



# Intel<sup>®</sup> Integrated RAID Controller SRCZCR

Low profile U160/U320 MROMB Controller  
Technical Product Specification, Version 1.1.5

---

November 2002



Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Intel® Integrated RAID may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com>.

Copyright © Intel Corporation, 2002

\*Other names and brands may be claimed as the property of others.

## Contents

---

1	Introduction.....	5
1.1	Acronyms .....	5
1.2	Documentation Conventions .....	6
1.3	Product Overview .....	6
1.4	Operating System Support.....	6
1.5	List of Features.....	7
2	Hardware.....	8
2.1	Physical Layout .....	8
2.2	Major Components .....	8
2.2.1	Intel® 80303 I/O Microprocessor .....	8
2.2.2	Intel Smart 3 FlashFile™ Flash Memory .....	9
2.2.3	SDRAM (Cache).....	9
2.2.4	PCI Interface .....	10
2.2.5	Motherboard Requirements.....	10
2.2.6	Jumpers and Switches .....	10
2.2.7	Diagnostic Features .....	10
2.3	Architecture Features (HW).....	11
2.4	Electrical Characteristics.....	11
2.5	Environmental Specifications .....	11
2.6	Supported Hard Drive Technology .....	12
2.7	Support for Non-Hard-Disk-Drive SCSI Devices (Non-Direct-Access Devices) ..	13
3	Software .....	14
3.1	Software Architecture Overview .....	14
3.1.1	User Interface .....	15
3.1.2	System Management .....	15
3.1.3	Common Layers .....	15
3.1.4	RAID Firmware.....	16
4	RAID Functionality and Features .....	17
4.1	Hierarchy .....	17
4.1.1	RAID Host/Array Drive Statuses .....	19
4.1.2	Logical Drive Statuses.....	20
4.1.3	RAID Controller Drive Limitations (Host, Array, Logical, and Physical) .	20
4.2	Utilities and Tools .....	22
4.3	RAID Features.....	23
4.3.1	RAID Level Support.....	23
4.3.2	Caching .....	23
4.3.3	Hot Fix (Spare) Disk Drives.....	25
4.3.4	Hot Plug Disk Drive Support.....	26
4.3.5	Auto-detection of Hot Plug Disk Drives in Non SAF-TE Drive Enclosures. 26	
4.3.6	Auto-declare Hot Fix (Spare) Drive .....	26
4.3.7	RAID Array Drive Roaming .....	26

	4.3.8	Online RAID Array Configurations.....	27
	4.3.9	Background Initialization and Instant Availability (RAID 4 and 5).....	28
	4.3.10	Configurable Disk Drive SCSI Parameters.....	29
	4.3.11	Configurable I/O Controller SCSI Parameters.....	29
5		Certifications and Supported Technologies .....	30
	5.1	OS Certifications .....	30
	5.2	Electronic Regulatory Agencies Certifications (Hardware) .....	31
	5.3	Supported Specifications and Standards .....	33
6		Technical Drawings and Diagrams .....	34
7		Appendices .....	35
	7.1	Referenced Documentation .....	35
	7.2	BIOS Boot Messages.....	35

# Introduction

# 1

## 1.1 Acronyms

The following acronyms are used throughout this document.

**Table 1-1. Acronyms (Sheet 1 of 2)**

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 <i>error correcting code</i> and <i>error correcting circuits</i> )
FRU	Flash Recovery Utility
FUU	Flash Update Utility
FW	Firmware
HBA	Host Bus Adapter
IIR	Intel Integrated RAID
IOP	I/O Processor, the Intel® 80303 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

Table 1-1. Acronyms (Sheet 2 of 2)

Acronym	Description
StorCon	Storage Console: A character-based, menu-driven tool used for setting up, monitoring and maintaining mass storage device subsystems based on IIR Controllers.
StorCon+	Storage Console Plus: A GUI-based, menu-driven tool used for setting up, monitoring and maintaining mass storage device subsystems based on IIR 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” and “the controller” are used interchangeably throughout this document. Each term represents the physical PCB that integrates all the components of the SRCZCR RAID PCI add-in card.

## 1.3 Product Overview

The SRCZCR is a low profile, 64bit/66MHz PCI, Modular RAID on Motherboard (MROMB) add-in controller. 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.

The major components of the SRCZCR RAID controller are:

- the Intel® 80303 I/O Processor
- 32MB (megabytes) of integrated SDRAM memory
- a 16Mb (2MB) Flash memory unit

See [Section 2.2.5](#) for additional motherboard requirements, including supported SCSI controllers.

## 1.4 Operating System Support

The following operating systems are fully validated and supported:

- Microsoft\* Windows\* 2000 Advanced Server, Service Pack 2
- Novell\* NetWare\* 6.0
- SCO\* Openserver\* 5.0.6a
- Red Hat\* Linux\* 7.2 (2.4 kernel)
- SCO UnixWare\* 8.0
- SuSe Linux 7.3 Professional

The following operating systems are supported with limited compatibility validation:

- Windows 2000 server
- Windows NT 4.0 Server and Enterprise Server, Service Pack 6a or higher
- Novell NetWare 5.1 (support pack 2a or higher)
- Linux (2.4 kernel)
  - Red Hat Linux 7.1
  - Debian\* Linux 2.2r6
  - Caldera Linux 3.1
  - Mandrake\* Linux 8.1
  - TurboLinux\* 7.0 Server
- SCO UnixWare 7.1.1
- FreeBSD\* 4.4 and 4.5
- Linux (2.2 kernel)
  - Red Hat Linux 7.0

## 1.5 List of Features


- Supports RAID levels 0, 1, 4, 5, and 10
- Supports selected Ultra160 or Ultra320 SCSI controllers. See [Section 1.3](#) for more information.
- 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)
- 32MB 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

## Hardware

## 2

### 2.1 Physical Layout

Figure 2-1. RAID Controller Physical Layout



### 2.2 Major Components

#### 2.2.1 Intel® 80303 I/O Microprocessor

The controller features the Intel 80303 I/O processor (IOP). The major components of the 80303 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. For more detailed information on the Intel 80303 IOP, visit the 80303 IOP's website at <http://developer.intel.com/design/iio/docs/iop303.htm>.

##### 2.2.1.1 80960JT-100® core

The 80303 uses the 80960JT-100 core. The core processor runs at 100 MHz with an internal 64-bit 100 MHz PCI bus. Among other features, it contains a 128-bit register bus, 16Kbyte two-way instruction cache, 4Kbyte direct-mapped data cache, 1Kbyte zero wait state data RAM, and single clock execution of most instructions.



### 2.2.1.2 PCI-to-PCI Bridge

This feature is not utilized by the SRCZCR controller.

The PCI-to-PCI bridge features fully independent PCI bus operation with independent clocks, dedicated data queues, 32-bit/33Mhz and 64-bit/66Mhz PCI bus support, and 64-bit Dual Address Cycle addressing.

### 2.2.1.3 Memory Controller Unit

The Memory Controller provides direct control of memory systems external to the 80960 core processor, including SDRAM and Flash. It features programmable chip selects, a wait state generator, ECC single-bit correction and double-bit error detection. The bus interface to the memory controller operates at 100 MHz. The SRCZCR uses 32MB of embedded memory.

### 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 80303 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, 16Mb (2MB) 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 2-1: J1B1](#)) 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 SDRAM (Cache)

The controller provides 32MB of 3.3 volt PC-100 ECC unbuffered CAS 2 latency SDRAM. The memory is embedded into the controller'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 controller has a 64bit/66MHz PCI interface. Although designed for this interface specification, it is compatible with both 32bit/33MHz and 64bit/33MHz PCI interfaces. The PCI interface is keyed for 3.3v slots and is PCI 2.2 specification compliant.

## 2.2.5 Motherboard Requirements

- **RAIDIOS-enabled expansion slot**—The SRCZCR controller requires a RAIDIOS enabled PCI slot on the motherboard as specified in the *Intel RAIDIOS Design Guide v. 1.3*, document order #273724-004. Other specifications, enabling circuits or modifications to the RAIDIOS specification (document order #273724-004) are not supported.
- **SCSI controller**—The SRCZCR requires one of the following compatible SCSI controllers on the motherboard:
  - SCSI Ultra160
    - Adaptec\* AIC-7899x dual-channel Ultra160 SCSI controller
    - Adaptec AIC-7892x single-channel Ultra160 SCSI controller
  - SCSI Ultra320
    - Adaptec AIC-7901 single-channel Ultra320 SCSI controller
    - Adaptec AIC-7902 dual-channel Ultra320 SCSI controller
- **SCSI connectors**—One LVD SCSI connector is required for each channel of the SCSI controller. Also, proper SCSI cabling is required to connect the RAID controller to the SCSI controller's channel.

## 2.2.6 Jumpers and Switches

**Table 2-1. Jumper Settings**

Jumper Block	Jumper Position	Definition
J1B1	Pins [1-2]	IOP Reset / FW Recovery Enabled
	Pins [2-3]	IOP Normal-Run / FW Update Enabled

**J1B1** – 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).

## 2.2.7 Diagnostic Features

The controller provides an LED ([Figure 2-1: DS2A1](#)) to indicate SCSI bus I/O activity.

## 2.3 Architecture Features (HW)

Table 2-2. Hardware Architecture

Component	Features
I/O Microprocessor	The 80303 uses the 80960JT-100 core. The core processor runs at 100 MHz with an internal 64-bit 100 MHz PCI bus. Among other features, it contains a 128-bit register bus, 16Kbyte two-way instruction cache, 4Kbyte direct-mapped data cache, 1Kbyte zero wait state data RAM, and single clock execution of most instructions.
Cache Memory	The 80303 IOP supports up to 512MB of 64bit 100 MHz ECC SDRAM operating at 66MHz. <i>The SRCZCR only utilizes 32MB of embedded memory.</i>
Flash Memory	3.3v, 16Mb (2MB) flash memory chip is used to store the RAID firmware
I/O interface (PCI)	RAIDIOS-enabled PCI 2.2 compliant
PCI Transfer Rate	528 MB/sec (Burst)
PCI Signaling	3.3 volts

## 2.4 Electrical Characteristics

Table 2-3. Electrical Specifications

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

## 2.5 Environmental Specifications

Table 2-4. Environmental Specifications

Environmental Stress Test	Required Conditions
Thermal: Non-Operating	Cycle: -40° C to 70° C for 50 hours Soak: 70° C for 22 hours
Humidity: Non-Operating	50% to 85% non-condensing at 25° C to 70° C for 168 hours
Operating Temperature & Voltage Shmoo	-0° C to 55° C, and +/- 5% Vcc 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 controller 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 SRCZCR; assuming a dual channel SCSI controller on the motherboard).

The controller 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 controller is designed to optimally utilize an Ultra160 or Ultra320 SCSI controller implementation on the motherboard; yet, it is backward compatible with older SCSI hard disk drive specifications.

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 2-5](#) for the maximum cable length distances that apply to each mode.

**Table 2-5. Supported SCSI Device Standards**

SCSI Drive Standard	Speed (MB/Sec)	Bus Width (Bits)	Maximum Cable Length (meters) <sup>(1)</sup>		Controller Maximum Devices per SCSI Channel
			Single-Ended	LVD	
SCSI-1 <sup>(2)</sup>	5	8	6	<sup>(3)</sup>	8
Fast SCSI <sup>(2)</sup>	10	8	3	<sup>(3)</sup>	8
Fast Wide SCSI	20	16	3	<sup>(3)</sup>	16
Ultra SCSI <sup>(2)</sup>	20	8	1.5	<sup>(3)</sup>	8
Ultra SCSI <sup>(2)</sup>	20	8	3	N/A	4
Wide Ultra SCSI	40	16	N/A	<sup>(3)</sup>	16
Wide Ultra SCSI	40	16	1.5	N/A	8
Wide Ultra SCSI	40	16	3	N/A	4
Ultra2 SCSI <sup>(2, 4)</sup>	40	8	<sup>(4)</sup>	12	8
Wide Ultra2 <sup>(4)</sup>	80	16	<sup>(4)</sup>	12	16
Ultra160 (Ultra3)	160	16	<sup>(4)</sup>	12	16
Ultra320	320	16	<sup>(4)</sup>	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 singled-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.

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)<sup>1</sup> of the RAID controller. The RAID controller passes through all control of these devices to the host operating system.

Types of supported non-Direct-Access SCSI devices (this does not cover specific vendors and models):

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

---

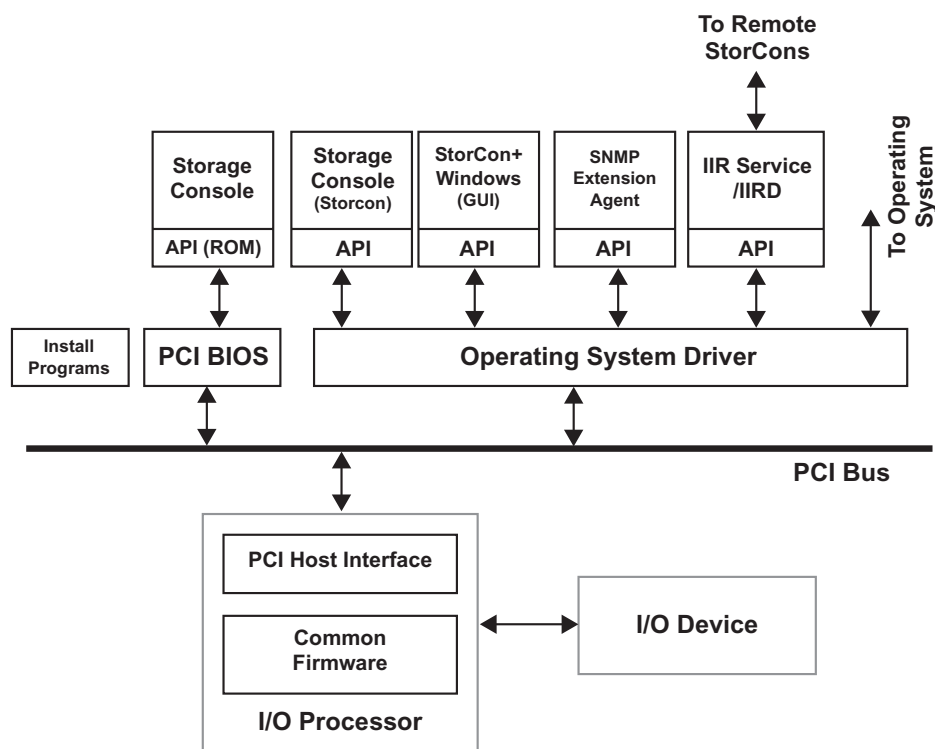
1. For MROMB controllers, the SCSI connectors are integrated into the system motherboard.

# Software

# 3

## 3.1 Software Architecture Overview

Figure 3-1. 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 3-1](#) 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 controller. It can be accessed via two methods. The first is during system boot time when entering the controller's BIOS by depressing the <Ctrl>+<G> keys when prompted. The utility accesses 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 user interface between both OS and pre-OS environments.

### 3.1.1.2 Storage Console+

Storage Console+ (Storage Console Plus) 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 IIR Management Information Base (MIB). It interacts with the host resident MIB as well as the IIR HBA (via the IIR 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 IIR initialization from host system memory during POST.

### 3.1.3.2 IIR API

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

### 3.1.3.4 RAID Configuration Service

The **IIR Service / IIRD** allows remote access to the IIR 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 IIR 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 IIR products both internal and external allowing for a common disk data format across the entire IIR product line.

The I/O device pictured next to the firmware ([Figure 3-1](#)) also communicates over an abstraction layer allowing for elegant integration of multiple I/O device types from different vendors.



# RAID Functionality and Features

---

# 4

## 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 Magneto Optical 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 hierarchical level drives that are detectable by the host operating system. The three lowest level hierarchical 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 destroy.

Within the Advanced Setup of 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 Statuses

Table 4-1 lists the available states of RAID Host and Array Drives

**Table 4-1. Array Drive Statuses**

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; its normal state.
Rebuild	RW	RAID 1 / 4 / 5 / 10	No	Array drive assumes this status after the automatic activation of a Hot Fix or after manual replacement (Hot Plug).
Expand	RW	RAID 4 / 5	Yes	This status indicates that the RAID level and or capacity are being migrated. The Expand status combines with the Ready and Fail statuses as shown below: <ul style="list-style-type: none"> <li>• Ready/Expand</li> <li>• Fail/Expand</li> </ul>
Fail	RW	RAID 1 / 4 / 5 / 10	No	This status indicates that a <i>logical drive</i> has failed.
Error	RO	RAID 4 / 5 / 10	No	If a component of a disk array is missing during startup and <i>fail mode</i> is not activated, the array is set to read only.
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. The Patch status combines with other statuses as shown below: <ul style="list-style-type: none"> <li>• Ready/Patch</li> <li>• Fail/Patch</li> <li>• Error/Patch</li> <li>• Idle/Patch</li> <li>• Rebuild/Patch</li> </ul>
<p>Notes:  RW = Read Write  RO = Read Only  The drive <i>attribute</i> indicates the access level that the host OS has to the drive.</p>				

## 4.1.2 Logical Drive Statuses

**Table 4-2. Host Drive Statuses**

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 within expected parameters.
Notes: RW = Read Write RO = Read Only The drive attribute indicates the level of host OS access to the drive.		

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

The following are limitations assuming the following:

- Ultra160 or Ultra320 speed, depending on the SCSI controller
- Dual channel SCSI controller
- Cabling that meets Ultra160 or Ultra320 specifications, depending on the SCSI controller

**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 SRCZCR RAID controller is **15**, see [Table 4-3](#). 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**.



RAID firmware limits the controller 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 4-3. RAID Controller Drive Maximum Limitations (SRCZCR)**

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

## 4.2 Utilities and Tools

Table 4-4. Utilities and Tools (Sheet 1 of 2)

Management/ Monitoring	Description	Actions
Storage Console (StorCon)	This is a text-based user interface 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 be ran locally or remotely using TCP/IP or SPX/IPX network protocols.
Storage Console+ (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 be ran locally or remotely using TCP/IP or SPX/IPX network protocols.
Monitoring/ Reporting	Description	Actions (using StorCon text-based)
View Statistics	Allows the viewing of I/O activity of host, logical, and physical drives and cache activity.	From the menu, choose Express or Advanced Setup -> View Statistics. User can adjust synch rate and enable and disable logging.
View Events	Allows the viewing and saving of all events regarding the controller since it's last boot up.	From the menu, choose Express or Advanced Setup -> View Events.
View Hard Disk Information	Lists all hard drives connected to the RAID controller with information about each hard drive.	You 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 controller.
Diagnostics	Description	Actions (using StorCon text-based)
Save Information	Gives the administrator the ability to save the configuration information of the RAID controller in an ASCII file for viewing.	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.	From the menu, choose Advanced Setup -> Configure Controller -> Controller Settings. User selected options are: No test, Standard, Double scan, and Intensive.
Parity Verify	Selecting this feature causes the controller to verify the parity on RAID level 4 and 5 arrays.	From the menu, choose 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 correctness.	From the menu, choose Advanced Setup -> Configure Physical Devices -> Select Physical Drive -> Check Surface. <b>Warning!</b> This action destroys all data on the selected drive.

**Table 4-4. Utilities and Tools (Sheet 2 of 2)**

Firmware	Description	Actions
Firmware Update	Utility that allows the updating of the controller's firmware using: 1. DOS utility for major and minor updates (e.g. major, 2.32. xx to 2.33.xx) or 2. StorCon for minor updates (e.g. minor, 2.32.xy to 2.32.xz).	1. Is accessed via DOS utility. 2. Is accessed via the Storage Console during boot up (<Ctrl>+<G>) or from within 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 that the user set the controller's J1B1 jumper to <i>Reset/Firmware Recovery</i> mode and the use of the DOS firmware recovery utility.

## 4.3 RAID Features

### 4.3.1 RAID Level Support

**Table 4-5. Supported RAID Levels**

RAID Level	Description	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)
<b>Disk Pass-Through to Host</b>		
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 controller - 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 controller. This mode of caching utilizes the memory that is located on the controller (SDRAM that is either embedded or an inserted DIMM module). 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 controller's memory.

**Warning:** In the event of power loss to the computer system, data located in 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.**

### 4.3.2.1 Controller Cache Selections

- Caching on/off selectable (sets all controller caching algorithms including Read Ahead to all on or all off)
- Delayed Write on/off selectable (Write Back)

### 4.3.2.2 Physical Disk Cache Selections

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

**Table 4-6. Supported Cache Settings**

	Cache Setting						Cache Configuration Description		
	Disk Cache			Controller Cache					
1	Write	_	Read	_	Cache	_	Delayed Write	•_	No cache
2	Write	_	Read	•	Cache	_	Delayed Write	•_	Disk Read
3	Write	•	Read	_	Cache	_	Delayed Write	•_	Disk Write
4	Write	•	Read	•	Cache	_	Delayed Write	•_	Disk Read Write
5	Write	_	Read	_	Cache	•	Delayed Write	_	Controller Read
6	Write	_	Read	•	Cache	•	Delayed Write	_	Disk Read, Controller Read
7	Write	•	Read	_	Cache	•	Delayed Write	_	Disk Write, Controller Read
8	Write	•	Read	•	Cache	•	Delayed Write	_	Disk Read Write, Controller Read
9	Write	_	Read	_	Cache	•	Delayed Write	•	Controller Read Write
10	Write	_	Read	•	Cache	•	Delayed Write	•	Disk Read, Controller Read Write
11	Write	•	Read	_	Cache	•	Delayed Write	•	Disk Write, Controller Read Write
12	Write	•	Read	•	Cache	•	Delayed Write	•	Disk Read Write, Controller Read Write



**Table 4-6. Supported Cache Settings**

	Cache Setting		Cache Configuration Description
	Disk Cache	Controller Cache	
Legend	<ul style="list-style-type: none"> <li>_ disabled/off</li> <li>• enabled/on</li> <li>•_ either</li> </ul>		

### 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 been configured with *Pooled Hot Fix Access* enabled.

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 4-7](#), if the pooled hot fix drive for the four RAID arrays is 9 GB, then it would only protect arrays 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 4-6](#)).

**Table 4-7. 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
<p><b>NOTE:</b> 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).</p>				

### 4.3.4 Hot Plug Disk Drive Support

This feature allows the ability to remove and replace SCA (single connect controller) 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 as detailed in the following sections.

### 4.3.5 Auto-detection of Hot Plug Disk Drives in Non SAF-TE Drive Enclosures

This is a special feature that allows the use of non-intelligent disk enclosures (requires truly hot pluggable 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 must control the RAID subsystems of the two computer systems (see list of compatible controllers in this section). 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.

**Warning:** †† The SRCZCR controller, with firmware 2.34.yy-Rzzz, is not compatible with all previous controllers and firmware versions. Do not attempt RAID Array Drive Roaming between RAID controllers that are not compatible with the SRCZCR

**controller. Unpredictable behavior may include, but is not limited to, data loss or corruption.**

**Note:** For a migrated RAID array to be recognized by the new host operating system, the host system may need to be rebooted.

Compatible controllers include:

- SRCU42L – Firmware version 2.34.yy-Rzzz
- SRCU31LA - Firmware version 2.3x.yy-Rzzz
- SRCU31A– Firmware version 2.3x.yy-Rzzz
- SRCU32U – Firmware version 2.3x.yy-Rzzz
- SRCZCR – Firmware version 2.3x.yy-Rzzz
- SRCMRU – Firmware version 2.3x.yy-Rzzz

**Note:** Compatibility refers only to how data and RAID configuration information are written to the hard disks. Compatible controllers have the same data file format.

## 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, then 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 be have a status of *ready*.

### 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:

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

### 4.3.9 Background Initialization and Instant Availability (RAID 4 and 5)

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.

There are two build modes available when creating RAID arrays, *Standard* (or non-destructive) and *Destructive*. *Destructive* mode 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 SRCZCR controller, background initialization does not continue during POST until after the OS driver is loaded.***

### 4.3.10 Configurable Disk Drive 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. Most settings are set automatically and can not be configured manually. To view or configure disk drive SCSI parameters from the Storage Console, choose Advanced Setup → Configure Physical Devices → Select Physical Drive → SCSI Parameter/Initialize.

**Table 4-8. Configurable Disk Drive SCSI Parameters**

Parameter	Setting/Value	Configuration	Description
Synch. Transfer	Enabled / Disabled	Automatic	This setting when enabled allows the controller to operate in synchronous transfer mode
Synch. Transfer Rate	Speed (MB/sec)	Automatic	Allows for the setting of the speed for the SCSI hard disks (up to 320 MB/sec for U320 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	Enabled / Disabled	Automatic	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	On / Off	Automatic	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	Configurable	For performance reasons, the <i>Read Ahead</i> and <i>Write</i> cache of the hard disk drives should always be on. This is the default setting.

### 4.3.11 Configurable I/O Controller SCSI Parameters

When the controller starts up, its SCSI parameters are automatically set to their optimal settings. Manual configuration is not required. However, you can customize I/O SCSI parameters for each channel on the controller by using the Storage Console menu Advanced Setup > Configure Controller > Controller Settings.

**Table 4-9. Configurable I/O Controller SCSI Parameters**

Parameter	Setting/Value	Description
Domain Validation	On / Off	When this is set to <i>On</i> , Domain validation allows for a cyclical check of the correct data transfer at a given rate for the respective channel.  For ultra320, the SRCZCR controller supports Basic, Enhanced, and Margined Domain Validation. For ultra160, the SRCZCR controller supports Basic and Enhanced Domain Validation.

# Certifications and Supported Technologies

## 5

### 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 5-1](#) 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 5-1. OS Certification Requirements**

OS Vendor	Details	Test Suite Version
Microsoft	The product shall be WHQL Certified as a RAID controller for the following: <ul style="list-style-type: none"> <li>Windows 2000 Advanced Server</li> </ul>	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: <ul style="list-style-type: none"> <li>NetWare 5.1</li> </ul>	Test kit ver. 3.1 or latest available
SCO	The product shall be certified to receive "Works with SCO" certification for the following: <ul style="list-style-type: none"> <li>SCO UnixWare 7.1.1</li> </ul>	<ul style="list-style-type: none"> <li>PLT ver. 8.13 and HBACert ver. 7.1 or latest available</li> </ul>
Red Hat Linux	The product shall be certified to receive "Red Hat Ready" certification for the following: <ul style="list-style-type: none"> <li>Red Hat Linux 7.1 (the 2.4 kernel)</li> </ul>	<ul style="list-style-type: none"> <li>Test suite ver. 1.6.9 or latest</li> </ul>

## 5.2 Electronic Regulatory Agencies Certifications (Hardware)

Table 5-2. Electronic Equipment Regulatory Certifications

Country/Region	Applicable Specification	Agency Certifications	Product Labeling	Manual Statements
U.S.A. (EMC)	FCC part 2; FCC part 15, subpart B, Class A	Class A verification testing only. No laboratory accreditation required.	FCC Class A statement on product, product packaging, or prominently located in the users manual, as follows:  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.	FCC Class A statements prominently located in users manual, as follows:  Note: 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 designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate  <b>radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.</b> Plus, the users manual shall caution the user that changes or modifications not expressly approved by Intel Corp. could void the user's authority to operate the equipment.
U.S.A. and Canada (safety)	UL Accessory Listing to UL60950, 3 <sup>rd</sup> ed. and CAN/CSA C22.2 No 60950-00	UL and UL-Canada safety certification	UL/UL-Canada Accessory Listing mark, includes Intel file no. E139761, and I.T.E. product category; UL pwb fabricator mark; and V-0 or V-1 flame rating mark.	No known UL required safety statements.
Canada (EMC)	ICES-003 Digital Apparatus, Class A. Note, FCC or CISPR 22:1993 limits accepted in lieu of ICES-003	Class A verification testing only. No laboratory accreditation required.	Industry Canada Class A statement on product, product packaging, or prominently located in the user's manual, as follows: <b>This Class A digital apparatus complies with Canadian ICES-003.</b> French version is optional, as follows: <b>Cet appareil numérique de la classe (*) est conforme à la norme NMB-003 du Canada.</b>	No additional statements.

European Union (EMC)	EN55022:1998 Class A (emissions)	Manufacturer's Declaration of Conformity (DoC) to the EMC directive. No laboratory accreditation required.	CE mark on board or alternately, CE mark on product packaging. Minimum 5mm in height.	Declaration of Conformity statement, plus CE mark, plus the following (including translations), <b>This product follows the provisions of the European Directive 89/336/EEC (EMC) and Directive 73/23/EEC (Safety/LVD).</b>
European Union (EMC)	EN55024:1998 (immunity), consists of IEC 61000-4-2 (ESD), IEC 61000-4-3 (radiated immunity), IEC 61000-4-4 (EFT), IEC 61000-4-5 (surge), IEC 61000-4-6 (conducted immunity), and IEC 61000-4-11 (voltage dips and interrupts)	Same as EN55022 (EMI) requirements above	Same as EN55022 (EMI) requirements above.	Same as EN55022 (EMI) requirements above.
European Union (safety)	EN 60950, 3 <sup>rd</sup> ed.	Manufacturer's Declaration of Conformity (DoC) to the Low Voltage directive. No laboratory accreditation required.	Same as EN55022 (EMI) requirements above.	Same as EN55022 (EMI) requirements above.
Australia, New Zealand (EMC)	AS/NZS 3458:1995 (CISPR 22:1993), Class A	Manufacturer's Declaration of Conformity (DoC)	C-Tick mark with Intel supplier code no. (N232) on board or product packaging.	None.
Korea (EMC)	Nat'l standards w/ emissions based on CISPR 22:1997 and immunity based on CISPR 24 with EN 55024:1998 type modifications	In-country testing required at an RRL (Radio Research Laboratory) accredited lab.	Korean MIC (Ministry of Information and Communication) logo mark on board or product packaging.	Korean MIC (Ministry of Information and Communication) logo mark plus Korean certification text
Taiwan (EMC)	CNS 13438 (CISPR22) Information Technology Equipment	Test at a BSMI accredited lab or an A2LA or NIST NVLAP accredited lab that is notified to BSMI by NIST	Taiwan registration DoC logo mark on board or product packaging.	Traditional Chinese BSMI Class A text in users manual.
Australia, New Zealand, Taiwan, Korea, other APAC (safety)	IEC 60950, 3 <sup>rd</sup> ed. (aka CB report w/ CB certificate)	Notified body test house required, e.g., UL, Nemko, etc.	None.	None.



## 5.3 Supported Specifications and Standards

**Table 5-3. Supported Standards and Specifications**

Specification/Standard	Details
SAF-TE Dpecification 1.0	The product supports SAF-TE Backplane processors, including: processing/reporting/logging of messages and issuing of commands per SAF-TE specification. <a href="http://www.nstor.com/support/whitepapers/safte_spec.doc">http://www.nstor.com/support/whitepapers/safte_spec.doc</a>
SES	The product supports SES Backplanes to include: processing/reporting/logging of messages and issuing of commands per SES specification. <a href="http://www.nowhere.net/~raster/SCSI3/ses-r08a.pdf">http://www.nowhere.net/~raster/SCSI3/ses-r08a.pdf</a>
ACPI Power Management (Windows 2000)	<ul style="list-style-type: none"> <li>• The product supports power states D0, and D3 as defined by Storage Device Class Power Management Reference Specification, Microsoft, Version 1.0A, February 1997.</li> <li>• The product supports power state D3 as follows:               <ul style="list-style-type: none"> <li>• Drive controller (i.e., interface and control electronics): not functional; context lost</li> <li>• Interface mode (i.e., communications timings): not preserved</li> <li>• Drive motor (i.e., spindle): stopped</li> </ul> </li> </ul>
Bootable CD-ROM	<ul style="list-style-type: none"> <li>• The product supports booting from a CD-ROM based on the "El Torito" Bootable CD-ROM Format Specification, version 1.0, <a href="http://www.phoenix.com/PlatSS/PDFs/specs-cdrom.pdf">http://www.phoenix.com/PlatSS/PDFs/specs-cdrom.pdf</a></li> <li>• The product supports a bootable CD-ROM from a PCI 2.2 Plug and Play compliant BIOS.</li> <li>• The product supports No Emulation mode with PnP BIOS's.</li> </ul>



# Appendices

# 7

## 7.1 Referenced Documentation

Table 7-1. Referenced Documentation

Document Title	Order #
SRCZCR Hardware Installation and User's Guide	273561-001
Software Installation and User's Guide	273479-001
RAIDIOS Design Guide v. 1.3	273724-004

## 7.2 BIOS Boot Messages

When booting the IIR controller, the boot message in Figure 7-1 appears. In this example, an IIR Controller has been detected in PCI slot 9, and it has 32MB of RAM installed on the controller. A detailed explanation of the boot message follows.

Figure 7-1. BIOS Boot Message

```

PCI Storage RAID Controller BIOS Version 7.00C
Copyright (C) 1991-2002 by Intel Corporation Apr 24 2002
All rights reserved!
BIOS at 0x000CA000-0x000C27FF
1 Controller(s) found, Selftests OK, scanning SCSI Bus ...
[PCI 1/9] DPMEM (B) at 0xFA000000-0xFA003FFF, INT A = IRQ 11.
[PCI 1/9] SRCMRU - HWL0 - 32 MB SDRAM/ECC - 2048KB Flash-RAM
[PCI 1/9] SN 00FE2211 - Firmware-FW-Version 2.34.00-R01C - Apr 24 2002..
[PCI 1/9] ECC support enabled.
[PCI 1/9] SCSI-A ID:1 LUN:0 -- RBD HyperFLASH KIISCA 320MB/s.
[PCI 1/9] SCSI-A ID:2 LUN:0 -- RBD HyperFLASH KIISCA 320MB/s
[PCI 1/9] SCSI-A ID:3 LUN:0 -- RBD HyperFLASH KIISCA 320MB/s
[PCI 1/9] SCSI-A ID:4 LUN:0 -- RBD HyperFLASH KIISCA async
[PCI 1/9] SCSI-A ID:5 LUN:0 -- RBD HyperFLASH KIISCA async..
[PCI 1/9] Private RAID-5 Host Drive 0 installed (ready)
>>> Press <Ctrl><G> to enter Intel(R) Storage Console <<<

```

**Note:** Because each SCSI channel has its own scan process running in parallel, the POST scan of SCSI devices does not display in sequential order.

BIOS at 0x000CA000 - 0x000C27FF

Unlike ISA or EISA computers where the BIOS address of a peripheral expansion card is set manually (ISA, jumpers) or with the help of a configuration file (EISA, cfg file) and the address space is determined by the user, the PCI system-BIOS automatically maps the BIOS of a PCI compatible peripheral expansion card to a memory address. At each cold or



warm boot, it determines which address space to assign to the BIOS of an expansion card. The message shown in [Figure 7-1](#) reports the physical address occupied by the IIR firmware BIOS.

```
[PCI 1/9]
```

PCI device, bus system 1, slot 9. The PCI 2.x specification allows several PCI bus systems to be present in one PCI computer. All IIR Controllers have been designed to support multiple PCI bus system computers. The slot number indicated in the message above does not refer to the 9th PCI slot, but indicates that the IIR Controller is plugged into a slot which is the 9th one the PCI chipset that the PCI computer can access. To determine which physical PCI slot this corresponds to, consult the system manual of your PCI computer.

```
DPMEM (B) at 0xFA000000 - 0xFA003FFF INT A = IRQ 11
```

DPMEM stands for dual ported memory. The IIR Controller needs this address space of the PCI Computer for the command communication. As with the IIR Controller BIOS, this mapping is also automatically carried out by the PCI system-BIOS. This information is essential when installing Expanded Memory Managers under DOS and Windows. The IIR DPMEM address space has to be excluded from the control of such a manager. Furthermore, this message tells us that the PCI INT A of the IIR Controller has been assigned to the system IRQ 11. This assignment is also carried out automatically if the PCI system-BIOS is 100% PCI 2.X compatible.

```
SRCZCR -- HWL0 -- 32 MB SDRAM / ECC - 2048kB Flash-RAM
```

SRCU42L -- HWL0 stands for the type of IIR Controller found by the IIR firmware BIOS. (HWL means hardware level.) 64 MB SDRAM /ECC indicates that the installed DIMM is a 64MB ECC SDRAM module. 4096kB indicates the size of the installed Flash-RAM. Depending on the size of the installed RAM the following messages are possible (xx = 64, 128, 256):

```
xx MB SDRAM/ECC
```

```
xx MB ECC-SDRAM-Module
```

```
async
```

SCSI devices are indicated as 'async' until they have been initialized by the IIR controller.

```
SCSI-A
```

Indicates SCSI devices which have been detected on the controller's SCSI channels.

```
Press <CTRL><G> to enter INTEL(R) Storage Console
```

After pressing <Ctrl>+<g>, the message *Intel® Storage Console to start after POST* appears. After POST is complete, another message says *Please wait to start Intel® Storage Console....* Soon afterwards, the IO bus scan is completed and the built-in StorCon configuration program is loaded. You can use this program to manage the RAID controller and RAID subsystem.

1	Introduction.....	5
1.1	Acronyms.....	5
1.2	Documentation Conventions.....	6
1.3	Product Overview .....	6
1.4	Operating System Support .....	6
1.5	List of Features.....	7
2	Hardware.....	8
2.1	Physical Layout.....	8
2.2	Major Components .....	8
2.2.1	Intel® 80303 I/O Microprocessor .....	8
2.2.2	Intel Smart 3 FlashFile™ Flash Memory .....	9
2.2.3	SDRAM (Cache).....	9
2.2.4	PCI Interface.....	10
2.2.5	Motherboard Requirements.....	10
2.2.6	Jumpers and Switches.....	10
2.2.7	Diagnostic Features.....	10
2.3	Architecture Features (HW).....	11
2.4	Electrical Characteristics .....	11
2.5	Environmental Specifications.....	11
2.6	Supported Hard Drive Technology .....	12
2.7	Support for Non-Hard-Disk-Drive SCSI Devices (Non-Direct-Access Devices) .	13
3	Software .....	14
3.1	Software Architecture Overview .....	14
3.1.1	User Interface .....	15
3.1.2	System Management.....	15
3.1.3	Common Layers .....	15
3.1.4	RAID Firmware .....	16
4	RAID Functionality and Features .....	17
4.1	Hierarchy .....	17
4.1.1	RAID Host/Array Drive Statuses.....	19
4.1.2	Logical Drive Statuses .....	20
4.1.3	RAID Controller Drive Limitations (Host, Array, Logical, and Physical) .	20
4.2	Utilities and Tools .....	22
4.3	RAID Features.....	23
4.3.1	RAID Level Support.....	23
4.3.2	Caching.....	23
4.3.3	Hot Fix (Spare) Disk Drives .....	25
4.3.4	Hot Plug Disk Drive Support.....	26
4.3.5	Auto-detection of Hot Plug Disk Drives in Non SAF-TE Drive Enclosures	26
4.3.6	Auto-declare Hot Fix (Spare) Drive.....	26
4.3.7	RAID Array Drive Roaming.....	26
4.3.8	Online RAID Array Configurations .....	27
4.3.9	Background Initialization and Instant Availability (RAID 4 and 5) .....	28
4.3.10	Configurable Disk Drive SCSI Parameters .....	29
4.3.11	Configurable I/O Controller SCSI Parameters .....	29



5	Certifications and Supported Technologies .....	30
5.1	OS Certifications.....	30
5.2	Electronic Regulatory Agencies Certifications (Hardware) .....	31
5.3	Supported Specifications and Standards .....	33
6	Technical Drawings and Diagrams .....	34
7	Appendices .....	35
7.1	Referenced Documentation.....	35
7.2	BIOS Boot Messages .....	35