

6Gb/s SAS RAID Cards

ARC-12x4 Series - 4/8 ports

(PCIe 2.0 to 6Gb/s SAS RAID Controllers)

Quick Start Guide

Version: 2.1

Issue Date: May, 2016

Copyright and Trademarks

The information regarding products in this manual is subject to change without prior notice and does not represent a commitment on the part of the vendor, who assumes no liability or responsibility for any errors that may appear in this manual. All brands and trademarks are the properties of their respective owners. This manual contains materials protected under International Copyright Conventions. All rights reserved. No part of this manual may be reproduced in any form or by any means, electronic or mechanical, including photocopying, without the written permission of the manufacturer and the author.

FCC Statement

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against 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.

Manufacturer's Declaration for CE Certification

We confirm ARC-12x4 has been tested and found compliant with the requirements in the council directive relating to the EMC Directive 2004/108/EC. Regarding to the electromagnetic compatibility, the following standards were applied:

EN 55022: 2006, Class B
EN 61000-3-2: 2006
EN 61000-3-3: 1995+A1: 2001+A2: 2005

EN 55024:1998+A1:2001=A2:2003
IEC61000-4-2: 2001
IEC61000-4-3: 2006
IEC61000-4-4: 2004
IEC61000-4-5: 2005
IEC61000-4-6: 2006
IEC61000-4-8: 2001
IEC61000-4-11: 2004

Contents

1. Introduction	6
1.1 Overview	6
1.2 Features	8
2. Hardware Installation	11
2.1 Before You First Installing.....	11
2.2 Board Layout.....	11
2.3 Installation.....	14
2.4 Cables.....	21
2.4.1 Internal Min SAS 4i to SATA Cable	21
2.4.2 Internal Mini SAS 4i (SFF-8087) to 4xSFF-8482 Cable	22
2.4.3 Internal Min SAS 4i (SFF-8087) to Internal Min SAS 4i (SFF-8087) cable	23
2.5 LED Cables.....	23
2.6 Hot-plug Drive Replacement	27
2.6.1 Recognizing a Drive Failure	27
2.6.2 Replacing a Failed Drive	27
2.7 Summary of the installation.....	27
Appendix A.....	31
RAID Concept	31
RAID Set.....	31
Volume Set.....	31
Ease of Use Features.....	32
• Foreground Availability/Background Initialization	32
• Online Array Roaming	32
• Online Capacity Expansion.....	32
• Online RAID Level and Stripe Size Migration	34
• Online Volume Expansion	35
High Availability	35
• Global/Local Hot Spares	35
• Hot-Swap Disk Drive Support.....	36
• Auto Declare Hot-Spare	36
• Auto Rebuilding	37
• Adjustable Rebuild Priority.....	37
High Reliability	38
• Hard Drive Failure Prediction.....	38
• Auto Reassign Sector.....	38
• Consistency Check	39

Data Protection	39
• Battery Backup	39
• Recovery ROM	40
Appendix B	41
Understanding RAID	41
RAID 0	41
RAID 1	42
RAID 10 (1E)	43
RAID 3	43
RAID 5	44
RAID 6	45
RAID x0	45
JBOD	46
Single Disk (Pass-Through Disk)	46
Summary of RAID Levels	47

INTRODUCTION

1. Introduction

This section presents a brief overview of the 6Gb/s SAS RAID controller, ARC-1214-4I/1224-8I. (PCIe 2.0 to 6Gb/s SAS RAID controllers)

1.1 Overview

ARC-1214-4I/1224-8I internal PCIe 2.0 host RAID controllers are a cost-effective solutions for connecting up to 4/8 6Gb/s SAS peripheral devices. The RAID controllers are based on the same RAID kernel of field-proven internal/external RAID controller and same device driver architecture with widely used 3Gb/s and 6Gb/s SAS/SATA RAID controller. Applications that benefit most features from these controllers include NAS, server RAID solutions, web servers, near-line backup, security systems and streaming applications. ARC-1214-4I/1224-8I support directly attached 4/8 internal 6Gb/s SAS ports via 1/2 SFF-8087 connector.

ARC-1214-4I/1224-8I 6Gb/s RAID controllers are low-profile PCI cards, ideal for 1U and 2U rack-mount systems. These controllers utilize the same RAID kernel that has been field-proven in existing external RAID controller products, allowing Areca to quickly bring stable and reliable PCIe 2.0 6Gb/s SAS RAID controllers to the market.

Unparalleled Performance

Embedded with ARM-based storage I/O processor makes those products a pure hardware RAID controller and raise the standard to higher performance levels with several enhancements including 6Gb/s SAS ports, on-board 1GB SDRAM memory and high performance PCIe 2.0 x8 lane host interface bus interconnection. The optional battery backup module provides power to the cache if it contains data not yet written to the drives when power is lost. With several port configuration options including 8 internal and 4 internal, Areca 6Gb/s SAS RAID adapters deliver the ideal price/performance and connectivity solution for entry-level server platforms, workstations and desktop systems. The test result is against overall performance compared to other 6Gb/s SAS RAID controllers.

INTRODUCTION

Unsurpassed Data Availability

As storage capacities continue to rapidly increase, users need greater level of disk drive fault tolerance, which can be implemented without doubling the investment in disk drives. The RAID 6 can offer fault tolerance greater than RAID 1 or RAID 5 but only consumes the capacity of 2 disk drives for distributed parity data. With hardware RAID 6 engine, the 6Gb/s SAS RAID controllers provide the highest RAID 6 feature to meet this requirement. The controller can concurrently compute two parity blocks and get very similar RAID 5 operation performance.

Configuring two or more hard drives in a RAID setup can speed up hard-drive performance and provide automatic protection against data loss. The controllers can provide RAID levels 0, 1, 1E, 3, 5, 6, 10, 30, 50, 60, Single Disk or JBOD for maximum configuration flexibility. Its high data availability and protection derives from the following capabilities: Online RAID Capacity Expansion, Array Roaming, Online RAID Level / Stripe Size Migration, Global Online Spare, Automatic Drive Failure Detection, Automatic Failed Drive Rebuilding, Disk Hot-Swap, Online Background Rebuilding, Instant Availability/Background Initialization, Auto Reassign Sector, Redundant Flash Image and Battery Backup Module. Greater than Two TB Support allows for very large volume set application in 64-bit environment such as data-mining and managing large databases.

Maximum Interoperability

The SAS RAID adapter supports a complete suite of broad operating system including Linux (Open Source), FreeBSD (Open Source), Windows 10/8/2012/7/2008/Vista/2003/XP, VMware, Solaris (Open Source), Mac OS X and more, along with key system monitoring features such as enclosure management (Serial bus & SGPIO) and SNMP function. Our products and technology are based on extensive testing and validation process; same as Areca SAS RAID controllers field-proven compatibility with operating systems, motherboards, applications and device drives.

INTRODUCTION

Easy RAID Management

The controllers contain an embedded McBIOS RAID manager that can access via hot key at M/B BIOS boot-up screen. This pre-boot McBIOS RAID manager can use to simplify the setup and management of RAID controller. The controller firmware also contains a browser-based McRAID storage manager which can be accessed through the Ethernet port or ArchHTTP proxy server in Windows, Linux, FreeBSD, Mac and more environments. The McRAID storage manager allows local and remote to create and modify RAID set, volume set, and monitor RAID status from standard web browser. The Single Admin Portal (SAP) monitor utility can support one application to scan multiple RAID units in the network.

1.2 Features

Controller Architecture

- ARM 900MHz storage I/O processor
- PCIe 2.0 x8 lanes host interface
- 1GB on-board SDRAM with ECC
- Support up to 4/8 internal 6Gb/s SAS physical links
- ARC-1214-4I/1224-8I supports up to 4/8 x 6Gb/s SAS/SATA HDDs/SSD
- Multi-adapter support for large storage requirements
- BIOS boot support for greater fault tolerance
- BIOS PnP (plug and play) and BBS (BIOS boot specification) support
- Support EFI BIOS for Mac Pro and UEFI for PC
- NVRAM for RAID event & transaction log
- Redundant flash image for controller availability
- Battery backup module (BBM) ready (optional)
- RoHS compliant

RAID Features

- RAID level 0, 1, 10(1E), 3, 5, 6, 30, 50, 60, Single Disk or JBOD
- Multi-level RAID 0 and RAID 10 (R00 and R100)
- Support up to 1MB stripe size
- Multiple RAID selection
- Online array roaming
- Online RAID level/stripe size migration

INTRODUCTION

- Online capacity expansion and RAID level migration simultaneously
- Online volume set growth
- SED (self-encrypting drives) function support
- Instant availability and background initialization
- Support global and dedicated hot spare
- Automatic drive insertion/removal detection and rebuilding
- Greater than 2TB capacity per disk drive support
- Greater than 2TB per volume set (64-bit LBA support)
- Support intelligent power management to save energy and extend service life
- Multiple pairs SSD/HDD disk clone function
- SSD automatic monitor clone (AMC) support

Monitors/Notification

- System status indication through global HDD activity/fault connector, individual fault connector, LCD/serial bus connector and alarm buzzer
- SMTP support for email notification
- SNMP support for remote manager
- Enclosure management (Serial bus and SGPIO) ready

RAID Management

- Field-upgradeable firmware in flash ROM

In-Band Manager

- Hot key "boot-up" McBIOS RAID manager via M/B BIOS
- Web browser-based McRAID storage manager via ArchHTTP proxy server for all operating systems
- Support Command Line Interface (CLI)
- API library for customer to write monitor utility
- Single Admin Portal (SAP) monitor utility

Out-of-Band Manager

- Firmware-embedded web browser-based McRAID storage manager, SMTP manager, SNMP agent and Telnet function via Ethernet port
- API library for customer to write monitor utility
- Support push button and LCD display panel (optional)

INTRODUCTION

Operating System

- Windows 10/8/2012/7/2008/Vista/XP/2003
- Linux
- FreeBSD
- VMware (Driver 6.x support CLI in-band management utility)
- Solaris 10/11 x86/x86_64
- Mac OS 10.4.x/10.5.x/10.6.x/10.7.x/10.8.x/10.9.x

6Gb/s SAS RAID controllers		
Model name	ARC-1214-4I	ARC-1224-8I
I/O Processor	ARM 900MHz Storage I/O Processor	
Form Factor (H x L)	Low Profile: 64.4 x 169.5 mm	
Host Bus Type	PCIe 2.0 x8 Lanes	
Driver Connector	1xSFF-8087	2xSFF-8087
Drive Support	Up to 4x6Gb/s SAS/SATA HDDs/SSD	Up to 8x6Gb/s SAS/SATA HDDs/SSD
RAID Level	0, 1, 1E, 3, 5, 6, 10, Single Disk, and JBOD	0, 1, 1E, 3, 5, 6, 10, 30, 50, 60, Single Disk, and JBOD
On-Board Cache	1GB on-board SDRAM with ECC	
Management Port	In-Band: PCIe Out-of-Band: LCD, and LAN Port	
Enclosure Ready	Individual Faulty Header, SGPIO, and Serial bus	

Note:

Low-profile bracket has included on the low profile board shipping package.

2. Hardware Installation

This section describes the procedures for installing the 6Gb/s SAS RAID controllers.

2.1 Before You First Installing

Thanks for purchasing the 6Gb/s SAS RAID controller as your RAID data storage subsystem. This user manual gives simple step-by-step instructions for installing and configuring the 6Gb/s SAS RAID controller. To ensure personal safety and to protect your equipment and data, reading the following information package list carefully before you begin installing.

Package Contents

If your package is missing any of the items listed below, contact your local dealers before you install. **(Disk drives and disk mounting brackets are not included)**

- 1 x 6Gb/s SAS RAID controller in an ESD-protective bag
- 1 x Installation CD – containing driver, relative software, an electronic version of this manual and other related manual
- 1 x Quick start guide
- 1 x Low-profile bracket

2.2 Board Layout

The RAID controllers can support a family SAS interface included 4/8 internal ports with 6Gb/s capability. This section provides the board layout and connector/jumper for the 6Gb/s SAS RAID controller.

HARDWARE INSTALLATION




Figure 2-1, ARC-1214-4I/1224-8I 6Gb/s SAS RAID controller

Connector	Type	Description
1. (J8)	Ethernet Port	RJ45
2. (J7)	Individual Fault LED Header	4-pin header
3. (J6)	Individual Activity (HDD) LED Header	4-pin header
4. (J4)	Global Fault/Activity LED	4-pin header
5. (J2)	Battery Backup Module Connector	12-pin box header
6. (J1)	Manufacture Purpose Port	10-pin header
7. (J3)	I ² C/LCD Connector	8-pin header
8. (SCN1)	SAS 1-4 Ports	SFF-8087
9. (SCN2)	SAS 5-8 Ports	SFF-8087

Table 2-1, ARC-1214-4I/1224-8I connectors

Note:

1. Each SFF-8087 connector uses one port multiplier.
2. Controller's performance is limited by the port multiplier numbers on the controller.
3. Applications will get the best performance when you attach the max number HDDs or rearrange the HDD position to use the max port multiplier on the RAID controller. Each SFF-8087 connector uses one port multiplier chip.

HARDWARE INSTALLATION

Tools Required

An ESD grounding strap or mat is required. Also required are standard hand tools to open your system's case.

System Requirement

The 6Gb/s SAS RAID controller can be installed in an universal PCIe slot and requires a motherboard that:

ARC-1214-4I/1224-8I 6Gb/s SAS RAID controller requires:

- Comply with the PCIe 2.0 x8 lanes
It can work on the PCIe 2.0 x1, x4, x8, and x16 signal with x8 or x16 slot M/B.
- Backward-compatible with PCIe 1.0 x1, x4, x8, and x16 signal with x8 or x16 slot M/B.

Installation Tools

The following items may be needed to assist with installing the 6Gb/s SAS RAID controller into an available PCIe expansion slot.

- Small screwdriver
- Host system hardware manuals and manuals for the disk or enclosure being installed.

Personal Safety Instructions

Use the following safety instructions to help you protect your computer system from potential damage and to ensure your own personal safety.

- Always wear a grounding strap or work on an ESD-protective mat.
- Before opening the system cover, turn off power switches and unplug the power cords. Do not reconnect the power cords until you have replaced the covers.

Electrostatic Discharge

Static electricity can cause serious damage to the electronic components on this 6Gb/s SAS RAID controller. To avoid damage

HARDWARE INSTALLATION

caused by electrostatic discharge, observe the following precautions:

- Do not remove the 6Gb/s SAS RAID controller from its anti-static packaging until you are ready to install it into a computer case.
- Handle the 6Gb/s SAS RAID controller by its edges or by the metal mounting brackets at its each end.
- Before you handle the 6Gb/s SAS RAID controller in any way, touch a grounded, anti-static surface, such as an unpainted portion of the system chassis, for a few seconds to discharge any built-up static electricity.

Warning:

High voltages may be found inside computer equipment. Before installing any of the hardware in this package or removing the protective covers of any computer equipment, turn off power switches and disconnect power cords. Do not reconnect the power cords until you have replaced the covers.

2.3 Installation

Use the following instructions below to install a PCIe 2.0 6Gb/s SAS RAID controller.

Step 1. Unpack

Unpack and remove the PCIe 2.0 6Gb/s SAS RAID controller from the package. Inspect it carefully, if anything is missing or damaged, contact your local dealer.

Step 2. Power PC/Server Off

Turn off computer and remove the AC power cord. Remove the system's cover. For the instructions, please see the computer system documentation.

Step 3. Install the PCIe 6Gb/s SAS RAID Cards

To install the 6Gb/s SAS RAID controller, remove the mounting screw and existing bracket from the rear panel behind the selected

HARDWARE INSTALLATION

PCIe 2.0 slot. Align the gold-fingered edge on the card with the selected PCIe 2.0 slot. Press down gently but firmly to ensure that the card is properly seated in the slot, as shown on Figure 2-2. Then, screw the bracket into the computer chassis. ARC-1214-4I/1224-8I controllers require a PCIe 2.0 x8 slot for better performance.




Figure 2-2, Insert 6Gb/s SAS RAID controller into a PCIe slot

Step 4. Mount the Drives

You can connect the SAS drives to the controller through direct cable and backplane solutions. In the direct connection, SAS/SATA drives are directly connected to 6Gb/s SAS RAID controller port with SAS cables. The 6Gb/s SAS RAID controller can support up to 8 ports. Remove the front bezel from the computer chassis and install the cages or SAS/SATA drives in the computer chassis. Loading drives to the drive tray if cages are installed. Be sure that the power is connected to either the cage backplane or the individual drives.

In the backplane solution, SAS/SATA drives are directly connected to 6Gb/s SAS system backplane. The number of SAS/SATA drives is limited to the number of slots available on the system backplane.

HARDWARE INSTALLATION

Step 5. Install SATA Cable

This section describes SATA cable how to connect on controller.




Figure 2-3, SATA cable connect to Hdds




Figure 2-4, SAS cable connect to backplane

HARDWARE INSTALLATION

Step 6. Install the LED Cable (Optional)

The preferred I/O connector for server backplanes is the internal SFF-8087 connector. This connector has eight signal pins to support four SAS/SATA drives and six pins for the SGPIO (Serial General Purpose Input/Output) side-band signals. The SGPIO bus is used for efficient LED management and for sensing drive Locate status. See SFF 8485 for the specification of the SGPIO bus. For backplane without SGPIO supporting, Please refer to Section 2.6 LED cables for fault/activity LED cable installation.

LED Management: The backplane may contain LEDs to indicate drive status. Light from the LEDs could be transmitted to the outside of the server by using light pipes mounted on the SAS drive tray. A small microcontroller on the backplane, connected via the SGPIO bus to a 6Gb/s SAS RAID controller, could control the LEDs. Activity: blinking 5 times/second and Fault: solid illuminated

Drive Locate Circuitry: The location of a drive may be detected by sensing the voltage level of one of the pre-charge pins before and after a drive is installed.


The following signals define the SGPIO assignments for the Min SAS 4i internal connector (SFF-8087) in the 6Gb/s SAS RAID controller.

PIN	Description	PIN	Description
SideBand0	SClock (Clock signal)	SideBand1	SLoad (Last clock of a bit stream)
SideBand2	Ground	SideBand3	Ground
SideBand4	SDataOut (Serial data output bit stream)	SideBand5	SDataIn (Serial data input bit stream)
SideBand6	Reserved	SideBand7	Reserved

The SFF-8087 to 4 SATA with sideband cable which follows SFF-8448 specification. The SFF-8448 sideband signals cable is reserved for the backplane with header on it. The following signal defines the sideband connector which can work with Areca sideband cable on its SFF-8087 to 4 SATA cable.

The sideband header is located at backplane. For SGPIO to work properly, please connect Areca 8-pin sideband cable to the sideband header as shown below. See the table for pin definitions.

HARDWARE INSTALLATION



Sideband Pin Definitions

Note:

For latest release versions of drivers, please download from <http://www.areca.com.tw/support/main.htm>

Step 7. Adding a Battery Backup Module (Optional)

Please refer to Appendix B Battery Backup Module (ARC6120BA-T121) for installing the BBM in your 6Gb/s SAS RAID controller.

Step 8. Re-check Fault LED Cable Connections (Optional)

Be sure that the proper failed drive channel information is displayed by the fault LEDs. An improper connection will tell the user to "Hot Swap" the wrong drive. This can result in removing the wrong disk (one that is functioning properly) from the controller. This can result in failure and loss of system data.

Step 9. Power up the System

Thoroughly check the installation, reinstall the computer cover, and reconnect the power cord cables. Turn on the power switch at the rear of the computer (if equipped) and then press the power button at the front of the host computer.

Step 10. Install the Controller Driver

For a new system:

- Driver installation usually takes place as part of operating system installation. Please refer to Chapter 4 "Diver Installation" for the detailed installation procedure.

HARDWARE INSTALLATION

For an existing system:

- To install the controller driver into the existing operating system. For the detailed installation procedure, please refer to the Chapter 4, "Driver Installation".

Step 11. Install ArcHTTP Proxy Server

ARC-1214-4I/1224-8I firmware has embedded the web-browser McRAID storage manager. ArcHTTP proxy server will launch the web-browser McRAID storage manager. It provides all of the creation, management and monitor ARC-1214-4I/1224-8I RAID controller status. Please refer to the Chapter 5 for the detail "ArcHTTP Proxy Server Installation". For SNMP agent function, please see the "SNMP Operation & Installation" section in the Appendix C.

Step 12. Configure Volume Set

With Areca series RAID cards, there are 4 methods to manage your ARC-1214-4I/1224-8I RAID controller. It can be configured by using the LCD with keypad, McBIOS RAID manager (terminal emulation) or McRAID storage manager (via LAN port and ArcHTTP proxy utility).

- **Method 1: Internal PCIe Connection (McBIOS RAID Manager)**
The ARC-1214-4I/1224-8I RAID controller can be configured via a BIOS start up McBIOS manager. The McBIOS RAID manager is firmware-based and is used to configure RAID sets and volume sets. Because the utility resides in the ARC-1214-4I/1224-8I RAID controller firmware, operation is independent of any operating systems on your computer. For additional information on using the BIOS on-screen to configure the RAID subsystem see the Chapter 3 of "BIOS Configuration".
- **Method 2: Internal PCIe Connection (McRAID Storage Manager)**
You're now ready to use the McRAID storage manager to set up RAID volumes. Your ARC-1214-4I/1224-8I RAID controller can be configured by using McRAID storage manager (launched by ArcHTTP proxy server). ARC-1214-4I/1224-8I RAID controller has embedded the TCP/IP & web browser-based RAID manager in the firmware. User can use the standard web browsers to manage the RAID controller using ArcHTTP proxy server in-

HARDWARE INSTALLATION

stalled. For additional information on using the McRAID storage manager to configure the RAID controller see the Chapter 6 of “Web Browser-Based Configuration”.

- **Method 3: LAN Port Connection (McRAID Storage Manager)**
The ARC-1214-4I/1224-8I RAID controller has embedded the TCP/IP & web browser-based RAID manager in the firmware. User can remote manage the RAID controller without adding any user specific software (platform independent) via standard web browsers directly connected to the 10/100Mbps RJ45 LAN port. For additional information on using the LAN port to configure the RAID subsystem see the Chapter 6 of “Web Browser-Based Configuration”.
- **Method 4: Front LCD Panel with Keypad (Optional)**
You can use LCD front panel and keypad function to simply create the RAID volume. The LCD status panel also informs you of the disk array’s current operating status at a glance. For additional information on using the LCD to configure the RAID controller see the ARC1000_LCD manual on the shipping CD. The LCD provides a system of screens with areas for information, status indication, or menus. The LCD screen displays up to two lines at a time of menu items or other information.

Step 13. Format, Partition and Mount the ARC-1214-4I/1224-8I RAID Controller Volumes

After the volume set is ready for system accesses, it needs to be partitioned, formatted, and mounted by the operating system. There are various steps, depending on what operating system you are using (Windows, Linux, FreeBSD or Mac, etc.). Detailed steps for each operating system are provided on their disk utility. After that, the ARC-1214-4I/1224-8I RAID controller can be fully used.

Step 14. Determining the Boot Sequences

For PC system:

6Gb/s SAS RAID controller is a bootable controller. If your system already contains a bootable device with an installed operating system, you can set up your system to boot a second operating system from the new controller. To add a second bootable controller, you may need to enter setup of motherboard BIOS and change the

HARDWARE INSTALLATION

device boot sequence so that the 6Gb/s SAS RAID controller heads the list. If the system BIOS setup does not allow this change, your system may be not configurable to allow the 6Gb/s SAS RAID controller to act as a second boot device.

For Apple Mac Pro system:

Areca controller has supported the EFI BIOS on the PCIe 2.0 6Gb/s SAS RAID controller. You have other alternatively to add volumes on the Intel-based Mac bootable device listing. You can follow the following procedures to add 6Gb/s SAS RAID controller on the Mac bootable device listing.

1. Set the BIOS selection in System Controls: Advance Configuration to "EFI" option for Intel_based MacPro boot.
2. Ghost (such as Carbon Copy Cloner ghost utility) the Mac OS X system disk on the Intel-based Mac to the external 6Gb/s SAS RAID controller volume set. Carbon Copy Cloner is an archival type of back up software. You can take your whole Mac OS X system and make a carbon copy or clone to Areca volume similar as an other hard drive.
3. Power up Intel-based Mac, it will take about 30 seconds for controller firmware ready. Areca volume will be added in the bootable device automatically.

2.4 Cables

You can connect the end devices to each other through direct cables or through the SATA backplane connections. The following is an example of some internal SATA cables.

2.4.1 Internal Min SAS 4i to SATA Cable

The Min SAS 4i to SATA cables are used for connection between the 6Gb/s SAS RAID controller internal connectors and connectors on the SAS/SATA connector backplane. The 6Gb/s SAS controllers have 1-6 Min SAS 4i (SFF-8087) internal connectors, each of them can support up to four SAS/SATA drives.

HARDWARE INSTALLATION

These controllers can be installed in a server RAID enclosure with standard SATA connectors backplane. The following diagram shows the picture of Min SAS 4i to 4*SATA cables. Backplane supports SGPIO header can leverage the SGPIO function on the 6Gb/s SAS/SATA RAID controller through the sideband cable.




Figure 2-5, Internal Min SAS 4i to 4x SATA cable

The SFF-8448 sideband signals cable is reserved for the backplane with header on it.

2.4.2 Internal Mini SAS 4i (SFF-8087) to 4xSFF-8482 Cable

These controllers can be installed in a server RAID enclosure with out a backplane. The kind of cable will attach directly to the SAS disk drives. The following diagram shows the picture of Mini SAS 4i (SFF-8087) to 4xSFF-8482 cables.




Figure 2-6, Min SAS 4i to 4xSFF-8482 cable

HARDWARE INSTALLATION

2.4.3 Internal Min SAS 4i (SFF-8087) to Internal Min SAS 4i (SFF-8087) cable

The 6Gb/s SATA RAID controllers have 1-6 Min SAS 4i internal SFF-8087 connectors, each of them can support up to four SATA signals. These controllers can be installed in a server RAID enclosure with Min SAS 4i internal connectors backplane. This Min SAS 4i cable has eight signal pins to support four SATA drives and six pins for the SGPIO (Serial General Purpose Input/Output) side-band signals. The SGPIO bus is used for efficient LED management and for sensing drive Locate status.



Figure 2-7, Internal Min SAS 4i to Min SAS 4i cable


2.5 LED Cables

There is no SGPIO supported in the most of old version SATA backplane. The 6Gb/s SAS RAID controller also provides two kinds of alternative LED cable header to support the fault/activity status for those backplanes. The global indicator connector is used by the server/desktop system global indicator LED.

The following electronics schematic is the 6Gb/s SAS RAID controller logical of fault/activity header. The signal from EPLD CTL output pin is cathode (-) side.

The following diagrams and descriptions describe each type of connector.

HARDWARE INSTALLATION



Note:

A cable for the global indicator comes with your computer system. Cables for the individual drive LEDs may come with a drive cage, or you may need to purchase them.

A: Individual Fault LED and Global Activity/Fault Indicator Connector

Most of the backplanes have supported the HDD activity from the HDD. The 6Gb/s SAS RAID controller also provides the fault activity for fault LED. Connect the cables for the drive fault LEDs between the backplane of the cage and the respective connector on the 6Gb/s SAS RAID controller.


The following table is the fault LED signal behavior.

LED	Normal Status	Problem Indication
Fault LED	<p>When the fault LED is solid illuminated, there is no disk present.</p> <p>When the fault LED is off, then disk is present and status is normal.</p>	<p>When the fault LED is slow blinking (2 times/sec), that disk drive has failed and should be hot-swapped immediately.</p> <p>When the activity LED is illuminated and fault LED is fast blinking (10 times/sec) there is rebuilding activity on that disk drive.</p>

If the system will use only a single global indicator, attach the LED to the two pins of the global activity/fault connector. The global fault pin pair connector is the overall fault signal. This signal will light up in any disk drive failure.

HARDWARE INSTALLATION

Figure 2-8, ARC-1214-4I/
1224-8I individual and global
LED indicator



B: Areca Serial Bus Connector

You can also connect the Areca serial bus interface to a proprietary SATA backplane enclosure. This can reduce the number of activity LED and/or fault LED cables. The serial bus (I²C) interface can also cascade to another SATA backplane enclosure for the additional channel status display.





Figure 2-9, Activity/Fault LED serial bus connector connected between ARC-1214-4i/ARC-1224-8i & 4 SATA HDD backplane.

The following picture and table is the serial bus signal name description for the LCD/I²C Module Connector (J3).



HARDWARE INSTALLATION

PIN	Description	PIN	Description
1	Power (+5V)	2	GND
3	LCD Module Interrupt	4	Protect Key
5	LCD Module Serial Data	6	Fault/Activity Clock
7	Fault/Activity Serial Data	8	LCD Module Clock

You can use one optional LCD front panel and keypad function to simply create the RAID volume. The LCD status panel also informs you the disk array's current operating status at a glance. The LCD configuration is described in a separate manual: RAID Card_LCD manual. The LCD housed in a 5¼-inch half-height or 3.5-inch canister.




Figure 2-10, Connect to LCD Status Panel

You can use one optional Card LED indicator connected serial bus interface to display fault/activity status on I/O bracket or 3.5-inch canister. This LED provides indications about the operational state of the HDD on the RAID controller.




Figure 2-11, Connect to LED Indicator

2.6 Hot-plug Drive Replacement

The RAID controller supports the ability of performing a hot-swap drive replacement without powering down the system. A disk can be disconnected, removed, or replaced with a different disk without taking the system off-line. The RAID rebuilding will be processed automatically in the background. When a disk is hot swap, the RAID controller may no longer be fault tolerant. Fault tolerance will be lost until the hot swap drive is subsequently replaced and the rebuild operation is completed.

2.6.1 Recognizing a Drive Failure

A drive failure can be identified in one of the following ways:

1. An error status message lists failed drives in the event log.
2. A fault LED illuminates on the front of RAID subsystem if failed drives are inside.

2.6.2 Replacing a Failed Drive

With RAID subsystem drive tray, you can replace a defective physical drive while your computer is still operating. When a new drive has been installed, data reconstruction will be automatically started to rebuild the contents of the disk drive. The controller always uses the smallest hotspare that "fits". If a hotspare is used and the defective drive is exchanged on-line, the new inserted HDD will automatically assign as a hotspare HDD.


2.7 Summary of the installation

The flow chart below describes the installation procedures for 6Gb/s SAS RAID controllers. These procedures includes hardware installation, the creation and configuration of a RAID volume through the McBIOS/McRAID manager, OS installation and installation of 6Gb/s SAS RAID controller software.

The software components configure and monitor the 6Gb/s SAS RAID controllers as following table.

HARDWARE INSTALLATION

Configuration Utility	Operating System Supported
McBIOS RAID Manager	OS-Independent
McRAID Storage Manager (Via ArchHTTP proxy server)	Windows 10/8/Server 2012/7/2008/ Vista/XP/2003, Linux, FreeBSD, Solaris and Mac
McRAID Storage Manager (Via Ethernet port)	OS-Independent
SAP Monitor (Single Admin Portal to scan for multiple RAID units in the net- work, via ArchHTTP proxy server)	Windows 10/8/Server 2012/7/2008/ Vista/XP/2003
SNMP Manager Console Integration	Windows 10/8/Server 2012/7/2008/ Vista/XP/2003, Linux and FreeBSD



McRAID Storage Manager

Before launching the firmware-embedded web server, McRAID storage manager through the PCIe bus, you need first to install the ArchHTTP proxy server on your server system. If you need additional information about installation and start-up of this function, see the McRAID Storage Manager section in Chapter 6.

ArchHTTP Proxy Server

ArchHTTP has to be installed for GUI RAID console (McRAID storage manager) to run. It is used to launch the web browser McRAID storage manager. It also runs as a service or daemon in the background that allows capturing of events for mail and SNMP traps notification. If you need additional information about installation and start-up of this function, see the ArchHTTP Proxy Server Installation section in Chapter 5.

BIOS CONFIGURATION

CLI Utility

CLI (Command Line Interface) lets you set up and manage RAID controller through a command line interface. CLI performs many tasks at the command line. You can download CLI manual from Areca website or software CD <CDROM>\DOCS directory.

SNMP Manager Console Integration

There are two ways to transport SNMP data on the 6Gb/s SAS RAID controller: in-band PCIe host bus interface or out-of-band built-in LAN interface. Enter the "SNMP Trap IP Address" option on the firmware-embedded SNMP configuration function for user to select the SNMP data agent-side communication from the out-of-band built-in LAN interface. To use in-band PCIe host bus interface, keep blank on the "SNMP Trap IP Address" option.

● **Out of Band-Using LAN Port Interface**

Out-of-band interface refers to transport SNMP data of 6Gb/s SAS controllers from a remote station connected to the controller through a network cable. Before launching the SNMP manager on the clinet, you need firstly to enable the firmware-embedded SNMP agent function and no additional agent software inquired on your server system. If you need additional information about installation and start-up this function, see the section 6.8.4 SNMP Configuration.

● **In-Band-Using PCIe Host Bus Interface**

In-band interface refers to management of the SNMP data of 6Gb/s SAS controllers from a PCIe host bus. In-band interface is simpler than out-of-band interface for it requires less hardware in its configuration. Since the 6Gb/s SAS RAID controller is already installed in the host system, no extra connection is necessary. Just load the necessary in-band Areca SNMP extension agent for the controllers. Before launching the SNMP agent in the sever, you need first to enable the firmware-embedded SNMP community configuration and install Areca SNMP extension agent in your server system. If you need additional information about installation and start-up the function, see the SNMP Operation & Installation section in the Appendix C.

BIOS CONFIGURATION

Single Admin Portal (SAP) Monitor

This utility can scan for multiple RAID units on the network and monitor the controller set status. For additional information, see the utility manual (SAP) in the packaged CD or download it from the web site <http://www.areca.com.tw>.

Appendix A


RAID Concept

RAID Set

A RAID set is a group of disks connected to a RAID controller. A RAID set contains one or more volume sets. The RAID set itself does not define the RAID level (0, 1, 1E, 3, 5, 6, 10, 30, 50, 60, etc); the RAID level is defined within each volume set. Therefore, volume sets are contained within RAID sets and RAID Level is defined within the volume set. If physical disks of different capacities are grouped together in a RAID set, then the capacity of the smallest disk will become the effective capacity of all the disks in the RAID set.

Volume Set

Each volume set is seen by the host system as a single logical device (in other words, a single large virtual hard disk). A volume set will use a specific RAID level, which will require one or more physical disks (depending on the RAID level used). RAID level refers to the level of performance and data protection of a volume set. The capacity of a volume set can consume all or a portion of the available disk capacity in a RAID set. Multiple volume sets can exist in a RAID set. For the RAID controller, a volume set must be created either on an existing RAID set or on a group of available individual disks (disks that are about to become part of a RAID set). If there are pre-existing RAID sets with available capacity and enough disks for the desired RAID level, then the volume set can be created in the existing RAID set of the user's choice.



APPENDIX

In the illustration, volume 1 can be assigned a RAID level 5 of operation while volume 0 might be assigned a RAID level 1E of operation. Alternatively, the free space can be used to create volume 2, which could then be set to use RAID level 5.

Ease of Use Features

- **Foreground Availability/Background Initialization**

RAID 0 and RAID 1 volume sets can be used immediately after creation because they do not create parity data. However, RAID 3, 5, 6, 30, 50 or 60 volume sets must be initialized to generate parity information. In Background Initialization, the initialization proceeds as a background task, and the volume set is fully accessible for system reads and writes. The operating system can instantly access the newly created arrays without requiring a reboot and without waiting for initialization to complete. Furthermore, the volume set is protected against disk failures while initialing. If using Foreground Initialization, the initialization process must be completed before the volume set is ready for system accesses.

- **Online Array Roaming**

RAID controllers store RAID configuration information on the disk drives. The controller therefore protects the configuration settings in the event of controller failure. Online array roaming allows the administrators the ability to move a complete RAID set to another system without losing RAID configuration information or data on that RAID set. Therefore, if a server fails, the RAID set disk drives can be moved to another server with an Areca RAID controllers and the disks can be inserted in any order.

- **Online Capacity Expansion**


Online Capacity Expansion makes it possible to add one or more physical drives to a volume set without interrupting server operation, eliminating the need to backup and restore after reconfiguration of the RAID set. When disks are added to a RAID set, unused capacity is added to the end of the

RAID set. Then, data on the existing volume sets (residing on the newly expanded RAID set) is redistributed evenly across all the disks. A contiguous block of unused capacity is made available on the RAID set. The unused capacity can be used to create additional volume sets.

A disk, to be added to a RAID set, must be in normal mode (not failed), free (not spare, in a RAID set, or passed through to host) and must have at least the same capacity as the smallest disk capacity already in the RAID set.


Capacity expansion is only permitted to proceed if all volumes on the RAID set are in the normal status. During the expansion process, the volume sets being expanded can be accessed by the host system. In addition, the volume sets with RAID level 1, 10, 3, 5 or 6 are protected against data loss in the event of disk failure(s). In the case of disk failure, the volume set changes from "migrating" state to "migrating+degraded" state. When the expansion is completed, the volume set would then transition to "degraded" mode. If a global hot spare is present, then it further changes to the "rebuilding" state.

The expansion process is illustrated as following figure.



RAID controller redistributes the original volume set over the original and newly added disks, using the same fault-tolerance configuration. The unused capacity on the expand RAID set can then be used to create an additional volume set, with a different fault tolerance setting (if required by the user.)

APPENDIX



- **Online RAID Level and Stripe Size Migration**

For those who wish to later upgrade to any RAID capabilities, a system with online RAID level/stripe size migration allows a simplified upgrade to any supported RAID level without having to reinstall the operating system.

The RAID controllers can migrate both the RAID level and stripe size of an existing volume set, while the server is online and the volume set is in use. Online RAID level/stripe size migration can prove helpful during performance tuning activities as well as when additional physical disks are added to the RAID controller. For example, in a system using two drives in RAID level 1, it is possible to add a single drive and add capacity and retain fault tolerance. (Normally, expanding a RAID level 1 array would require the addition of two disks). A third disk can be added to the existing RAID logical drive and the volume set can then be migrated from RAID level 1 to 5. The result would be parity fault tolerance and double the available capacity without taking the system down. A fourth disk could be added to migrate to RAID level 6. It is only possible to migrate to a higher RAID level by adding a disk; disks in an existing array can't be reconfigured for a higher RAID level without adding a disk.

Online migration is only permitted to begin, if all volumes to be migrated are in the normal mode. During the migration process, the volume sets being migrated are accessed by the host system. In addition, the volume sets with RAID level 1, 1E, 10, 3, 5 or 6 are protected against data loss in the event of disk failure(s). In the case of disk failure, the volume set

transitions from migrating state to (migrating+degraded) state. When the migration is completed, the volume set transitions to degraded mode. If a global hot spare is present, then it further transitions to rebuilding state.

● **Online Volume Expansion**

Performing a volume expansion on the controller is the process of growing only the size of the latest volume. A more flexible option is for the array to concatenate an additional drive into the RAID set and then expand the volumes on the fly. This happens transparently while the volumes are online, but, at the end of the process, the operating system will detect free space at after the existing volume.

Windows, NetWare and other advanced operating systems support volume expansion, which enables you to incorporate the additional free space within the volume into the operating system partition. The operating system partition is extended to incorporate the free space so it can be used by the operating system without creating a new operating system partition.

You can use the Diskpart.exe command line utility, included with Windows Server 2003 or the Windows 2000 Resource Kit, to extend an existing partition into free space in the dynamic disk.

Third-party software vendors have created utilities that can be used to repartition disks without data loss. Most of these utilities work offline. Partition Magic is one such utility.

High Availability

● **Global/Local Hot Spares**

A hot spare is an unused online available drive, which is ready for replacing the failure disk. The hot spare is one of the most important features that RAID controllers provide to deliver a high degree of fault-tolerance. A hot spare is a spare physical drive that has been marked as a hot spare and therefore is not a member of any RAID set. If a disk drive used in a

APPENDIX

volume set fails, then the hot spare will automatically take its place and the data previously located on the failed drive is reconstructed on the hot spare.

Dedicated hot spare is assigned to serve one specified RAID set. Global hot spare is assigned to serve all RAID set on the RAID controller. Dedicated hot spare has higher priority than the global hot spare. For this feature to work properly, the hot spare must have at least the same capacity as the drive it replaces. The hot spare function only works with RAID level 1, 1E, 3, 5, 6, 10, 30, 50, or 60 volume set.

The "Create Hot Spare" option gives you the ability to define a global/dedicated hot spare disk drive. To effectively use the hot spare feature, you must always maintain at least one drive that is marked as a global hot spare.

Important:

The hot spare must have at least the same capacity as the drive it replaces.

● **Hot-Swap Disk Drive Support**

The RAID controller chip includes a protection circuit that supports the replacement of SAS/SATA hard disk drives without having to shut down or reboot the system. A removable hard drive tray can deliver "hot swappable" fault-tolerant RAID solutions. This feature provides advanced fault tolerant RAID protection and "online" drive replacement.

● **Auto Declare Hot-Spare**

If a disk drive is brought online into a system operating in degraded mode, the RAID controllers will automatically declare the new disk as a spare and begin rebuilding the degraded volume. The Auto Declare Hot-Spare function requires that the smallest drive contained within the volume set in which the failure occurred.

In the normal status, the newly installed drive will be reconfigured an online free disk. But, the newly-installed drive is

automatically assigned as a hot spare if any hot spare disk was used to rebuild and without new installed drive replaced it. In this condition, the Auto Declare Hot-Spare status will be disappeared if the RAID subsystem has since powered off/on.

The Hot-Swap function can be used to rebuild disk drives in arrays with data redundancy such as RAID level 1, 1E, 3, 5, 6, 10, 30, 50 and 60.

● **Auto Rebuilding**

If a hot spare is available, the rebuild starts automatically when a drive fails. The RAID controllers automatically and transparently rebuild failed drives in the background at user-definable rebuild rates.

If a hot spare is not available, the failed disk drive must be replaced with a new disk drive so that the data on the failed drive can be automatically rebuilt and so that fault tolerance can be maintained.

RAID controllers will automatically restart the system and rebuilding process if the system is shut down or powered off abnormally during a reconstruction procedure condition.

When a disk is hot swapped, although the system is functionally operational, the system may no longer be fault tolerant. Fault tolerance will be lost until the removed drive is replaced and the rebuild operation is completed.

During the automatic rebuild process, system activity will continue as normal, however, the system performance and fault tolerance will be affected.

● **Adjustable Rebuild Priority**

Rebuilding a degraded volume incurs a load on the RAID subsystem. The RAID controllers allow the user to select the rebuild priority to balance volume access and rebuild tasks appropriately. The Background Task Priority is a relative indication of how much time the controller devotes to a back-

APPENDIX

ground operation, such as rebuilding or migrating. RAID controller allows user to choose the task priority (Ultra Low (5%), Low (20%), Medium (50%), High (80%)) to balance volume set access and background tasks appropriately. For high array performance, specify an Ultra Low value. Like volume initialization, after a volume rebuilds, it does not require a system reboot.

High Reliability

● Hard Drive Failure Prediction

In an effort to help users avoid data loss, disk manufacturers are now incorporating logic into their drives that acts as an "early warning system" for pending drive problems. This system is called SMART. The disk integrated controller works with multiple sensors to monitor various aspects of the drive's performance, determines from this information if the drive is behaving normally or not, and makes available status information to 6Gb/s SATA RAID controller firmware that probes the drive and look at it.

The SMART can often predict a problem before failure occurs. The controllers will recognize a SMART error code and notify the administer of an impending hard drive failure.

● Auto Reassign Sector

Under normal operation, even initially defect-free drive media can develop defects. This is a common phenomenon. The bit density and rotational speed of disks is increasing every year, and so are the potential of problems. Usually a drive can internally remap bad sectors without external help using cyclic redundancy check (CRC) checksums stored at the end of each sector.

The RAID controller drives perform automatic defect re-assignment for both read and write errors. Writes are always completed - if a location to be written is found to be defective, the drive will automatically relocate that write command to a new location and map out the defective location. If there is

a recoverable read error, the correct data will be transferred to the host and that location will be tested by the drive to be certain the location is not defective. If it is found to have a defect, data will be automatically relocated, and the defective location is mapped out to prevent future write attempts.

In the event of an unrecoverable read error, the error will be reported to the host and the location will be flagged as being potentially defective. A subsequent write to that location will initiate a sector test and relocation should that location prove to have a defect. Auto Reassign Sector does not affect disk subsystem performance because it runs as a background task. Auto Reassign Sector discontinues when the operating system makes a request.

- **Consistency Check**

A consistency check is a process that verifies the integrity of redundant data. To verify RAID 3, 5, 6, 30, 50 or 60 redundancy, a consistency check reads all associated data blocks, computes parity, reads parity, and verifies that the computed parity matches the read parity.

Consistency checks are very important because they detect and correct parity errors or bad disk blocks in the drive. A consistency check forces every block on a volume to be read, and any bad blocks are marked; those blocks are not used again. This is critical and important because a bad disk block can prevent a disk rebuild from completing. We strongly recommend that you run consistency checks on a regular basis—at least once per week. Note that consistency checks degrade performance, so you should run them when the system load can tolerate it.

Data Protection

- **Battery Backup**

The RAID controllers are armed with a Battery Backup Module (BBM). While a Uninterruptible Power Supply (UPS) protects most servers from power fluctuations or failures, a BBM pro-

APPENDIX

vides an additional level of protection. In the event of a power failure, a BBM supplies power to retain data in the RAID controller's cache, thereby permitting any potentially dirty data in the cache to be flushed out to secondary storage when power is restored.

The batteries in the BBM are recharged continuously through a trickle-charging process whenever the system power is on. The batteries protect data in a failed server for up to three or four days, depending on the size of the memory module. Under normal operating conditions, the batteries last for three years before replacement is necessary.

- **Recovery ROM**

RAID controller firmware is stored on the flash ROM and is executed by the I/O processor. The firmware can also be updated through the RAID controllers PCIe 2.0 bus port or Ethernet port without the need to replace any hardware chips. During the controller firmware upgrade flash process, it is possible for a problem to occur resulting in corruption of the controller firmware. With our Redundant Flash Image feature, the controller will revert back to the last known version of firmware and continue operating. This reduces the risk of system failure due to firmware crash.

Appendix B

Understanding RAID


RAID is an acronym for Redundant Array of Independent Disks. It is an array of multiple independent hard disk drives that provides high performance and fault tolerance. The RAID controller implements several levels of the Berkeley RAID technology. An appropriate RAID level is selected when the volume sets are defined or created. This decision should be based on the desired disk capacity, data availability (fault tolerance or redundancy), and disk performance. The following section discusses the RAID levels supported by the RAID controller.

The RAID controllers makes the RAID implementation and the disks' physical configuration transparent to the host operating system. This means that the host operating system drivers and software utilities are not affected, regardless of the RAID level selected. Correct installation of the disk array and the controller requires a proper understanding of RAID technology and the concepts.

RAID 0


RAID 0, also referred to as striping, writes stripes of data across multiple disk drives instead of just one disk drive. RAID 0 does not provide any data redundancy, but does offer the best High-speed data throughput. RAID 0 breaks up data into smaller blocks and then writes a block to each drive in the array. Disk striping enhances performance because multiple drives are accessed simultaneously; the reliability of RAID level 0 is less because the entire array will fail if any one disk drive fails.

APPENDIX




RAID 1

RAID 1 is also known as “disk mirroring”; data written on one disk drive is simultaneously written to another disk drive. Read performance will be enhanced if the array controller can, in parallel, access both members of a mirrored pair. During writes, there will be a minor performance penalty when compared to writing to a single disk. If one drive fails, all data (and software applications) are preserved on the other drive. RAID 1 offers extremely high data reliability, but at the cost of doubling the required data storage capacity.



RAID 10 (1E)


RAID 10(1E) is a combination of RAID 0 and RAID 1, combining striping with disk mirroring. RAID Level 10 combines the fast performance of Level 0 with the data redundancy of level 1. In this configuration, data is distributed across several disk drives, similar to Level 0, which are then duplicated to another set of drive for data protection. RAID 10 has been traditionally implemented using an even number of disks, some hybrids can use an odd number of disks as well. Illustration is an example of a hybrid RAID 10(1E) array comprised of five disks; A, B, C, D and E. In this configuration, each strip is mirrored on an adjacent disk with wrap-around. Areca RAID 10 offers a little more flexibility in choosing the number of disks that can be used to constitute an array. The number can be even or odd.



RAID 3


RAID 3 provides disk striping and complete data redundancy through a dedicated parity drive. RAID 3 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks, and then writes the blocks to all but one drive in the array. The parity data created during the exclusive-or is then written to the last drive in the array. If a single drive fails, data is still available by computing the exclusive-or of the contents corresponding strips of the surviving member disk. RAID 3 is best for applications that require very fast data-transfer rates or long data blocks.

APPENDIX




RAID 5

RAID 5 is sometimes called striping with parity at byte level. In RAID 5, the parity information is written to all of the drives in the controllers rather than being concentrated on a dedicated parity disk. If one drive in the system fails, the parity information can be used to reconstruct the data from that drive. All drives in the array system can be used for seek operations at the same time, greatly increasing the performance of the RAID system. This relieves the write bottleneck that characterizes RAID 4, and is the primary reason that RAID 5 is more often implemented in RAID arrays.



RAID 6

RAID 6 provides the highest reliability. It is similar to RAID 5, but it performs two different parity computations or the same computation on overlapping subsets of the data. RAID 6 can offer fault tolerance greater than RAID 1 or RAID 5 but only consumes the capacity of 2 disk drives for distributed parity data. RAID 6 is an extension of RAID 5 but uses a second, independent distributed parity scheme. Data is striped on a block level across a set of drives, and then a second set of parity is calculated and written across all of the drives.




RAID x0

RAID level-x0 refers to RAID level 00, 100, 30, 50 and 60. RAID x0 is a combination multiple RAID x volume sets with RAID 0 (striping). Striping helps to increase capacity and performance without adding disks to each RAID x array. The operating system uses the spanned volume in the same way as a regular volume. Up to one drive in each sub-volume (RAID 3 or 5) may fail without loss of data. Up to two drives in each sub-volume (RAID 6) may fail without loss of data. RAID level x0 allows more physical drives in an array. The benefits of doing so are larger volume sets, increased performance, and increased reliability.

The following illustration is an example of a RAID level x0 logical drive.

APPENDIX



Important:

RAID level 00, 100, 30, 50 and 60 can support up to eight RAID set. If volume is RAID level 00, 100, 30, 50, or 60, you can't change the volume to another RAID level. If volume is RAID level 0, 1, 10(1E), 3, 5, or 6, you can't change the volume to RAID level 00, 100, 30, 50, or 60.

JBOD

(Just a Bunch Of Disks) A group of hard disks in a RAID box are not set up as any type of RAID configuration. All drives are available to the operating system as an individual disk. JBOD does not provide data redundancy.

Single Disk (Pass-Through Disk)

Pass through disk refers to a drive that is not controlled by the RAID firmware and thus can not be a part of a RAID volume. The drive is available to the operating system as an individual disk.

Summary of RAID Levels

6Gb/s SATA RAID controller supports RAID Level 0, 1, 10(1E), 3, 5, 6, 30, 50 and 60. The following table provides a summary of RAID levels.

RAID Level Comparison			
RAID Level	Description	Disks Requirement (Minimum)	Data Availability
0	Also known as striping. Data distributed across multiple drives in the array. There is no data protection.	1	No data Protection
1	Also known as mirroring. All data replicated on 2 separated disks. N is almost always 2. Due to this is a 100 % duplication, so is a high costly solution.	2	Up to one disk failure
10(1E)	Also known as mirroring and striping. Data is written to two disks simultaneously, and allows an odd number or disk. Read request can be satisfied by data read from wither one disk or both disks.	3	Up to one disk failure in each sub-volume
3	Also known Bit-Interleaved Parity. Data and parity information is subdivided and distributed across all data disks. Parity information normally stored on a dedicated parity disk.	3	Up to one disk failure
5	Also known Block-Interleaved Distributed Parity. Data and parity information is subdivided and distributed across all disk. Parity information normally is interspersed with user data.	3	Up to one disk failure
6	RAID 6 provides highest reliability, but not widely used. Similar to RAID 5, but does two different parity computations or the same computation on overlapping subsets of the data. The RAID 6 can offer fault tolerance greater that RAID 1 or RAID 5 but only consumes the capacity of 2 disk drives for distributed parity data.	4	Up to two disk failure

APPENDIX

30	RAID 30 is a combination multiple RAID 3 volume sets with RAID 0 (striping)	6	Up to one disk failure in each sub-volume
50	RAID 50 is a combination multiple RAID 5 volume sets with RAID 0 (striping)	6	Up to one disk failure in each sub-volume
60	RAID 60 is a combination multiple RAID 6 volume sets with RAID 0 (striping)	8	Up to two disk failure in each sub-volume