PLT ver. 8.13 and HBACert ver. 7.1 or latest available • PLT ver. 8.13 and O5HBACert ver. 7.1 or latest available
Red Hat Linux	The product shall be certified to receive “redhat READY” certification for the following: • Red Hat Linux 7.0 (the 2.2 kernel) • Red Hat Linux 7.1 (the 2.4 kernel)	• Test suite ver. 1.5.9 or latest available • Test suite ver. 1.6.9 or latest available

5.2 Product Regulation Compliance Information

This product will comply with the following Product Safety, EMC, and Immunity requirements.

5.2.1 Product Safety Compliance

- UL 1950 - CSA 950 (US/Canada)
- EN 60 950 (European Union)
- IEC60 950 (International)
- 73/23/EEC - Low Voltage Directive (European Union - CE)
- EMKO-TSE (74-SEC) 207/94 (Nordics)
- AS/NZS 3562 (New Zealand)

5.2.2 Product EMC / Immunity Compliance

- Title 47 CFR Part 15 (Class A) (USA)
- ICES-003 (Class A) (Canada)
- CISPR 22, 3rd Edition (Class A) & CISPR 24 (International)
- EN 55022:1994 (Class A) & EN55024: 1998 (European Union)
- AS/NZS 3548:1995 (Class A) (Australia/New Zealand)
- CNS 13438 (Class A) (Taiwan)
- MIC Notices No. 1997-42 (Class A) & 1997-41 (Korea)

5.3 Product Regulatory Compliance Markings

This product will be marked with the following Product Certification / Regulation Markings pending room availability:

	USA / Canada – UL Recognition Mark
	European CE Directive Mark
CANADA ICES-003 CLASS A CANADA NMB-003 CLASSE A	Industry Canada Class A Marking
 N232	Australian C-tick Mark
檢磁 3902I918	Taiwan BSMI Certification Number
警告使用者： 這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策	Taiwan BSMI Class A EMC Warning
	Korean RRL MIC Mark (Logo)

5.4 Regional EMC Compliance Information

5.4.1 USA – FCC Verification Notice (Class A)

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

For questions related to the EMC performance of this product, contact:

Intel Corporation
5200 N.E. Elam Young Parkway
Hillsboro, OR 97124
1-800-628-8686

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

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and the receiver.
- Connect the equipment to an outlet on a circuit other than the one to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

5.4.2 CANADA – INDUSTRY CANADA (Class A)

This Class A digital apparatus complies with Canadian ICES-003.

Cet appereil numérique de la classe A est conforme à la norme NMB-003 du Canada.

5.4.3 EUROPE – CE Declaration of Conformity

This product has been tested in accordance too, and complies with the European Low Voltage Directive (73/23/EEC) and European EMC Directive (89/336/EEC). The product has been marked with the CE Marking to illustrate its compliance

5.4.4 TAIWAN – BSMI Certification Information

BSMI Certification Number is:

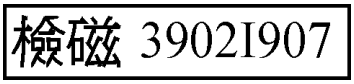

BSMI EMC Warning:

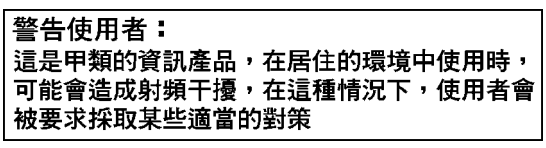

5.5 Supported Specifications and Standards

Table 17. Supported Standards and Specifications

Specification/Standard	Details
SAF-TE specification 1.0	The product supports SAF-TE Backplane processors including: processing/reporting/logging of messages and issuing of commands per SAF-TE specification. *http://www.nstor.com/support/whitepapers/safte_spec.doc*
SES	The product supports SES Backplanes to include: processing/reporting/logging of messages and issuing of commands per SES specification. *http://www.nowhere.net/~raster/SCSI3/ses-r08a.pdf*
ACPI Power Management (Windows 2000)	• The product supports power states D0 and D3 as defined by Storage Devices Class Power Management Reference Specification, Microsoft, Version 1.0A, February 1997. • The product supports power state D3 as follows: — Drive controller (interface and control electronics): not functional; context lost — Interface mode (communications timings): not preserved — Drive motor (spindle): stopped
Bootable CD-ROM	• The product supports booting from a CD-ROM based on the “El Torito” Bootable CD-ROM Format Specification, version 1.0, *http://www.phoenix.com/PlatSS/PDFs/specs-cdrom.pdf* • The product supports a bootable CD-ROM from a PCI 2.2 Plug and Play compliant BIOS. • The product supports “No Emulation” mode with PnP BIOS.

6 Technical Drawings and Diagrams

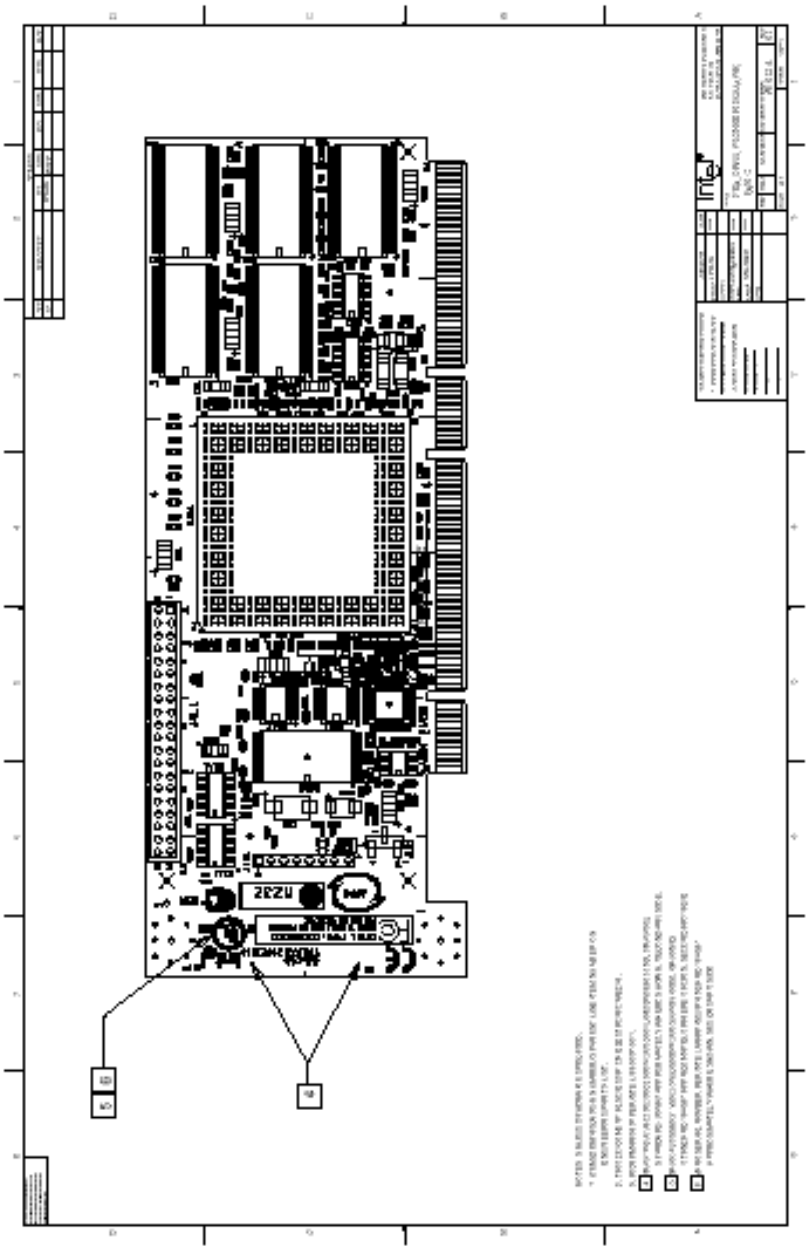

Figure 3. Adapter PBA Drawing

7 Appendices

7.1 Referenced Documentation

Table 18. Referenced Documentation

Document Title	Order #
Intel RAID Controller SRCMR Installation	A75270-001
Intel RAID Controller SRCMR User's Guide	A75271-001