

Intel® RAID Software User's Guide:

- Intel® Embedded Server RAID Technology 2
- Intel® IT/IR RAID
- Intel® Integrated Server RAID
- Intel® RAID Controllers using the Intel® RAID Software Stack 3

Revision 17.0

November, 2011

Intel Order Number: D29305-017

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

Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

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

Copyright © 2011 by Intel Corporation. Portions Copyright 2005-2011 by LSI Logic Corporation. All rights reserved.

Table of Contents

Chapter 1: Overview	1
Supported Hardware	1
Software	3
RAID Terminology	4
Fault Tolerance	4
Enclosure Management	6
Performance	6
Chapter 2: RAID Levels	9
Summary of RAID Levels	9
Selecting a RAID Level	10
RAID 0 - Data Striping	10
RAID 1 - Disk Mirroring/Disk Duplexing	11
RAID 5 - Data Striping with Striped Parity	11
RAID 6 - Distributed Parity and Disk Striping	12
RAID IME	13
RAID 10 - Combination of RAID 1 and RAID 0	14
RAID 50 - Combination of RAID 5 and RAID 0	15
RAID 60 - Combination of RAID 0 and RAID 6	16
RAID Configuration Strategies	19
Maximizing Fault Tolerance	19
Maximizing Performance	20
Maximizing Storage Capacity	21
RAID Availability	23
RAID Availability Concept	23
Spare Drives	23
Rebuilding	23
Copyback	23
Configuration Planning	24
Dimmer Switch Feature	24
Number of Physical Disks	25
MegaRAID Fast Path	25
4K Sector Drive Support	25
Larger than 2TB Drive Support	25
Power Save settings	26
Array Purpose	26
Chapter 3: RAID Utilities	27
Intel® Embedded Server RAID Technology 2 BIOS Configuration Utility	27
LSI MPT* SAS BIOS Configuration Utility	27
Intel® RAID BIOS Console 2 Configuration Utility for Intelligent RAID	28

Intel® RAID Web Console 2 Configuration and Monitoring Utility	29
Drive Hierarchy within the RAID Firmware	30
Intel® Intelligent RAID Controller Features	30
Enterprise Features	30
Fault Tolerant Features	31
Cache Options and Settings	32
Background Tasks	32
Error Handling	33
Audible Alarm	33
Chapter 4: Intel® RAID Drivers	35
RAID Driver Installation for Microsoft Windows*	35
Installation in a New Microsoft Windows* Operating System	35
Installation in an Existing Microsoft Windows* Operating System	36
RAID Driver Installation for Red Hat* Enterprise Linux	36
RAID Driver Installation for SuSE* Linux	37
RAID Driver Installation for Novell NetWare*	37
Installation in a New Novell Netware* System	37
Installation in an Existing Novell Netware* System	38
RAID Driver Installation for Solaris* 10	38
Installation in a New Solaris* System	38
Installation in an Existing Solaris* System	38
Chapter 5: Intel® Embedded Server RAID BIOS Configuration Utility	39
Creating, Adding or Modifying a Virtual Drive Configuration	40
Setting the Write Cache and Read Ahead Policies	42
Working with a Global Hot-spare Drive	42
Adding a Hot-spare Drive	42
Removing a Hot-spare Drive	43
Rebuilding a Drive	43
Auto Rebuild and Auto Resume	43
Checking Data Consistency	44
Viewing and Changing Device Properties	44
Forcing a Drive Online or Offline	45
Configuring a Bootable Virtual Drive	45
Deleting (Clearing) a Storage Configuration	45
Chapter 6: Intel® IT/IR RAID Configuration	47
IM and IME Configuration Overview	47
Features	47
Creating IM and IME Volumes	48
Creating an IM Volume	49
Creating an IME Volume	50
Creating a Second IM or IME Volume	51
Managing Hot Spares	51
Other Configuration Tasks	53

Viewing Volume Properties	53
Synchronizing an Array	53
Activating an Array	53
Deleting an Array	54
Locating a Drive or Multiple Drives in a Volume	54
Selecting a Boot Disk	55
IS Configuration Overview	55
Creating IS Volumes	56
Creating a Second IS Volume	57
Other Configuration Tasks	58
Viewing IS Volume Properties	58
Activating an Array	58
Deleting an Array	58
Locating a Disk Drive, or Multiple Disk Drives in a Volume	59
Selecting a Boot Disk	59
Chapter 7: Intel® RAID BIOS Console 2 Utility	61
Quick Configuration Steps	61
Detailed Configuration Steps using the Intel® RAID BIOS Console 2	62
Start the Intel® RAID BIOS Console 2 Utility	62
Screen and Option Descriptions	62
Setting Up a RAID Array Using the Configuration Wizard	68
Creating RAID 0, 1, 5, or 6 using Intel® RAID BIOS Console 2 (detailed)	71
Creating RAID 10, RAID 50, and RAID 60 using Intel® RAID BIOS Console 2	75
Setting Drive Parameters	80
Creating a Hot Spare	81
Viewing Event Details	84
Chapter 8: Intel® RAID Web Console 2	87
Configuration Functions	87
Monitoring Functions	87
Maintenance Functions	88
Hardware and Software Requirements	88
Installing the Intel® RAID Web Console 2 on a Microsoft Windows* Operating System	88
Installing the Intel® RAID Web Console 2 on Linux or SuSE* Linux Enterprise Server	90
Intel® RAID Web Console 2 Support and Installation on VMWare	90
Installing Intel® RAID Web Console 2 for VMWare Classic	90
Uninstalling Intel® RAID Web Console 2 for VMWare	90
Installing Intel® RAID Web Console 2 Support on the VMWare ESX	91
Starting the Intel® RAID Web Console 2	96
Intel® RAID Web Console 2 Screens	96
Physical/Virtual View Panel	99
Properties/Operations/Graphical View Panel	100
Event Log Panel	101
Menu Bar/Manage Menu	102

Menu Bar/Go To Menu	102
File Menu/Log Menu	102
File Menu/Tool Menu	102
File Menu/Help Menu	102
Drive Configuration Tasks	102
Create Virtual Drive	103
Creating a Virtual Drive Using Simple Configuration	103
Creating a Virtual Drive Using Advanced Configuration	107
Creating a Spanned Disk Group	114
Creating Hot Spares	115
Setting Adjustable Task Rates	117
Adding a Drive to a Virtual Disk	118
Removing a Drive or Changing the RAID Level of a Virtual Disk	121
Changing Virtual Disk Properties	122
Deleting a Virtual Disk	123
Managing Configurations	124
Saving a Configuration to Disk	124
Clearing a Configuration from a Controller	126
Adding a Configuration from a File	127
Monitoring System Events and Devices	128
Monitoring System Events	128
Monitoring Controllers	129
Monitoring Disk Drives and Other Physical Devices	130
Monitoring Virtual Disks	133
Monitoring Enclosures	133
Monitoring Battery Backup Units	134
Battery Learn Cycle	135
Monitoring Rebuilds and Other Processes	137
Maintaining and Managing Storage Configurations	138
Initializing a Virtual Disk	138
Running a Consistency Check	139
Scanning for New Drives	140
Rebuilding a Drive	140
Removing a Drive	141
Flashing the Firmware	142
Enabling RAID Premium Features	142
Enabling Full Disk Encryption feature	142
Enabling Snapshot feature	151
Enabling Super Sized Cache	161
Appendix A: Creating a Virtual Drive Using Advanced Configuration	165
Appendix B: Events and Messages	173

List of Figures

Figure 1. RAID 0 - Data Striping.....	10
Figure 2. RAID 1 - Disk Mirroring/Disk Duplexing	11
Figure 3. RAID 5 - Data Striping with Striped Parity.....	12
Figure 4. Example of Distributed Parity across Two Blocks in a Stripe (RAID 6).....	13
Figure 5. Integrated Mirroring Enhanced with Three Disks	14
Figure 6. RAID 10 - Combination of RAID 1 and RAID 0	15
Figure 7. RAID 50 - Combination of RAID 5 and RAID 0	16
Figure 8. RAID 60 Level Virtual Drive	18
Figure 9. Intel® Embedded Server RAID BIOS Configuration Utility Screen	40
Figure 10. Adapter Properties Screen.....	49
Figure 11. Create New Array Screen	50
Figure 12. Manage Array Screen	52
Figure 13. Adapter Properties Screen.....	56
Figure 14. Create New Array Screen	57
Figure 15. Intel® RAID BIOS Console 2 Menu.....	63
Figure 16. Intel® RAID BIOS Console 2 - Controller Selection	64
Figure 17. Controller Properties	64
Figure 18. Additional Controller Properties	65
Figure 19. Intel® RAID BIOS Console 2 - Configuration Types	68
Figure 20. Selecting Configuration	69
Figure 21. Intel® RAID BIOS Console 2 - Configuration Methods	70
Figure 22. Intel® RAID BIOS Console 2 - Add Physical Drives to Array	71
Figure 23. Intel® RAID BIOS Console 2 - Set Array Properties	72
Figure 24. Intel® RAID BIOS Console 2 - Confirm Configuration.....	73
Figure 25. Intel® RAID BIOS Console 2 - Initialization Speed Setting	74
Figure 26. Intel® RAID BIOS Console 2 – Multiple Disk Groups for RAID 10, 50, or 60.....	75
Figure 27. Intel® RAID BIOS Console 2 – Spanning Multiple Arrays.....	76
Figure 28. Intel® RAID BIOS Console 2 – Viewing Completed Settings.....	77
Figure 29. Intel® RAID BIOS Console 2 – Initialization Settings.....	78
Figure 30. Intel® RAID BIOS Console 2 – RAID 10 Final Screen	78
Figure 31. Intel® RAID BIOS Console 2 – RAID 10 Properties Screen	79
Figure 32. Intel® RAID BIOS Console 2 – RAID 50 Properties Screen	79
Figure 33. Intel® RAID BIOS Console 2 – Choosing a Hot Spare Drive	81
Figure 34. Intel® RAID BIOS Console 2 – Setting a Hot Spare Drive	82
Figure 35. Intel® RAID BIOS Console 2 – Viewing Hot Spare	82
Figure 36. Intel® RAID BIOS Console 2 – Main Screen showing Hot Spare Drive.....	83
Figure 37. Intel® RAID BIOS Console 2 – Event Information Screen	84
Figure 38. Intel® RAID BIOS Console 2 – Selecting Events to View	85
Figure 39. Intel® RAID BIOS Console 2 – Viewing an Event.....	85
Figure 40. Intel® RAID Web Console 2 – Customer Information Screen	89
Figure 41. Setup Type Screen	89
Figure 42. Intel® RAID Web Console 2 – Select Server Screen	96

Figure 43. Intel® RAID Web Console 2 – Login Screen	97
Figure 44. Intel® RAID Web console 2 dashboard	98
Figure 45. Intel® RAID Web Console 2 – Main Screen	99
Figure 46. Intel® RAID Web Console 2 – Operations Tab.....	100
Figure 47. Intel® RAID Web Console 2 – Graphical Tab (Optional feature)	101
Figure 48. Virtual Drive Creation Menu.....	104
Figure 49. Virtual Drive Creation Mode.....	104
Figure 50. Create Virtual Drive Screen	105
Figure 51. Create Virtual Drive - Summary Window	106
Figure 52. Option to Create Additional Virtual Drives	106
Figure 53. Option to Close the Configuration Wizard	107
Figure 54. Virtual Drive Creation Menu.....	108
Figure 55. Virtual Drive Creation Mode.....	108
Figure 56. Create Drive Group Settings Screen	109
Figure 57. Span 0 of Drive Group 0	110
Figure 58. Span 0 and Span 1 of Drive Group 0.....	110
Figure 59. Virtual Drive Settings Window	111
Figure 60. New Virtual Drive 0.....	112
Figure 61. Create Virtual Drive Summary Window	112
Figure 62. Option to Create Additional Virtual Drives	113
Figure 63. Option to Close the Configuration Wizard	113
Figure 64. Assign Global Hotspare	115
Figure 65. Assign Dedicated Hotspare	116
Figure 66. Select Hotspare Drive.....	116
Figure 67. Set Adjustable Task Rates	117
Figure 68. Starting Modify Drive Group	119
Figure 69. Select RAID level to migrate.....	119
Figure 70. Selecting Drives to Add	120
Figure 71. Changing RAID Level	120
Figure 72. Selecting Drives to Remove	121
Figure 73. Changing RAID Level	122
Figure 74. Set Virtual Disk Properties.....	123
Figure 75. Save Configuration to File	124
Figure 76. Save Configuration Dialog Box.....	125
Figure 77. Clear Configuration.....	126
Figure 78. Add Saved Configuration.....	127
Figure 79. Event Information Window	128
Figure 80. Controller Information	129
Figure 81. Physical Drive Information	130
Figure 82. Locating a Physical Drive	131
Figure 83. Patrol Read Configuration	132
Figure 84. Virtual Drive Properties.....	133
Figure 85. Enclosure Information.....	134
Figure 86. Battery Backup Unit Information	135
Figure 87. Battery Backup Unit Operations	136
Figure 88. Group Show Progress Window.....	137
Figure 89. Selecting Initialize	138

Figure 90. Group Consistency Check Window.....	139
Figure 91. Scan for Foreign Configuration	140
Figure 92. Preparing Drive for Removal.....	141
Figure 93. Check Controller Security status.....	143
Figure 94. Check Drive Security status.....	144
Figure 95. Enable Drive Security.....	144
Figure 96. Start Security Wizard	145
Figure 97. Enter Security Key Identifier.....	145
Figure 98. Enter Security Key	146
Figure 99. Enter Pass Phrase	147
Figure 100. Confirm Enable Drive Security	147
Figure 101. Check Drive Security Enabled status.....	148
Figure 102. Select Full Disk Encryption	148
Figure 103. Create RAID Virtual Drive with FDE enabled.....	149
Figure 104. Instant Secure Erase.....	150
Figure 105. Confirm Secure Erase.....	150
Figure 106. Enable MegaRAID Recovery	152
Figure 107. Enter the Capacity for Snapshot Respository	152
Figure 108. Confirm Enable Snapshot	153
Figure 109. Snapshot Base is shown.....	154
Figure 110. Enter Snapshot Name.....	154
Figure 111. Create Snapshot	155
Figure 112. Create View.....	155
Figure 113. Set MegaRAID Recovery Properties.....	156
Figure 114. Disable MegaRAID Recovery	157
Figure 115. Confirm Disable Snapshots.....	157
Figure 116. Adapter Selection.....	158
Figure 117. Selecting Snapshot Base	158
Figure 118. Selecting Advanced Operations.....	159
Figure 119. Selecting Rollback.....	159
Figure 120. Selecting a snapshot.....	160
Figure 121. Confirm Page	160
Figure 122. Rollback operation is done.....	161
Figure 123. Create SSC from Dashboard	162
Figure 124. Create SSC Drive Group.....	162
Figure 125. Create SSCD name	163
Figure 126. SSC Summary	163
Figure 127. SSCD status shown	164
Figure 128. Delete SSCD.....	164
Figure 129. Virtual Drive Creation Menu	165
Figure 130. Virtual Drive Creation Mode	166
Figure 131. Create Drive Group Settings Screen.....	167
Figure 132. Span 0 of Drive Group 0	168
Figure 133. Span 0 and Span 1 of Drive Group 0.....	169
Figure 134. Virtual Drive Settings Window.....	170
Figure 135. New Virtual Drive 0	171
Figure 136. Create Virtual Drive Summary Window.....	171

Figure 137. Option to Create Additional Virtual Drives	172
Figure 138. Option to Close the Configuration Wizard	172

List of Tables

Table 1. RAID 0 Overview	10
Table 2. RAID 1 Overview	11
Table 3. RAID 5 Overview	12
Table 4. RAID 6 Overview	12
Table 5. RAID 1E Overview	14
Table 6. RAID 10 Overview	15
Table 7. RAID 50 Overview	16
Table 8. RAID 60 Overview	17
Table 9. RAID Levels and Fault Tolerance	19
Table 10. RAID Levels and Performance.....	20
Table 11. RAID Levels and Capacity	22
Table 12. Factors to Consider for Array Configuration	26
Table 13. Intel® RAID BIOS Console 2 Toolbar Icon Descriptions	62
Table 14. MFI Messages	174

1 Overview

The software described in this document is designed for use with Intel® RAID controllers, and with on-serverboard RAID solutions that use the Intel® RAID Software Stack 3 (driver package names begin with “ir3”), Embedded Server RAID Technology 2 (driver package names begin with ESRT2) or Intel® IT/IR RAID.

Supported Hardware

This manual covers the software stack that is shared by multiple Intel® server products:

- Intel® Embedded Server RAID Technology 2 (ESRT2) on the Intel® Enterprise South Bridge 2 (ESB2) in the chipset, the Intel® I/O Controller Hub 9R (ICH9R), the Intel® 3420 PCH chipset, Intel® C200 series chipset and Intel® C600 series chipset used in the following:
 - Intel® Server Board S1200BTL/S1200BTS
 - Intel® Server Boards based on the Intel® S5000 and S7000 chipsets
 - Intel® Server Boards based on the Intel® 5500/5520 chipset with the Intel® I/O Controller Hub 10R (ICH10R)
 - Intel® Server Boards that include the LSI* 1064e SAS (Serially attached SCSI) controller and some that include the LSI* 1068 SAS controller
 - Intel® Server Boards S3420GP
 - Intel® Server Boards S3200SH and X38ML
 - Intel® SAS Entry RAID Module AXX4SASMOD (when the module is in ESRTII mode)
 - Intel® RAID Controller SASMF8I

Intel® Embedded Server RAID Technology 2 provides driver based RAID modes 0,1, and 10 with an optional RAID 5 mode provided by Intel® RAID C600 Upgrade Key RKSATA4R5, RKSATA8R5, RKSAS4R5 or RKSAS8R5.

Intel® Embedded Server RAID Technology 2 provides driver-based RAID modes 0, 1, and 10 with an optional RAID 5 mode provided by the Intel® RAID Activation Key AXXRAKSW5 on the ESB2 and LSI* 1064e on some models of Intel® server boards. ESB2 supports SATA only.

LSI* SAS 1064e and 1068 provide SATA (Serial ATA) and SAS support. Not all 1068 SAS boards provide Intel® Embedded Server RAID Technology 2 modes.

Intel® Embedded Server RAID Technology 2 must be enabled in the server system BIOS before it is available. Intel® Embedded Server RAID Technology 2 is limited to a maximum of eight drives including hot spare(s). Expander devices are not yet supported by ESRT2.
- Intel® IT/IR RAID solutions with below Intel® IT/IR RAID controllers:
 - Intel® RAID Controller SASWT4I
 - Intel® RAID Controller SASUC8I

- Intel® RAID SAS Riser Controller AFCSASRISER in Intel® Server System S7000FC4UR without Intel® SAS RAID Activation Key AXXRAKSAS2 installed.
- Intel® RAID SAS Riser Controller AFCSASRISER in Intel® Server System S7000FC4UR without Intel® SAS RAID Activation Key AXXRAKSAS2 installed
- Intel® SAS Entry RAID Module AXX4SASMOD
- Intel® 6G SAS PCIe Gen2 RAID Module RMS2LL080 and RMS2LL040
- Intel® Integrated RAID Technology on the Intel® ROMB solutions. Server boards and systems include:
 - Intel® Server Board S5000PSL (Product code: S5000PSLROMB)
 - Intel® Server System SR1550AL (Product code: SR1550ALSAS)
 - Intel® Server System SR2500 (Product code: SR2500LX)
 - Intel® Server System SR4850HW4s
 - Intel® Server System SR6850HW4s
 - Intel® Server System S7000FC4UR with a SAS riser card
 - Intel® Server Boards S3420GP, S5520HC/S5500HCV, S5520UR, S5520SC, and S5500WB12V/S5500WB with the Intel® Integrated RAID Controller SROMBSASMR

Systems using the Intel® RAID Controller SROMBSAS18E provide XOR RAID modes 0, 1, 5, 10, and 50 when the optional Intel® RAID Activation Key AXXRAK18E and a DDR2 400 MHz ECC DIMM are installed.

Systems using the Intel® RAID Controller SROMBSASFC or SROMBSASMP2 require the optional Intel® RAID Activation Key AXXRAKSAS2 and a DDR2 667 MHz ECC DIMM to provide RAID modes 0, 1, 5, 6, 10, 50, and 60.

The Intel® Integrated RAID Controller SROMBSASMR has a specially designed connector that only fits Intel® Server Boards S5520HC/S5500HCV, S5520UR, S5520SC, and S5500WB12V/S5500WB.

Note: This manual does not include the software RAID modes provided by the SAS riser card on the Intel® Server System S7000FC4UR. This manual does not include the RAID modes provided by the FALSASMP2 without Intel® RAID Activation Key AXXRAKSAS2.

- Intel® Intelligent RAID used on the discrete Intel® RAID controllers RS25AB080, RS25SB008, RS25DB080, RS25NB008, RS2VB080, RS2VB040, RT3WB080, RS2SG244, RS2WG160, RMS2MH080, RMS2AF080, RMS2AF040, RS2BL080, RS2BL040, RS2BL080DE, RS2BL080SNGL, RS2PI008, RS2PI008DE, RS2MB044, RS2WC080, RS2WC040, SROMBSASMR, SRCASAWB, SRCASARB, SRCASAJV, SRCASABB8I, SRCASASLS4I, SRCASASPH16I, SROMBSASFC, SROMBSASMP2, SROMBSAS18E, SRCASAS18E and SRCASAS144E.
 - The first generation SAS controllers (SRCASAS18E, SRCASAS144E, SROMBSAS18E) provide XOR RAID modes 0, 1, 5, 10, and 50
 - The second generation SAS controller (including SRCASAWB, SRCASARB, SRCASAJV, SRCASABB8I, SRCASASLS4I, SRCASASPH16I, SROMBSASFC, SROMBSASMP2, SROMBSASMR) provides XOR RAID modes 0, 1, 5, 6, 10, 50, and 60.
 - The Intel® 6G SAS PCIe Gen 2 RAID Controllers (including RS25AB080, RS25SB008, RS25DB080, RS25NB008, RS2VB080, RS2VB040, RT3WB080, RS2SG244, RS2WG160, RS2BL080, RS2BL080SNGL, RS2BL080DE,

RS2BL040, RS2PI008, RS2PI008DE, RS2MB044, RS2WC080, RS2WC040, RMS2MH080, RMS2AF080 and RMS2AF040) support SAS 2.0 new features with XOR RAID modes 0, 1, 5, 6, 10, 50, and 60. (RS2WC080 and RS2WC040 are entry level hardware RAID controllers and do not support RAID 6 and 60; RMS2AF080 and RMS2AF040 are entry level hardware RAID controllers and do not support RAID 10, 6 and 60.)

For more details, refer to the *Technical Product Specification (TPS)* or *Hardware User's Guide (HWUG)* for the RAID controllers.

Note: *The Intel® RAID Controllers RMS2AF080, RMS2AF040, RS2WC080, and RS2WC040 only support strip sizes of 8KB, 16KB, 32KB, and 64KB. Also, their Cache Policy only supports Write Through, Direct I/O, and Normal RAID (No Read Ahead). For more details, refer to their Hardware User's Guide (HWUG).*

This manual does not include information about native SATA or SAS-only modes of the RAID controllers.

Two versions of the Intel® RAID Controller RS2BL080 are available - RS2BL080, RS2BL080DE.

All features on RS2BL080 are supported on RS2BL080DE. In addition, RS2BL080DE provides one more feature of FDE (Full Disk Encryption) that RS2BL080 doesn't support.

Two versions of the Intel® RAID Controller RS2PI008 are available - RS2PI008, RS2PI008DE.

All features on RS2PI008 are supported on RS2PI008DE. In addition, RS2PI008DE provides one more feature of FDE (Full Disk Encryption) that RS2PI008 doesn't support.

Caution: *Some levels of RAID are designed to increase the availability of data and some to provide data redundancy. However, installing a RAID controller is not a substitute for a reliable backup strategy. It is highly recommended you back up data regularly via a tape drive or other backup strategy to guard against data loss. It is especially important to back up all data before working on any system components and before installing or changing the RAID controller or configuration.*

Software

Intel® Embedded Server RAID Technology 2, Intel® IT/IR RAID and Intel® Integrated Server RAID controllers include a set of software tools to configure and manage RAID systems. These include:

- Intel® RAID controller software and utilities: The firmware installed on the RAID controller provides pre-operating system configuration.
 - For Intel® Embedded Server RAID Technology 2, press <Ctrl> + <E> during the server boot to enter the BIOS configuration utility.
 - For Intel® IT/IR RAID, press <Ctrl> + <C> during the server boot to enter the LSI MPT* SAS BIOS Configuration Utility
 - For Intel® Integrated Server RAID, press <Ctrl> + <G> during the server boot to enter the RAID BIOS Console II.
- Intel® RAID Controller Drivers: Intel provides software drivers for the following operating systems.
 - Microsoft Windows 2000*, Microsoft Windows XP*, and Microsoft Windows Server 2003* (32-bit and 64-bit editions)
 - Red Hat* Enterprise Linux 3.0, 4.0, and 5.0 (with service packs; X86 and X86-64)

- SuSE* Linux Enterprise Server 9.0, SuSE* Linux Enterprise Server 10, and SuSE* Linux Enterprise Server 11(with service packs; X86 and X86-64)
- VMWare* ESX 4i

- Note: Only the combinations of controller, driver, and Intel® Server Board or System listed in the Tested Hardware and Operating System List (THOL) were tested. Check the supported operating system list for both your RAID controller and your server board to verify operating system support and compatibility.*
- Intel® RAID Web Console 2: A full-featured graphical user interface (GUI) utility is provided to monitor, manage, and update the RAID configuration.

RAID Terminology

RAID is a group of physical disks put together to provide increased I/O (Input/Output) performance (by allowing multiple, simultaneous disk access), fault tolerance, and reliability (by reconstructing failed drives from remaining data). The physical drive group is called an array, and the partitioned sets are called virtual disks. A virtual disk can consist of a part of one or more physical arrays, and one or more entire arrays.

Using two or more configured RAID arrays in a larger virtual disk is called spanning. It is represented by a double digit in the RAID mode/type (10, 50, 60).

Running more than one array on a given physical drive or set of drives is called a sliced configuration.

The only drive that the operating system works with is the virtual disk, which is also called a virtual drive. The virtual drive is used by the operating system as a single drive (lettered storage device in Microsoft Windows*).

The RAID controller is the mastermind that must configure the physical array and the virtual disks, and initialize them for use, check them for data consistency, allocate the data between the physical drives, and rebuild a failed array to maintain data redundancy. The features available per controller are highlighted later in this document and in the hardware guide for the RAID controller.

The common terms used when describing RAID functions and features can be grouped into two areas: fault tolerance (data protection and redundancy) and performance.

Fault Tolerance

Fault tolerance describes a state in which even with a drive failure, the data on the virtual drive is still complete and the system is available after the failure and during repair of the array. Most RAID modes are able to endure a physical disk failure without compromising data integrity or processing capability of the virtual drive.

RAID mode 0 is not fault tolerant. With RAID 0, if a drive fails, then the data is no longer complete and no longer available. Backplane fault tolerance can be achieved by a spanned array where the arrays are on different backplanes.

True fault tolerance includes the automatic ability to restore the RAID array to redundancy so that another drive failure will not destroy its usability.

Hot Spare

True fault tolerance requires the availability of a spare disk that the controller can add to the array and use to rebuild the array with the data from the failed drive. This spare disk is called a hot spare. It must be a part of the array before a disk failure occurs. A hot-spare drive is a physical drive that is maintained by the RAID controller but not actually used for data storage in the array unless another drive fails. Upon failure of one of the array's physical drives, the hot-spare drive is used to hold the recreated data and restore data redundancy.

Hot-spare drives can be global (available to any array on a controller) or dedicated (only usable by one array). There can be more than one hot spare per array and the drive of the closest capacity is used. If both dedicated and global hot-spare drives are available, then the dedicated drive is used first. If the hot swap rebuild fails, then that hot spare is also marked failed. Since RAID 0 is not redundant, there is no hot spare value.

If a hot-spare drive is not an option, then it is possible to perform a hot or cold swap of the failed drive to provide the new drive for rebuild after the drive failure. A swap is the manual substitution of a replacement drive in a disk subsystem. If a swap is performed while the system is running, it is a hot swap. A hot swap can only be performed if the backplane and enclosure support it. If the system does not support hot-swap drives, then the system must be powered down before the drive swap occurs. This is a cold swap.

In all cases (hot spare, hot swap, or cold swap), the replacement drive must be at least as large as the drive it replaces. In all three cases, the failed drive is removed from the array. If using a hot spare, then the failed drive can remain in the system. When a hot spare is available and an automatic rebuild starts, the failed drive may be automatically removed from the array before the utilities detect the failure. Only the event logs show what happened.

If the system is shut down during the rebuild, all rebuilds should automatically restart on reboot.

Note: *If running a sliced configuration (RAID 0, RAID 5, and RAID 6 on the same set of physical drives), then the rebuild of the spare will not occur until the RAID 0 array is deleted.*

On Intel® RAID Controller RS2WC080 and RS2WC040, if Virtual Drive is in degrade mode due to failed physical drive, auto rebuild is not supported for hot-plugged drive until a manual selection is made by users. As part of JBOD implementation for Intel® RAID Controller RS2WC080 and RS2WC040, all new drives that are hot-plugged will automatically become JBOD. Users need to manually move the JBOD drive to Unconfigured Good and auto rebuild starts after that. For more details, refer to Hardware User's Guide (HWUG) for above controllers.

Data Redundancy

Data redundancy is provided by mirroring or by disk striping with parity stripes.

- Disk mirroring is found only in RAID 1 and 10. With mirroring, the same data simultaneously writes to two disks. If one disk fails, the contents of the other disk can be used to run the system and reconstruct the failed array. This provides 100% data redundancy but uses the most drive capacity, since 50% of the total capacity is available. Until a failure occurs, both mirrored disks contain the same data at all times. Either drive can act as the operational drive.
- Parity is the ability to recreate data by using a mathematical calculation derived from multiple data sets. Parity is basically a checksum of all the data known as the

“ABCsum”. When drive A fails, the controller uses the ABCsum to calculate what remains on drives B+C. The remainder must be recreated onto new drive A.

Parity can be dedicated (all parity stripes are placed on the same drive) or distributed (parity stripes are spread across multiple drives). Calculating and writing parity slows the write process but provides redundancy in a much smaller space than mirroring. Parity checking is also used to detect errors in the data during consistency checks and patrol reads.

RAID 5 uses distributed parity and RAID 6 uses dual distributed parity (two different sets of parity are calculated and written to different drives each time.) RAID modes 1 and 5 can survive a single disk failure, although performance may be degraded, especially during the rebuild. RAID modes 10 and 50 can survive multiple disk failures across the spans, but only one failure per array. RAID mode 6 can survive up to two disk failures. RAID mode 60 can sustain up to two failures per array.

Data protection is also provided by running calculations on the drives to make sure data is consistent and that drives are good. The controller uses consistency checks, background initialization, and patrol reads. You should include these in regular maintenance schedules.

- The consistency check operation verifies that data in the array matches the redundancy data (parity or checksum). This is not provided in RAID 0 in which there is no fault tolerance.
- Background initialization is a consistency check that is forced five minutes after the creation of a virtual disk. Background initialization also checks for media errors on physical drives and ensures that striped data segments are the same on all physical drives in an array.
- Patrol read checks for physical disk errors that could lead to drive failure. These checks usually include an attempt at corrective action. Patrol read can be enabled or disabled with automatic or manual activation. This process starts only when the RAID controller is idle for a defined period of time and no other background tasks are active, although a patrol read check can continue to run during heavy I/O processes.

Enclosure Management

Enclosure management is the intelligent monitoring of the disk subsystem by software or hardware usually within a disk enclosure. It increases the ability for the user to respond to a drive or power supply failure by monitoring those sub systems.

Performance

Performance improvements come from multiple areas including disk striping and disk spanning, accessing multiple disks simultaneously, and setting the percentage of processing capability to use for a task.

Disk Striping

Disk striping writes data across all of the physical disks in the array into fixed size partitions or stripes. In most cases, the stripe size is user-defined. Stripes do not provide redundancy but improve performance since striping allows multiple physical drives to be accessed at the same time. These stripes are interleaved in a repeated sequential manner and the controller knows where data is stored. The same stripe size should be kept across RAID arrays.

Terms used with strip sizing are listed below.

- Strip size: One disk section
- Stripe size: Total of one set of strips across all data disks, not including parity stripes
- Stripe width: The number of disks involved

Disk Spanning

Disk spanning allows more than one array to be combined into a single virtual drive. The spanned arrays must have the same stripe size and must be contiguous. Spanning alone does not provide redundancy but RAID modes 10, 50, and 60 all have redundancy provided in their pre-spanned arrays through RAID 1, 5, or 6.

Note: *Spanning two contiguous RAID 0 drives does not produce a new RAID level or add fault tolerance. It does increase the size of the virtual volume and improves performance by doubling the number of spindles. Spanning for RAID 10, RAID 50, and RAID 60 requires two to eight arrays of RAID 1, 5, or 6 with the same stripe size and that always uses the entire drive.*

CPU Usage

Resource allocation provides the user with the option to set the amount of compute cycles to devote to various tasks, including the rate of rebuilds, initialization, consistency checks, and patrol read. Setting resource to 100% gives total priority to the rebuild. Setting it at 0% means the rebuild will only occur if the system is not doing anything else. The default rebuild rate is 30%.

2 RAID Levels

The RAID controller supports RAID levels 0, 1E, 5, 6, 10, 50, and 60. The supported RAID levels are summarized below. In addition, it supports independent drives (configured as RAID 0). This chapter describes the RAID levels in detail.

Summary of RAID Levels

- **RAID 0:** Uses striping to provide high data throughput, especially for large files in an environment that does not require fault tolerance. In Intel® IT/IR RAID, RAID 0 is also called Integrated Striping (IS), which supports striped arrays with two to ten disks.
- **RAID 1:** Uses mirroring so that data written to one disk drive simultaneously writes to another disk drive. This is good for small databases or other applications that require small capacity but complete data redundancy. In Intel® IT/IR RAID, RAID 1 is also called Integrated Mirroring (IM) which supports two-disk mirrored arrays and hot-spare disks.
- **RAID 5:** Uses disk striping and parity data across all drives (distributed parity) to provide high data throughput, especially for small random access.
- **RAID 6:** Uses distributed parity, with two independent parity blocks per stripe, and disk striping. A RAID 6 virtual disk can survive the loss of two disks without losing data.
- **RAID 1ME:** Integrated Mirroring Enhanced (IME) which supports mirrored arrays with three to ten disks, plus hot-spare disks. This is implemented in Intel® IT/IR RAID.
- **RAID 10:** A combination of RAID 0 and RAID 1, consists of striped data across mirrored spans. It provides high data throughput and complete data redundancy but uses a larger number of spans.
- **RAID 50:** A combination of RAID 0 and RAID 5, uses distributed parity and disk striping and works best with data that requires high reliability, high request rates, high data transfers, and medium-to-large capacity.

***Note:** It is not recommended to have a RAID 0, RAID 5, and RAID 6 virtual disk in the same physical array. If a drive in the physical array has to be rebuilt, the RAID 0 virtual disk will cause a failure during the rebuild.*

- **RAID 60:** A combination of RAID 0 and RAID 6, uses distributed parity, with two independent parity blocks per stripe in each RAID set, and disk striping. A RAID 60 virtual disk can survive the loss of two disks in each of the RAID 6 sets without losing data. It works best with data that requires high reliability, high request rates, high data transfers, and medium-to-large capacity.

Selecting a RAID Level

To ensure the best performance, select the optimal RAID level when the system drive is created. The optimal RAID level for a disk array depends on a number of factors:

- The number of physical drives in the disk array
- The capacity of the physical drives in the array
- The need for data redundancy
- The disk performance requirements

RAID 0 - Data Striping

RAID 0 provides disk striping across all drives in the RAID array. RAID 0 does not provide any data redundancy, but does offer the best performance of any RAID level. RAID 0 breaks up data into smaller segments, and then stripes the data segments across each drive in the array. The size of each data segment is determined by the stripe size. RAID 0 offers high bandwidth.

Note: RAID level 0 is not fault tolerant. If a drive in a RAID 0 array fails, the whole virtual disk (all physical drives associated with the virtual disk) will fail.

By breaking up a large file into smaller segments, the RAID controller can use both SAS drive and SATA drives to read or write the file faster. RAID 0 involves no parity calculations to complicate the write operation. This makes RAID 0 ideal for applications that require high bandwidth but do not require fault tolerance.




Figure 1. RAID 0 - Data Striping

Table 1. RAID 0 Overview

Uses	Provides high data throughput, especially for large files. Any environment that does not require fault tolerance.
Strong Points	Provides increased data throughput for large files. No capacity loss penalty for parity.
Weak Points	Does not provide fault tolerance or high bandwidth. If any drive fails, all data is lost.
Drives	1 to 32

RAID 1 - Disk Mirroring/Disk Duplexing

In RAID 1, the RAID controller duplicates all data from one drive to a second drive. RAID 1 provides complete data redundancy, but at the cost of doubling the required data storage capacity. [Table 2](#) provides an overview of RAID 1.

Table 2. RAID 1 Overview

Uses	Use RAID 1 for small databases or any other environment that requires fault tolerance but small capacity.
Strong Points	Provides complete data redundancy. RAID 1 is ideal for any application that requires fault tolerance and minimal capacity.
Weak Points	Requires twice as many disk drives. Performance is impaired during drive rebuilds.
Drives	2 to 32 (must be an even number of drives)




Figure 2. RAID 1 - Disk Mirroring/Disk Duplexing

RAID 5 - Data Striping with Striped Parity

RAID 5 includes disk striping at the block level and parity. Parity is the data's property of being odd or even, and parity checking detects errors in the data. In RAID 5, the parity information is written to all drives. RAID 5 is best suited for networks that perform a lot of small I/O transactions simultaneously.

RAID 5 addresses the bottleneck issue for random I/O operations. Because each drive contains both data and parity, numerous writes can take place concurrently.

[Table 3](#) provides an overview of RAID 5.

Table 3. RAID 5 Overview

Uses	Provides high data throughput, especially for large files. Use RAID 5 for transaction processing applications because each drive can read and write independently. If a drive fails, the RAID controller uses the parity drive to recreate all missing information. Use also for office automation and online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates.
Strong Points	Provides data redundancy, high read rates, and good performance in most environments. Provides redundancy with lowest loss of capacity.
Weak Points	Not well suited to tasks requiring lot of writes. Suffers more impact if no cache is used (clustering). If a drive is being rebuilt, disk drive performance is reduced. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
Drives	3 to 32

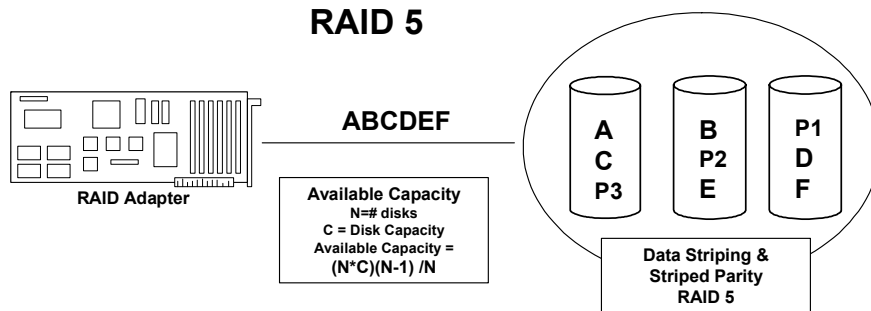


Figure 3. RAID 5 - Data Striping with Striped Parity

RAID 6 - Distributed Parity and Disk Striping

RAID 6 is similar to RAID 5 (disk striping and parity), but instead of one parity block per stripe, there are two. With two independent parity blocks, RAID 6 can survive the loss of two disks in a virtual disk without losing data.


Table 4 provides an overview of RAID 6.

Table 4. RAID 6 Overview

Uses	<p>Provides a high level of data protection through the use of a second parity block in each stripe. Use RAID 6 for data that requires a high level of protection from loss.</p> <p>In the case of a failure of one drive or two drives in a virtual disk, the RAID controller uses the parity blocks to recreate the missing information. If two drives in a RAID 6 virtual disk fail, two drive rebuilds are required, one for each drive. These rebuilds do not occur at the same time. The controller rebuilds one failed drive at a time.</p> <p>Use for office automation and online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates.</p>
-------------	--

Strong Points	Provides data redundancy, high read rates, and good performance in most environments. Can survive the loss of two drives or the loss of a drive while another drive is being rebuilt. Provides the highest level of protection against drive failures of all of the RAID levels. Read performance is similar to that of RAID 5.
Weak Points	Not well suited to tasks requiring lot of writes. A RAID 6 virtual disk has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Disk drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes. RAID 6 costs more because of the extra capacity required by using two parity blocks per stripe.
Drives	3 to 32

The following figure shows a RAID 6 data layout. The second set of parity drives are denoted by *Q*. The *P* drives follow the RAID 5 parity scheme.



Parity is distributed across all drives in the array. When only three hard drives are available for RAID 6, the situation has to be that P equals Q equals original data, which means that the original data has three copies across the three hard drives.

Figure 4. Example of Distributed Parity across Two Blocks in a Stripe (RAID 6)

RAID IME

An IME volume can be configured with up to ten mirrored disks (one or two global hot spares can also be added). [Figure 5](#) shows the logical view and physical view of an Integrated Mirroring Enhanced (IME) volume with three mirrored disks. Each mirrored stripe is written to a disk and mirrored to an adjacent disk. This type of configuration is also called RAID 1E.




Figure 5. Integrated Mirroring Enhanced with Three Disks

Table 5. RAID 1E Overview

Uses	Use RAID 1E for small databases or any other environment that requires fault tolerance but small capacity.
Strong Points	Provides complete data redundancy. RAID 1E is ideal for any application that requires fault tolerance and minimal capacity.
Weak Points	Requires twice as many disk drives. Performance is impaired during drive rebuilds.
Drives	3 to 10

RAID 10 - Combination of RAID 1 and RAID 0

RAID 10 is a combination of RAID 0 and RAID 1. RAID 10 consists of stripes across mirrored drives. RAID 10 breaks up data into smaller blocks and then mirrors the blocks of data to each RAID 1 RAID set. Each RAID 1 RAID set then duplicates its data to its other drive. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set. RAID 10 supports up to eight spans.

Table 6 provides an overview of RAID 10.

Table 6. RAID 10 Overview

Uses	Appropriate when used with data storage that requires 100 percent redundancy of mirrored arrays and that needs the enhanced I/O performance of RAID 0 (striped arrays). RAID 10 works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate to medium capacity.
Strong Points	Provides both high data transfer rates and complete data redundancy.
Weak Points	Requires twice as many drives as all other RAID levels except RAID 1.
Drives	4 - 240




Figure 6. RAID 10 - Combination of RAID 1 and RAID 0

RAID 50 - Combination of RAID 5 and RAID 0

RAID 50 provides the features of both RAID 0 and RAID 5. RAID 50 includes both parity and disk striping across multiple arrays. RAID 50 is best implemented on two RAID 5 disk arrays with data striped across both disk groups.

RAID 50 breaks up data into smaller blocks and then stripes the blocks of data to each RAID 5 disk set. RAID 5 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks and then writes the blocks of data and parity to each drive in the array. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.

RAID level 50 supports up to eight spans and tolerates up to eight drive failures, though less than total disk drive capacity is available. Though multiple drive failures can be tolerated, only one drive failure can be tolerated in each RAID 1 level array.

Table 7 provides an overview of RAID 50.

Table 7. RAID 50 Overview

Uses	Appropriate when used with data that requires high reliability, high request rates, high data transfer, and medium to large capacity.
Strong Points	Provides high data throughput, data redundancy, and very good performance.
Weak Points	Requires 2 to 8 times as many parity drives as RAID 5.
Drives	6 to 32

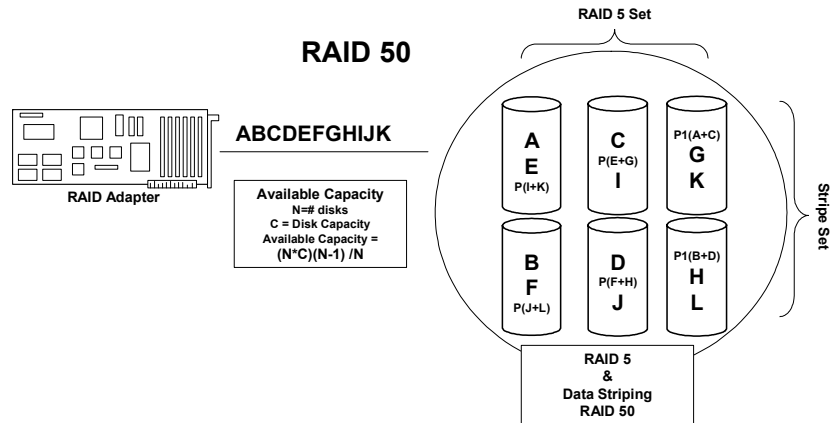


Figure 7. RAID 50 - Combination of RAID 5 and RAID 0

RAID 60 - Combination of RAID 0 and RAID 6

RAID 60 provides the features of both RAID 0 and RAID 6, and includes both parity and disk striping across multiple arrays. RAID 6 supports two independent parity blocks per stripe.

A RAID 60 virtual disk can survive the loss of two disks in each of the RAID 6 sets without losing data. RAID 60 is best implemented on two RAID 6 disk groups with data striped across both disk groups.

RAID 60 breaks up data into smaller blocks, and then stripes the blocks of data to each RAID 6 disk set. RAID 6 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks and then writes the blocks of data and parity to each drive in the array. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.

RAID 60 supports up to 8 spans and tolerates up to 16 drive failures, though less than total disk drive capacity is available. Each RAID 6 level can tolerate two drive failures.


Table 8 provides an overview of RAID 60.

Table 8. RAID 60 Overview

Uses	<p>Provides a high level of data protection through the use of a second parity block in each stripe. Use RAID 60 for data that requires a very high level of protection from loss.</p> <p>In the case of a failure of one drive or two drives in a RAID set in a virtual disk, the RAID controller uses the parity blocks to recreate all the missing information. If two drives in a RAID 6 set in a RAID 60 virtual disk fail, two drive rebuilds are required, one for each drive. These rebuilds do not occur at the same time. The controller rebuilds one failed drive, and then the other failed drive.</p> <p>Use for office automation, online customer service that requires fault tolerance or for any application that has high read request rates but low write request rates.</p>
Strong Points	<p>Provides data redundancy, high read rates, and good performance in most environments. Each RAID 6 set can survive the loss of two drives or the loss of a drive while another drive is being rebuilt. Provides the highest level of protection against drive failures of all of the RAID levels. Read performance is similar to that of RAID 50, though random reads in RAID 60 might be slightly faster because data is spread across at least one more disk in each RAID 6 set.</p>
Weak Points	<p>Not well suited to tasks requiring a lot of writes. A RAID 60 virtual disk has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Disk drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes. RAID 6 costs more because of the extra capacity required by using two parity blocks per stripe.</p>
Drives	<p>A minimum of 6.</p>

The following figure shows a RAID 6 data layout. The second set of parity drives are denoted by *Q*. The *P* drives follow the RAID 5 parity scheme.

Note: *When only three hard drives are available for RAID 6, the situation has to be that P equals Q equals original data, which means that the three hard drives have the same original data, which can afford two disk failures.*



NOTE: Parity is distributed across all drives in the array.

Figure 8. RAID 60 Level Virtual Drive

RAID Configuration Strategies

The most important factors in RAID array configuration are:

- Virtual disk availability (fault tolerance)
- Virtual disk performance
- Virtual disk capacity

You cannot configure a virtual disk that optimizes all three factors, but it is easy to choose a virtual disk configuration that maximizes one factor at the expense of another factor. For example, RAID 1 (mirroring) provides excellent fault tolerance, but requires a redundant drive. The following subsections describe how to use the RAID levels to maximize virtual disk availability (fault tolerance), virtual disk performance, and virtual disk capacity.

Maximizing Fault Tolerance

Fault tolerance is achieved through the ability to perform automatic and transparent rebuilds using hot-spare drives and hot swaps. A hot-spare drive is an unused online available drive that the RAID controller instantly plugs into the system when an active drive fails. After the hot spare is automatically moved into the RAID array, the failed drive is automatically rebuilt on the spare drive. The RAID array continues to handle requests while the rebuild occurs.

A hot swap is the manual substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is running hot-swap drives. Auto-Rebuild in the WebBIOS Configuration Utility allows a failed drive to be replaced and automatically rebuilt by “hot swapping” the drive in the same drive bay. The RAID array continues to handle requests while the rebuild occurs, providing a high degree of fault tolerance and zero downtime.

Table 9. RAID Levels and Fault Tolerance

RAID Level	Fault Tolerance
0	Does not provide fault tolerance. All data is lost if any drive fails. Disk striping writes data across multiple disk drives instead of just one disk drive. It involves partitioning each drive storage space into stripes that can vary in size. RAID 0 is ideal for applications that require high bandwidth but do not require fault tolerance.
1 or IME	Provides complete data redundancy. If one drive fails, the contents of the other drive can be used to run the system and reconstruct the failed drive. The primary advantage of disk mirroring is that it provides 100 percent data redundancy. Since the contents of the drive are completely written to a second drive, no data is lost if one of the drives fails. Both drives contain the same data at all times. RAID 1 or IME is ideal for any application that requires fault tolerance and minimal capacity.
5	Combines distributed parity with disk striping. Parity provides redundancy for one drive failure without duplicating the contents of entire disk drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In RAID 5, this method is applied to the entire drive or stripes across all disk drives in an array. Using distributed parity, RAID 5 offers fault tolerance with limited overhead.

RAID Level	Fault Tolerance
6	Combines distributed parity with disk striping. RAID 6 can sustain two drive failures and still maintain data integrity. Parity provides redundancy for two drive failures without duplicating the contents of entire disk drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In RAID 6, this method is applied to entire drives or stripes across all drives in an array. Using distributed parity, RAID 6 offers fault tolerance with limited overhead.
10	Provides complete data redundancy using striping across spanned RAID 1 arrays. RAID 10 works well for any environment that requires the 100 percent redundancy offered by mirrored arrays. RAID 10 can sustain a drive failure in each mirrored array and maintain drive integrity.
50	Provides data redundancy using distributed parity across spanned RAID 5 arrays. RAID 50 includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to recreate all missing information. RAID 50 can sustain one drive failure per RAID 5 array and still maintain data integrity.
60	Provides data redundancy using distributed parity across spanned RAID 6 arrays. RAID 60 can sustain two drive failures per RAID 6 array and still maintain data integrity. It provides the highest level of protection against drive failures of all of the RAID levels. RAID 60 includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to recreate all missing information.

Maximizing Performance

A RAID disk subsystem improves I/O performance. The RAID array appears to the host computer as a single storage unit or as multiple virtual units. I/O is faster because drives can be accessed simultaneously. [Table 10](#) describes the performance for each RAID level.

Table 10. RAID Levels and Performance

RAID Level	Performance
0	RAID 0 (striping) offers the best performance of any RAID level. RAID 0 breaks up data into smaller blocks, then writes a block to each drive in the array. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size from 8 KB to 128 KB. These stripes are interleaved in a repeated sequential manner. Disk striping enhances performance because multiple drives are accessed simultaneously.
1 or IME	With RAID 1 or IME (mirroring), each drive in the system must be duplicated, which requires more time and resources than striping. Performance is impaired during drive rebuilds.

RAID Level	Performance
5	<p>RAID 5 provides high data throughput, especially for large files. Use this RAID level for any application that requires high read request rates, but low write request rates, such as transaction processing applications, because each drive can read and write independently. Since each drive contains both data and parity, numerous writes can take place concurrently. In addition, robust caching algorithms and hardware based exclusive-or assist make RAID 5 performance exceptional in many different environments.</p> <p>Parity generation can slow the write process, making write performance significantly lower for RAID 5 than for RAID 0 or RAID 1. Disk drive performance is reduced when a drive is being rebuilt. Clustering can also reduce drive performance. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.</p>
6	<p>RAID 6 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. However, RAID 6 is not well suited to tasks requiring a lot of writes. A RAID 6 virtual disk has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Disk drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.</p>
10	<p>RAID 10 works best for data storage that need the enhanced I/O performance of RAID 0 (striped arrays), which provides high data transfer rates. Spanning increases the size of the virtual volume and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases (the maximum number of spans is eight). As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 5 array.</p>
50	<p>RAID 50 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the size of the virtual volume and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases (the maximum number of spans is eight). As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 5 array.</p>
60	<p>RAID 60 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the size of the virtual volume and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases (the maximum number of spans is eight). As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 6 array.</p> <p>RAID 60 is not well suited to tasks requiring a lot of writes. A RAID 60 virtual disk has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Disk drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.</p>

Maximizing Storage Capacity

Storage capacity is an important factor when selecting a RAID level. There are several variables to consider. Striping alone (RAID 0) requires less storage space than mirrored data (RAID 1 or IME) or distributed parity (RAID 5 or RAID 6). RAID 5, which provides redundancy for one drive failure without duplicating the contents of entire disk drives, requires less space than RAID 1. [Table 11](#) explains the effects of the RAID levels on storage capacity.

Table 11. RAID Levels and Capacity

RAID Level	Capacity
0	RAID 0 (disk striping) involves partitioning each drive storage space into stripes that can vary in size. The combined storage space is composed of stripes from each drive. RAID 0 provides maximum storage capacity for a given set of physical disks.
1 or IME	With RAID 1 (mirroring), data written to one disk drive is simultaneously written to another disk drive, which doubles the required data storage capacity. This is expensive because each drive in the system must be duplicated.
5	RAID 5 provides redundancy for one drive failure without duplicating the contents of entire disk drives. RAID 5 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks, then writes the blocks of data and parity to each drive in the array. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.
6	RAID 6 provides redundancy for two drive failures without duplicating the contents of entire disk drives. However, it requires extra capacity because it uses two parity blocks per stripe. This makes RAID 60 more expensive to implement.
10	RAID 10 requires twice as many drives as all other RAID levels except RAID 1. RAID 10 works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate to medium capacity. Disk spanning allows multiple disk drives to function like one big drive. Spanning overcomes lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources.
50	RAID 50 requires two to four times as many parity drives as RAID 5. This RAID level works best when used with data that requires medium to large capacity.
60	RAID 60 provides redundancy for two drive failures in each RAID set without duplicating the contents of entire disk drives. However, it requires extra capacity because a RAID 60 virtual disk has to generate two sets of parity data for each write operation. This makes RAID 60 more expensive to implement.

RAID Availability

RAID Availability Concept

Data availability without downtime is essential for many types of data processing and storage systems. Businesses want to avoid the financial costs and customer frustration associated with failed servers. RAID helps you maintain data availability and avoid downtime for the servers that provide that data. RAID offers several features, such as spare drives and rebuilds, that you can use to fix any physical disk problems, while keeping the servers running and data available. The following subsections describe these features.

Spare Drives

You can use spare drives to replace failed or defective drives in an array. A replacement drive must be at least as large as the drive it replaces. Spare drives include hot swaps, hot spares, and cold swaps.

A hot swap is the manual substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is running (performing its normal functions). In order for the functionality to work, the backplane and enclosure must support hot swap.

Hot-spare drives are physical drives that power up along with the RAID drives and operate in a standby state. If a physical disk used in a RAID virtual disk fails, a hot spare automatically takes its place and the data on the failed drive is rebuilt on the hot spare. Hot spares can be used for RAID levels 1, IME, 5, 6, 10, 50, and 60.

***Note:** If a rebuild to a hot spare fails for any reason, the hot-spare drive will be marked as “failed”. If the source drive fails, both the source drive and the hot-spare drive will be marked as “failed”.*

Before you replace a defective physical disk in a disk subsystem, a cold swap requires that you power down the system.

Rebuilding

If a physical disk fails in an array that is configured as a RAID 1, IME, 5, 6, 10, 50, or 60 virtual disk, you can recover the lost data by rebuilding the drive. If you have configured hot spares, the RAID controller automatically tries to use them to rebuild failed arrays. A manual rebuild is necessary if there are no hot spares available with enough capacity to rebuild the failed array. Before rebuilding the failed array, you must install a drive with enough storage into the subsystem.

Copyback

The copyback feature allows you to copy data from a source drive of a virtual drive to a destination drive that is not a part of the virtual drive. Copyback is often used to create or restore a specific physical configuration for a drive group (for example, a specific arrangement of drive group members on the device I/O buses). Copyback can be run automatically or manually.

Typically, when a drive fails or is expected to fail, the data is rebuilt on a hot spare. The failed drive is replaced with a new disk. Then the data is copied from the hot spare to the new drive, and the hot spare reverts from a rebuild drive to its original hot spare status. The copyback operation runs as a background activity, and the virtual drive is still available online to the host.

Copyback is also initiated when the first Self-Monitoring Analysis and Reporting Technology (SMART) error occurs on a drive that is part of a virtual drive. The destination drive is a hot spare that qualifies as a rebuild drive. The drive with the SMART error is marked as "failed" only after the successful completion of the copyback. This avoids putting the drive group in degraded status.

Note: *During a copyback operation, if the drive group involved in the copyback is deleted because of a virtual drive deletion, the destination drive reverts to an Unconfigured Good state or hot spare state.*

Order of Precedence

In the following scenarios, rebuild takes precedence over the copyback operation:

1. If a copyback operation is already taking place to a hot spare drive, and any virtual drive on the controller degrades, the copyback operation aborts, and a rebuild starts. The rebuild changes the virtual drive to the optimal state.
2. The rebuild operation takes precedence over the copyback operation when the conditions exist to start both operations. For example:
 - Where the hot spare is not configured (or unavailable) in the system.
 - There are two drives (both members of virtual drives), with one drive exceeding the SMART error threshold, and the other failed.
 - If you add a hot spare (assume a global hot spare) during a copyback operation, the copyback is aborted, and the rebuild operation starts on the hot spare.

Configuration Planning

Factors to consider when planning a configuration are the number of physical disks the RAID controller can support, the purpose of the array, and the availability of spare drives.

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

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

Dimmer Switch Feature

Powering and cooling drives represents a major cost for data centers. The new MegaRAID Dimmer™ Switch reduces the power consumption of the devices connected to a MegaRAID controller. This helps to share resources more efficiently and lower costs.

With Dimmer Switch, any unconfigured drive connected to a MegaRAID controller is spun down after 30 minutes of inactivity, reducing its power usage. Spun down drives are spun up automatically when you create a configuration using those drives.

Number of Physical Disks

Your configuration planning for the SAS RAID controller depends in part on the number of physical disks that you want to use in a RAID array. The number of drives in an array determines the RAID levels that can be supported. Only one RAID level can be assigned to each virtual disk.

MegaRAID Fast Path

The MegaRAID Fast Path Premium Feature is a high performance IO Accelerator technology for Solid State Drive RAID Arrays connected to a MegaRAID adapter. There are two “Levels” to this feature. Level 1 Fast Path enabled by default without PFK. The “Standard” or Level 1 Fast Path pertains to general IO Path improvements to Write Through data transfers. Additionally, controller cache tuning has resulted in improvements to configurations leveraging write back mode.

Level 2 Fast Path enabled automatically with Intel® RAID Premium Feature Key AXXRPFKSSD installed. Level 2 Fast Path is SSD-centric. This is where the Premium Feature kicks in by supporting full optimization of SSD Virtual Disk groups. With this premium feature enabled, Solid State Drive configurations tuned for small, random block-size IO activity -- typical of transactional database applications -- can sustain higher numbers of IO READs Per Second, compared with Level 1 Fast Path. The performance levels reached with this solution are equivalent to those of much costlier Flash-based adapter card solutions.

Refer to *Intel® RAID Premium Feature Key AXXRPFKSSD, AXXRPFKDE, and AXXRPFKSNSh Installation Guide (E88588-00x)* for description of the Premium Feature Key (PFK)

4K Sector Drive Support

The disk drive industry is in transition to support disk drives with native formatting of 4K sectors. 4K formatted drives provide an opportunity to improve capacity and error correction efficiencies as capacities grow, and as improvements in error correction are implemented. Drives supporting 4K sectors will also support a 512 byte emulation mode which will allow the drive to be operated with legacy OS and hardware products that do not support 4K byte sectors. Intel® plans to implement 4K sector support on all new products, beginning with those designed to utilize the LSI 2208 and LSI2308 SAS products. Currently shipping Intel® RAID and SAS products will support 4K sector drives running in legacy 512b sector mode.

Larger than 2TB Drive Support

The disk drive industry is in transition to support disk drives with volume size larger than 2 Terabytes. The Intel® 6G SAS PCIe Gen2 RAID Controllers can fully recognize and configure the volume of these disk drives. For the other RAID types listed in “Supported Hardware” section, there is no problem for the first 2TB space of the disk drive to be recognized and configured.

Power Save settings

The controller conserves energy by placing certain unused drives into powersave mode. The controller will automatically spin up drives from powersave mode whenever necessary. Drives that can be set to power save mode are: Unconfigured Drives, Hot Spare Drives, Configured Drives. The setting can be made in either RAID BIOS Console 2, or RAID Web Console 2. There is also a way to disable this setting from CmdTool2, by typing:

```
CmdTool2 -LDSetPowerPolicy -None -Lall -a0
```

This command only has effect when one or more logical drives are created. The Power Save mode currently is only supported by Intel® Intelligent RAID products.

Array Purpose

Important factors to consider when creating RAID arrays include availability, performance, and capacity. Define the major purpose of the disk array by answering questions related to these factors, such as the following, which are followed by suggested RAID levels for each situation:

- Will this disk array increase the system storage capacity for general-purpose file and print servers? Use RAID 5, 6, 10, 50, or 60.
- Does this disk array support any software system that must be available 24 hours per day? Use RAID 1, IME, 5, 6, 10, 50, or 60.
- Will the information stored in this disk array contain large audio or video files that must be available on demand? Use RAID 0.
- Will this disk array contain data from an imaging system? Use RAID 0 or 10.

Fill out [Table 12](#) to help you plan the array configuration. Rank the requirements for your array, such as storage space and data redundancy, in order of importance, and then review the suggested RAID levels.

Table 12. Factors to Consider for Array Configuration

Requirement	Suggested RAID Level(s)
Storage space	RAID 0, RAID 5
Data redundancy	RAID 5, RAID 6, RAID 10, RAID 50, RAID 60
Physical disk performance and throughput	RAID 0, RAID 10
Hot spares (extra physical disks required)	RAID 1, RAID IME, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60

3 RAID Utilities

Intel® Embedded Server RAID Technology 2 BIOS Configuration Utility

With support for up to six SATA drives or eight SAS/SATA drives, depending on the server board or system, the embedded RAID BIOS has the following features:

- Support for interrupt 13 and Int19h.
- Support for SATA CD-ROM/DVD-ROM devices, including support for booting from a CD-ROM drive.
- POST (Power On Self Test) and run-time BIOS support for device insertion and removal.
- Support for a migration path from Intel® Embedded Server RAID Technology 2 to Intel Integrated Server RAID hardware.
- Automatic resume of rebuilding, check consistency, and initialization.
- Global hot spare support based on the virtual drive size.
- Support for RAID levels 0, 1, 5, and 10.
- Support for auto rebuild.
- Support for different capacity disks in the same array.
- Support for up to eight physical drives and eight virtual drives.
- Stripe size of 64 KB only.
- Support for disk coercion with options of None, 128 MB, or 1 GB.
- Ability to select a virtual drive as boot device. By default, virtual drive 0 is bootable.

LSI MPT* SAS BIOS Configuration Utility

You can use the LSI MPT* SAS BIOS Configuration Utility to create one or two IM/IME volumes on each Intel® IT/IR RAID Controller, with one or two optional global hot-spare disks. All disks in an IM/IME volume must be connected to the same Intel® IT/IR RAID Controller.

Although you can use disks of different size in IM and IME volumes, the smallest disk in the volume will determine the logical size of all disks in the volume. In other words, the excess space of the larger member disk(s) will not be used. For example, if you create an IME volume with two 100 GB disks and two 120 GB disks, only 100 GB of the larger disks will be used for the volume.

Integrated Mirroring and Integrated Mirroring Enhanced support the following features:

- Configurations of one or two IM or IME volumes on the same Intel® IT/IR RAID Controller. IM volumes have two mirrored disks; IME volumes have three to ten mirrored disks. Two volumes can have up to a total of 12 disks.

Note: *This feature requires IR RAID firmware v1.20.00 or above to be installed.*

- One or two global hot-spare disks per controller, to automatically replace failed disks in IM/IME volumes. The hot-spare drives are in addition to the 12-disk maximum for two volumes per Intel® IT/IR RAID Controller.

Note: *Support for two hot-spare disks requires IR RAID firmware v1.20.00 or above.*

- Mirrored volumes run in optimal mode or in degraded mode (if one mirrored disk fails).
- Hot-swap capability.
- Presents a single virtual drive to the OS for each IM/IME volume.
- Supports both SAS and SATA disks. The two types of disks cannot be combined in the same volume. However, a Intel® IT/IR RAID Controller can support one volume with SATA disks and a second volume with SAS disks.
- Fusion-MPT* architecture.
- Easy-to-use BIOS-based configuration utility.
- Error notification: The drivers update an OS-specific event log.
- LED status support.
- Write journaling, which allows automatic synchronization of potentially inconsistent data after unexpected power-down situations.
- Metadata used to store volume configuration on mirrored disks.
- Automatic background resynchronization while the host continues to process inputs/outputs (I/Os).
- Background media verification ensures that data on IM/IME volumes is always accessible.

Intel® RAID BIOS Console 2 Configuration Utility for Intelligent RAID

The Intel® RAID BIOS Console 2 configuration utility provides full-featured, GUI-based configuration and management of RAID arrays. The Intel® RAID BIOS Console 2 utility resides in the controller firmware and is independent of the operating system. The Intel® RAID BIOS Console 2 configuration utility lets you:

- Select an Intel® RAID controller
- Choose a configuration method for physical arrays, disk groups, and virtual drives
- Create drive arrays
- Define virtual drives
- Initialize virtual drives
- Access controllers, virtual drives, and physical arrays to display their properties

- Create hot-spare drives
- Rebuild failed drives
- Verify data redundancy in RAID 1, 5, 6, 10, 50, or 60 virtual drives

Intel® RAID Web Console 2 Configuration and Monitoring Utility

The Intel® RAID Web Console 2 is an operating system-based, object-oriented GUI utility that configures and monitors RAID systems locally or over a network. The Intel® RAID Web Console 2 runs on each of the supported Microsoft Windows* and Linux operating systems.

With the Intel® RAID Web Console 2, you can perform the same tasks as you can with the Intel® RAID BIOS Console 2 or with the Intel® Embedded Server RAID BIOS Configuration utility. In addition, the Intel® RAID Web Console 2 provides on-the-fly RAID migration, creating almost limitless adaptability and expansion of any virtual drive while the system remains operational.

The Intel® RAID Web Console 2 allows you to:

- Create and manage virtual drives
- Add a drive to a RAID virtual drive

The following situations assume RAID 0 has one drive and RAID 1 has two drives:

- convert from RAID 0 to RAID 1 by adding one additional drive.
- convert from RAID 0 to RAID 5 by adding two additional drives.
- convert from RAID 0 to RAID 6 by adding three additional drives.
- convert from RAID 1 to RAID 0.
- convert from RAID 1 to RAID 5 by adding one additional drive.
- convert from RAID 1 to RAID 6 by adding two additional drives.
- convert from RAID 5 to RAID 0.
- convert from RAID 5 to RAID 6 by adding one additional drive.
- convert from RAID 6 to RAID 0.
- convert from RAID 6 to RAID 5.
- convert a degraded RAID into RAID 0.
- remove physical drives from a virtual drive.

Note: *While you can apply RAID-level migration at any time, Intel® recommends that you do so when there are no reboots. Many operating systems issues I/O operations serially (one at a time) during boot. With a RAID-level migration running, a boot can often take more than 15 minutes.*

Drive Hierarchy within the RAID Firmware

The Intel[®] Integrated RAID firmware is based on three fundamental levels. Virtual drives are created from drive arrays that are created from physical drives.

- **Level 1** consists of the physical drives (hard drives and removable hard disks). The firmware identifies each drive by its physical ID and maps it to a virtual address. A virtual drive can be constructed of more than one physical drive.
- **Level 2** consists of the array(s) formed by firmware made of one or more disks and can be made into RAID 0, 1, 5, 6, 10, 50, or 60.
- **Level 3** consists of the virtual drives. These are the only drives that can be accessed by the operating system. These are the drives given drive letters (C, D, and so forth) under the Microsoft Windows* operating system. The firmware automatically transforms each newly installed drive array into a virtual drive. RAID 0, 1, 5, or 6 use a single array and RAID 10, 50, 60 use multiple arrays.

Intel[®] Intelligent RAID Controller Features

Enterprise Features

- Online capacity expansion (OCE). Add capacity to the virtual drive. The added capacity can be presented to the operating system as additional space for the operating system to partition it as an additional drive, or it may be added to an operating system drive, depending upon the capability of the operating system.
- Online RAID level migration allows for upgrading a RAID level. Options are to go from RAID 1 to RAID 0, RAID 5 to RAID 0, RAID 6 to RAID 0, RAID 6 to RAID 5. With OCE, options are to go from RAID 0 to RAID 1, RAID 0 to RAID 5, RAID 0 to RAID 6, RAID 1 to RAID 5, RAID 1 to RAID 6, RAID 5 to RAID 6.
 - You cannot migrate or perform OCE on a spanned RAID array or disk group (RAID 10, RAID 50, or RAID 60).
 - You cannot migrate to a smaller capacity configuration.
 - You cannot perform OCE when there is more than one virtual drive on a virtual array or disk group.
- Each controller allows 128 virtual drives.
- When five or more disks are used, Smart Initialization automatically checks consistency of virtual drives for RAID 5. This allows performance optimization by enabling read-modify-write mode of operation with five or more disks in a RAID 5 array or disk group. Peer read mode of operation is used when the RAID 5 array or disk group contains three or four physical drives.
- If the system shuts down, the initialization or rebuild process automatically resumes on the next boot. Auto resume must be enabled prior to virtual drive creation.
- Stripe size is user definable on a per drive basis and can be 8, 16, 32, 64, or 128 KB in size. The default is 256 KB, which is optimal for many data access types.

- Hot spares can be set as global or dedicated. A global hot spare automatically comes online to replace the first drive to fail on any array or disk group on the controller. A dedicated hot spare is assigned to a specific array or disk group and only comes online to rebuild this specific failed array or disk group. A hot spare only comes online if it is the same size or larger than the failing drive (see drive coercion below), and if a drive has been marked as failed. If a drive is removed (and marked as failed) within a virtual drive, the hot spare automatically comes online. However, there must be disk activity (I/O to the drive) in order for a missing drive to be marked as failed.
- Drive coercion refers to the ability of the controller to recognize the size of the physical drives connected and then force the larger drives to use only the amount of space available on the smallest drive. Drive coercion allows an option to map out a reserved space to compensate for slightly smaller drive sizes that may be added later. The default is set to 1 GB. The coercion algorithm options are:
 - None: No coercion of size.
 - 128 MB: The software rounds the drive capacity down to the next 128 MB boundary and then up to the nearest 10 MB until the coerced capacity is larger than the actual drive size. It is then reduced by 10 MB.
 - 1 GB: The software rounds the drive capacity down to the nearest 1 GB boundary and then down by 1 MB. This corresponds to the terms most drive manufacturers use.

Fault Tolerant Features

- Configuration on Disk (COD) and NVRAM (Non-volatile Random Access Memory) storage of array and disk group configuration information. Array and disk group configuration information is stored both on the hard drive (COD) and in NVRAM. This helps protect against loss of the configuration due to adapter and/or drive failure.
- Failed drives are automatically detected and a transparent rebuild of the failed array automatically occurs using a hot-spare drive.
- Support for SAF-TE (SCSI Accessed Fault-Tolerant Enclosure) enabled enclosures allows enhanced drive failure and rebuild reporting via enclosure LEDs (Light-Emitting Diodes); support also includes hot swapping of hard drives.
- A battery backup for cache memory is available as an option. RAID controller firmware automatically checks for the presence of the battery module, and if found, allows the write back cache option. The adapter continuously tracks the battery voltage and reports if the battery is low. If low, the battery is first given a fast charge to replenish the charge and is then given a trickle charge to keep it at an optimal power level. Adapters that support the battery module include a “dirty cache” LED; when power is lost to the system and data remains in the cache memory that has not been written to disk, the LED signals that this operation needs to be completed. Upon reboot, the data in memory can then write to the hard disk drive.
- Although I/O performance may be lower, hard disk drive write-back cache is disabled by default because data can potentially be lost if a power outage occurs. Enabling the HDD write-back cache may improve performance, but when enabled, you should use a UPS (Uninterruptible Power Supply) device to prevent data loss during power outages.
- Battery life is about three years. You should monitor the battery health and replace when needed.

- SMART (Self-Monitoring Analysis and Reporting Technology) technology is supported. This provides a higher level of predictive failure analysis of the hard disk drives by the RAID controller.

Cache Options and Settings

Cache options and settings can be unique for each virtual drive.

- Cache Write Policy
 - Write Through: I/O completion is signaled only after the data is written to hard disk.
 - Write Back with BBU: I/O completion is signaled when data is transferred to cache.
 - Always Write Back- Write back is enabled even if BBU is bad or missing
- Cache Policy
 - Direct I/O: When possible, no cache is involved for both reads and writes. The data transfers are directly from host system to the disk and from the disk to the host system.
 - Cached I/O: All reads first look at cache. If a cache hit occurs, the data is read from cache; if not, the data is read from disk and the read data is buffered into cache. All writes to drive are also written to cache.
- Read Policy
 - No Read Ahead: Provides no read ahead for the virtual drive.
 - Read Ahead: Reads and buffers additional consecutive stripes/lines into cache.
 - Adaptive: The read ahead automatically turns on and off depending upon whether the disk is accessed for sequential reads or random reads.

Background Tasks

- Rebuilding a failed drive is performed in the background. The rebuild rate is tunable from 0-100%.
 - The rebuild rate controls the amount of system resources allocated to the rebuild.

Caution: *It is not recommended to increase the rebuild rate to over 50%. A higher rebuild rate can result in operating system requests not being serviced in a timely fashion and causing an operating system error.*
- A consistency check scans the consistency of data on a fault-tolerant disk to determine if data has been corrupted.
- Background initialization is a background check of consistency. It has the same functionality as the check consistency option but is automatic and can be canceled only temporarily. If it is canceled, it starts again in a few minutes. Background initialization is only performed on redundant volumes.
- RAID level migration and online capacity expansion are completed in the background.
- Patrol Read is a user definable option available in the Intel® RAID Web Console 2 that performs drive reads in the background and maps out any bad areas of the drive.

Error Handling

- Most commands are retried four or more times. The firmware is programmed to provide the best effort to recognize an error and recover from it if possible.
- Failures are logged and stored in NVRAM. Operating system-based errors are viewable from the event viewer in the Web Console 2.
- RAID-related errors can be reported by the hard drive firmware, SAF-TE controller, or the RAID controller firmware. These errors may be reported to the operating system through RAID management software, through SMART monitoring, or through CIM management. Some errors may also be reported by the SAF-TE controller and logged in the system event log (SEL) for the Intel[®] server board. In addition, the operating system may report access errors. Depending on the RAID controller and drive enclosure, the error may be evident by the color of LEDs, the flashing of LEDs, or audible alarms.

Audible Alarm

The following list of beep tones is used on Intel[®] Intelligent RAID Controllers. These beeps usually indicate that a drive has failed.

- Degraded Array or Disk Group: Short tone, 1 second on, 1 second off
- Failed Array or Disk Group: Long tone, 3 seconds on, 1 second off
- Hot Spare Commissioned - Short tone, 1 second on, 3 seconds off

During a rebuild, the tone alarm stays on. After the rebuild completes, an alarm with a different tone will sound.

The disable alarm option in either the Intel[®] RAID BIOS Console 2 or Intel[®] Web Console 2 management utilities holds the alarm disabled after a power cycle. You must use the enable alarm option to re-enable the alarm.

The silence alarm option in either the Intel[®] RAID BIOS Console 2 or the Intel[®] Web Console 2 management utilities will silence the alarm until a power cycle or another event occurs.

4 Intel® RAID Drivers

The drivers that Intel provides for Intel® RAID Controllers are not compatible with SCSI or SATA-only RAID controllers. The RAID driver files are available on the *Resource CD* that accompanies the RAID controllers. The driver files are also available at <http://downloadcenter.intel.com>. If you need to transfer the driver files to another system, you can copy them to a floppy disk or a USB key.

Note: *Intel updates software frequently and updated drivers may provide additional features. Check for new software at the Intel Web site: <http://www.intel.com/support/motherboards/server/>. See the Readme file that accompanies the download for updated information. For operating systems that are not listed here, but are listed at the above Intel Web site see Readme file that accompanies the download for installation steps.*

RAID Driver Installation for Microsoft Windows*

Installation in a New Microsoft Windows* Operating System

This procedure installs the RAID device driver system during the Microsoft Windows 2003*, Microsoft Windows 2000*, or Microsoft Windows XP* operating system installation. The system must contain an Intel® RAID controller. Microsoft Windows 2003* automatically adds the driver to the registry and copies the driver to the appropriate directory.

1. Start the Microsoft Windows* installation by booting from the Microsoft Windows* CD-ROM disk. The system BIOS must support booting from a CD-ROM drive. You may need to change BIOS settings to allow CD-ROM booting. See your system documentation for instructions.
2. Press <F6> to install when the screen displays:
`Press F6 if you need to install...`

Note: *You must press <F6> for the system to recognize the new driver.*

3. Choose <S> to specify an additional device when the screen displays:
`Setup could not determine the type...`

Note: *If this screen is not displayed as the first user input, then the setup program did not register that the <F6> key was pressed. Reboot the system and return to step 2.*

4. When the system asks for the manufacturer-supplied hardware support disk, insert the Microsoft Windows* driver disk and press <Enter>.
5. Select the appropriate Microsoft Windows* driver from the menu by highlighting it. Press <Enter> to proceed. The driver is added to the registry and copied to the appropriate directory.
6. Continue with the Microsoft Windows* operating system installation procedure.

Installation in an Existing Microsoft Windows* Operating System

This procedure installs or upgrades the RAID device driver on an existing Microsoft Windows 2003*, Microsoft Windows 2000*, or Microsoft Windows XP* operating system. The system must contain an Intel® RAID controller.

1. Boot to the Microsoft Windows* operating system. The Found New Hardware Wizard is displayed. The program identifies the SAS controller and requests the driver disk.
2. Insert the Microsoft Windows* driver disk into the floppy drive.
3. For Microsoft Windows 2003* or Microsoft Windows XP*, choose Install Software Automatically. In Microsoft Windows 2000*, choose Search for a Suitable Driver.
4. Microsoft Windows 2000* only: Click the Specify location box and make sure the search location is the floppy drive.
5. Click Next.
6. A message that this driver is not digitally signed may display. This message informs you that a nonsigned driver is being installed. If you see this message, click Continue Anyway.
7. The system loads the driver from the Microsoft Windows* driver disk and copies the driver to the system disk. The Found New Hardware Wizard screen displays the message:
`The wizard has finished...`
8. Click Finish to complete the driver upgrade.

RAID Driver Installation for Red Hat* Enterprise Linux

This section describes the installation of the device driver on new Red Hat* Enterprise Linux 3, 4, or 5 systems. The following are general installation guidelines. Refer to the release notes that accompanied the driver for information on updating the driver on an existing Red Hat* Linux system.

1. Boot to the CD-ROM with Disk 1. Command: *linux dd*
2. Press <Enter> at the boot prompt on the Welcome screen.
3. Copy the Linux driver image from the *Resource CD* to a disk or USB key.
4. Insert the disk with driver image.
5. Select Yes.
6. Scroll down to select Intel® RAID adapter driver. The utility locates and loads the driver for your device.
7. Follow the Red Hat* Linux installation procedure to complete the installation.

RAID Driver Installation for SuSE* Linux

SuSE* Linux uses a program called YaST2 (Yet another System Tool) to configure the operating system during installation. For complex installations, you can select “Install Manually” at the first install screen and a different program, `linuxrc`, is used. This section assumes a straightforward installation using YaST2.

1. Insert CD-ROM disk 1 into the CD-ROM drive and the RAID controller driver diskette in the floppy drive.
1. Boot to the CD-ROM.
2. The operating system loads a minimal operating system from the CD-ROM onto a RAM disk. The operating system also loads any driver module found in the floppy drive.
3. At the Welcome to YaST2 screen, select your language and click Accept.
4. At the Installation Settings screen, set up the disk partitioning.
5. Continue with the SuSE* Linux installation procedure.

RAID Driver Installation for Novell NetWare*

Installation in a New Novell Netware* System

Follow the instructions in the Novell Netware* Installation Guide to install Novell Netware in the server. Perform the following steps to install Novell NetWare using your Intel® RAID controller as a primary adapter.

Note: *Drivers for Novell Netware* are not available on the CD-ROM. The latest drivers are available at <http://www.intel.com/support/motherboards/server/> or from your CDI account.*

1. Boot from Novell NetWare*.
2. Follow the instructions on the screen until you reach the Device Driver screen, which is used to modify drivers.
3. Select Modify and press <Enter>.
4. On the Storage Driver Support screen select Storage Adapters and press <Enter>.
5. Delete any existing Intel® RAID adapter listings.
6. Press <Insert> to add unlisted drivers.
7. Press <Insert> again.
A path is displayed.
8. Press <F3>.
9. Insert the driver disk into the floppy drive, and press <Enter>.
The system will locate the .HAM driver.
10. Press the <Tab> key.
11. Select the Driver Summary screen, and press <Enter>.

12. Continue the Novell NetWare installation procedure.

Installation in an Existing Novell Netware* System

Perform the following steps to add the Novell NetWare* driver to an existing Installation.

Note: *Drivers for Novell Netware* are not available on the CD-ROM. The latest drivers are available at <http://www.intel.com/support/motherboards/server/> or from your CDI account.*

1. Type `nwconfig` at the root prompt and press <Enter>. The Configuration Options screen loads.
2. Select Drive Options and press <Enter>.
3. Select Configure Disk and Storage Device Options, and press <Enter>.
4. Select one of the following options displayed in the window:
 - a. Discover and Load an Additional Driver - If you select this option, the system discovers the extra unit and prompts you to select a driver from the list. Press <Insert> to insert the driver. This completes the procedure.
 - b. Select an Additional Driver - If you select this option the Select a Driver screen displays. Press <Insert>. Follow the instructions that display. Insert a disk into the floppy drive, and press <Enter>. The system will find and install the driver. This completes the procedure.

RAID Driver Installation for Solaris* 10

Installation in a New Solaris* System

This updated driver can be applied using the normal operating system installation options.

Note: *Drivers for Solaris* 10 are not available on the CD-ROM. The latest drivers are available at <http://www.intel.com/support/motherboards/server/> or from your CDI account.*

Boot the target system from the Solaris* 10 OS DVD (starting with DVD #1).

1. Select Solaris from the GRUB menu.
2. After the initial kernel loads, select option 5 Apply driver updated.
3. Insert driver floppy or CD into USB floppy or DVD-ROM drive, respectively, on the target system.

Installation in an Existing Solaris* System

1. Create a temporary directory “/tmp” under current working directory. Command: `mkdir tmp`.
2. Depending on your platform, `untar i386.tar` or `x86_64.tar`. Command: `tar -xf i386.tar` or `tar -xf x86_64.tar`.
3. Depending on your platform, run `install.sh` or `install32.sh`. Command: `sh install` or `sh install32`.

5 Intel® Embedded Server RAID BIOS Configuration Utility

If the SATA RAID or SAS RAID options are enabled in the server BIOS, an option to enter the Intel® Embedded Server RAID BIOS Configuration utility displays during the server boot process. To enter the utility, press <Ctrl> + <E> when prompted.

The Intel® Embedded Server RAID BIOS Configuration utility allows a user to:

- Create, add, modify, and clear virtual drive configurations
- Initialize or rebuild the configured drives
- Set the boot drive
- Create a global hot-spare drive
- View physical and virtual drive parameters
- View and set adapter properties, including consistency check and auto-resume

SATA and SAS systems use different versions of the Intel® Embedded Server RAID BIOS Configuration utility, but both versions use the same keystrokes and contain identical menus. The utility menus show limited help at the bottom of the screen and selections are chosen with the arrow keys and the space bar. If no virtual drive is available to configure, a warning is displayed. Only the number of potential physical drives differs for the SAS and SATA versions of the utility.

The following menu and sub-menu options are available:




Figure 9. Intel® Embedded Server RAID BIOS Configuration Utility Screen

Creating, Adding or Modifying a Virtual Drive Configuration

To create, add, or modify a virtual drive configuration, follow these steps:

1. Boot the system.
 2. Press <Ctrl> + <E> when prompted to start the Intel® Embedded Server RAID BIOS Configuration utility.
 3. Select **Configure** from the Main Menu.
 4. Select a configuration method:
 - **Easy Configuration** does not change existing configurations but allows new configurations.
 - **New Configuration** deletes any existing arrays and virtual drives and creates only new configurations.
 - **View/Add Configuration** lets you view or modify an existing configuration.
- For each configuration method, a list of available physical drives is displayed. These drives are in the READY state. If you select a physical drive in the list, information about each drive is displayed.
5. Use the arrow keys to move to a drive and press the space bar to add it to the array.

Note: The utility limits each drive to the size of the smallest drive.

The status for each selected drive that is added to an array changes status from READY to ONLIN A[array#]-[drive#]. For example, ONLIN A00-01 means array 0, disk drive 1.

6. (Optional) Create a global hot-spare drive by highlighting a drive that is marked READY and press the <F4> key. Then select **Yes** from the pop-up menu.
7. Repeat step 5 and step 6 to create a second array if needed. When you have selected drives for all desired arrays, press the <F10> key.
8. Select an array by highlighting it. Press the <Enter> key to set the properties.
9. The virtual drive configuration screen is displayed, This screen shows the following:
 - Virtual drive number
 - RAID level
 - Virtual drive size
 - Number of stripes in the physical array
 - Stripe size
 - State of the virtual drive

To set these options, highlight a property and press the <Enter> key. The available parameters for that property are displayed for the selection.

10. Select a RAID level: Select 0, 1, or 10 depending upon number of drives and the purpose.
11. Consider whether you need to override the default virtual drive size. By default, all available space in the array is assigned to the current virtual drive. For RAID 10 arrays, only one virtual drive can be defined for the entire array.
12. (Optional) Change the default Write Cache and Read Ahead policies. See Setting the Write Cache and Read Ahead Policies.
13. When you have finished defining the current virtual drive, select **Accept** and press the <Enter> key.
14. Repeat step 8 through step 13 for all virtual drives.
15. Save the configuration when prompted, and press any key to return to the Main Menu.
16. Select **Initialize** and use the space bar to highlight the virtual drive to initialize.

Caution: All data on the virtual drive is erased during an initialization.

17. Press the <F10> key. Select **Yes** at the prompt and press the <Enter> key to begin the initialization. A graph shows the progress of the initialization.
18. After the initialization is complete, press the <Esc> key to return to the previous menu. Pressing the <Esc> key closes the current menu. If a process is running when you press the <Esc> key, you are given the following options:
 - **Abort:** When Abort is selected, the task is stopped and will not resume. If an initialization has started, Abort does not restore data.
 - **Stop:** When Stop is selected, the current task stops. Stop is available only if auto resume is enabled on the adapter. See AutoResume/AutoRestore for information.

- **Continue:** The task continues normally. Continue cancels the press of the <Esc> key. If AutoResume is enabled, the task resumes from the point at which it was stopped.

Setting the Write Cache and Read Ahead Policies

Read and write cache settings apply to all virtual drives in an array. They may show as on/off; enable/disable; or as initials of the desired state, such as WB for Write Back. They are in menus as Write Policy and Read Policy or as Write Cache (WC) and Read Ahead (RA). You can view these policies from the Adapter Properties or from the Virtual Drive's View/Update Parameters.

The following are the cache policies:

- If WC is on, the write cache of the physical drives, which makes the virtual drive turn on. In this mode, when the physical drive cache receives all the data, the I/O request is signaled as completed.

***Caution:** If power fails before the cached data is written to the drive, the data is lost.*

- If WC is off, only if the data is written to the media of the drive, the I/O request is signaled as completed.
- RA = ON allows the read ahead mode of the physical drives, which makes the virtual drive turn on. In this mode, the physical drive will read additional data and store that data into its cache. This improves performance on sequential reads.

To change cache policies, follow these steps:

1. Select **Objects | Virtual Drive | Virtual Drive n | View/Update Parameters**.
2. Use the arrow key to select the option to change. Press the <Enter> key.
3. Use the arrow key to select **Off** or **On**.
4. If asked to confirm the change, use the arrow key to select **Yes**. Press the <Enter> key to change the cache setting.

Working with a Global Hot-spare Drive

A global, but not dedicated, hot-spare drive can be created to automatically replace a failed drive in a RAID 1 or RAID 10 array. For new arrays, you should create the global hot-spare during the configuration process. See [“Creating, Adding or Modifying a Virtual Drive Configuration”](#) on page 40.

Adding a Hot-spare Drive

To add a hot-spare drive to an existing configuration, follow these steps:

1. Select **Objects** from the Main Menu.
2. Select **Physical Drive**. A list of physical drives is displayed.

3. Select an unused drive from the list, and select **Make Hot Spare**. The screen changes to indicate HOTSP.

Removing a Hot-spare Drive

To remove a hot-spare drive, follow these steps:

1. Select **Objects** from the Main Menu.
2. Select **Physical Drive**. A list of physical drives is displayed.
3. Select the disk that displays **HOTSP**, press the <Enter> key.
4. Select **Force Offline** and press the <Enter> key. The status of the drive changes to **READY**. The drive can be used in another array.

Rebuilding a Drive

The Intel® Embedded Server RAID BIOS Configuration utility includes a manual rebuild option that rebuilds a failed array due to an individual failed drive in a RAID 1 or 10 array. RAID 0 drives are not redundant and cannot be rebuilt. You can also rebuild a good drive (not physically failed) using the existing configuration data.

To rebuild a drive:

1. Select **Rebuild** from the Main Menu. The failed drives show the status **FAIL**.
2. Press the arrow keys to highlight the physical drive that you want to rebuild. Press the space bar to select the drive.
3. Press the <F10> key and select **Y** to confirm. As the rebuild process begins, the drive indicator shows **REBLD**.
4. When the rebuild is complete, press any key to continue.

Auto Rebuild and Auto Resume

To ensure data protection, enable Auto Rebuild and Auto Resume so that drives are automatically re-created to maintain redundancy.

- In a pre-boot environment, auto rebuild starts only when you enter the BIOS utility.

***Note:** Hot-plug support is not available in the pre-boot environment. For the system BIOS or the Intel® Embedded Server RAID BIOS Configuration utility to detect the physical drive, insert the drive when the system is off.*

- When the operating system is running, the auto rebuild starts if the system has a hot-spare drive or if you replace the failed drive with a new drive.

The Auto Rebuild and Auto Resume options are available in the Intel® Embedded Server RAID BIOS Configuration utility from the menu that is displayed after you select **Objects | Adapter**.

Checking Data Consistency

The Check Consistency feature can be used on RAID 1 or RAID 10 drives to verify the data consistency between the mirrored drives. It can be set to only report or to both report and automatically fix the data.

1. From the Main Menu, select **Check Consistency** and press the <Enter> key.
A list of configured virtual drives is displayed.
2. Use the arrow keys to choose the desired drive. Press the space bar to select the virtual drive to check for consistency. (RAID 1 or 10 only)
3. Press the <F10> key.
4. At the prompt, select **Yes** and then press the <Enter> key.

If the Report and Fix/Report options are not shown, select **Main Menu | Objects | Adapter | ChkCons** and set **Report only** or **Fix\Report**.

Viewing and Changing Device Properties

You can view adapter, virtual drive, and physical drive properties. You can also change some adapter properties and the Write Cache and Read Ahead for Virtual Drives.

1. From the Main Menu select **Objects**.
2. Choose **Adapter**, **Virtual Drive**, or **Physical Drive**.
3. Select the device from the list and view the properties.
 - For virtual drives, choose **View | Update Parameters**.
 - For physical drives, choose **Drive Properties**.

The numeric values of the rates settings are the percentage of system resources. FGI and BGI are abbreviations for foreground and background initialization rates.

4. To change a value, highlight the property and press the <Enter> key.

Note: Some values cannot be changed.

5. Select or type a different value for the property and press the <Enter> key.
6. When you are finished, press the <Esc> key until you return to the Main Menu.

Forcing Drives Online or Offline

A drive can be forced offline so that a hot-spare drive will replace it. Power failures may cause a drive to go offline and you must force it back online.

Forcing a Drive Online or Offline

You can force a drive offline so that a hot-spare replaces it. Power failures may cause a drive to go offline and you must force it back online. To force a drive online or offline, follow these steps:

1. On the Main Menu, select **Objects** and then **Physical Drive**.
2. Highlight a physical drive that is a member of an array and press the <Enter> key.
3. From the menu, choose one of the following:
 - **Force Offline** to take the drive off line. If the drive was online, its status changes to FAIL.
 - **Force Online** to bring the drive on line. If the drive was offline, its status changes to ONLINE.

Configuring a Bootable Virtual Drive

Follow these steps to configure a bootable virtual drive:

1. From the Main Menu, select **Configure | Select Boot Drive**.
2. Select a virtual drive from the list to make it the designated boot drive.

Note: You should also check the system BIOS Setup utility for the boot order setting. To access the BIOS Setup utility, press the <F2> key when prompted during POST.

Deleting (Clearing) a Storage Configuration

Caution: Before you clear a storage configuration, back up all the data you want to keep.

To clear a storage configuration, follow these steps:

1. On the Main Menu, select **Configure | Clear Configuration**.
2. When the message appears, select **Yes** to confirm. All virtual drives are deleted from the configuration.

6 Intel® IT/IR RAID Configuration

This chapter explains how to create Integrated Mirroring (IM), Integrated Mirroring Enhanced (IME), and Integrated Striping (IS) volumes using the LSI MPT* SAS BIOS Configuration Utility.

IM and IME Configuration Overview

You can use the LSI MPT* SAS BIOS Configuration Utility to create one or two IM/IME volumes on each Intel® IT/IR RAID Controller, with one or two optional global hot-spare disks. All disks in an IM/IME volume must be connected to the same Intel® IT/IR RAID Controller.

Although you can use disks of different size in IM and IME volumes, the smallest disk in the volume will determine the logical size of all disks in the volume. In other words, the excess space of the larger member disk(s) will not be used. For example, if you create an IME volume with two 100 GB disks and two 120 GB disks, only 100 GB of the larger disks will be used for the volume.

Features

Integrated Mirroring and Integrated Mirroring Enhanced support the following features:

- Configurations of one or two IM or IME volumes on the same Intel® IT/IR RAID Controller. IM volumes have two mirrored disks; IME volumes have three to ten mirrored disks. Two volumes can have up to a total of 12 disks.

Note: *This feature requires IR RAID firmware v1.20.00 or above to be installed.*

- One or two global hot-spare disks per controller, to automatically replace failed disks in IM/IME volumes. The hot-spare drives are in addition to the 12-disk maximum for two volumes per Intel® IT/IR RAID Controller.

Note: *Support for two hot-spare disks requires IR RAID firmware v1.20.00 or above.*

- Mirrored volumes run in optimal mode or in degraded mode (if one mirrored disk fails).
- Hot-swap capability.
- Presents a single virtual drive to the OS for each IM/IME volume.
- Supports both SAS and SATA disks. The two types of disks cannot be combined in the same volume. However, a Intel® IT/IR RAID Controller can support one volume with SATA disks and a second volume with SAS disks.
- Fusion-MPT* architecture.

- Easy-to-use BIOS-based configuration utility.
- Error notification: The drivers update an OS-specific event log.
- LED status support.
- Write journaling, which allows automatic synchronization of potentially inconsistent data after unexpected power-down situations.
- Metadata used to store volume configuration on mirrored disks.
- Automatic background resynchronization while the host continues to process inputs/outputs (I/Os).

Background media verification ensures that data on IM/IME volumes is always accessible.

Creating IM and IME Volumes

The LSI MPT* SAS BIOS Configuration Utility is part of the Fusion-MPT* BIOS. When the BIOS loads during boot and you see the message about the LSI MPT* SAS BIOS Configuration Utility, press <Ctrl> + <C> to start the utility. The message then changes to:

```
Please wait, invoking LSI SAS Configuration Utility...
```

After a brief pause, the main menu appears. On some systems, however, the following message appears:

```
Configuration Utility will load following initialization!
```

In this case, the LSI MPT* SAS BIOS Configuration Utility loads after the system has completed a power-on self test. You can configure one or two IM or IME volumes per Fusion-MPT* controller. You can also configure one IM/IME and one Integrated Striping (IS) volume on the same controller, up to a maximum of 12 physical disk drives for the two volumes. In addition, you can create one or two hot spares for the IM/IME array(s).

The following guidelines also apply when creating an IM or IME volume:

- All physical disks in a volume must be either SATA (with extended command set support) or SAS (with SMART support). SAS and SATA disks cannot be combined in the same volume. However, you can create one volume with SAS disks and a second volume with SATA disks on the same controller.
- Disks must have 512 byte blocks and must not have removable media.
- An IM volume must have two drives. An IME volume can have three to ten drives. In addition, one or two hot spares can be created for the IM/IME volume(s).

Note: *If a disk in an IM/IME volume fails, it is rebuilt on the global hot spare if one is available. Intel recommends that you always use hot spares with IM/IME volumes.*

Creating an IM Volume

To create an IM volume with the LSI MPT* SAS BIOS Configuration Utility, follow these steps:

1. On the Adapter List screen, use the arrow keys to select an adapter.
2. Press <Enter> to go to the Adapter Properties screen, shown in [Figure 10](#).




Figure 10. Adapter Properties Screen

3. On the Adapter Properties screen, use the arrow keys to select **RAID Properties** on the screen and press <Enter>.
4. When prompted to select a volume type, select **Create IM Volume**.
The Create New Array screen shows a list of disks available to be added to a volume.
5. Move the cursor to the RAID Disk column and select a disk. To add the disk to the volume, change the No to Yes by pressing the <+> or <-> keys, or the space bar.
When the first disk is added, the utility prompts you to either keep existing data or overwrite existing data.
6. Press <M> to keep the existing data on the first disk or press <D> to overwrite it.
If you keep the existing data, this is called a data migration. The first disk will be mirrored onto the second disk, so any data you want to keep must be on the first disk selected for the volume. Data on the second disk is overwritten. The first disk must have 512 KB available for metadata after the last partition.
As disks are added, the Array Size field changes to reflect the size of the new volume.
7. [Optional] Add one or two global hot spares by moving the cursor to the hot spare column and pressing the <+> or <-> keys, or the space bar.

[Figure 11](#) shows an IM volume configured with one global hot-spare disk.




Figure 11. Create New Array Screen

- When the volume has been fully configured, press <C> and select **Save Changes**. Then exit this menu to commit the changes.

The LSI MPT SAS BIOS Configuration Utility* pauses while the array is created.

Creating an IME Volume

To create an IME volume with the LSI MPT* SAS BIOS Configuration Utility, follow these steps:

- On the Adapter List screen, use the arrow keys to select an Intel® IT/IR RAID Controller.
- Press <Enter> to load the Adapter Properties screen, shown in [Figure 10](#).
- On the Adapter Properties screen, use the arrow keys to select **RAID Properties** on the screen and press <Enter>.
- When prompted to select a volume type, choose **Create IME Volume**.

The Create New Array screen shows a list of disks that can be added to a volume.

- Move the cursor to the RAID Disk column and select a disk. To add the disk to the volume, change the No to Yes by pressing the <+> or <-> keys, or the space bar.
- Repeat this step to select a total of three to ten disks for the volume.

All existing data on all the disks you select will be overwritten. As you add disks, the Array Size field changes to reflect the size of the new volume.

- [Optional] Add one or two global hot spares to the volume by moving the cursor to the hot spare column and pressing the <+> or <-> keys, or the space bar.
- When the volume has been fully configured, press <C> and then select **Save changes**. Exit the menu to commit the changes.

The utility pauses while the array is created.

Creating a Second IM or IME Volume

Intel® IT/IR RAID Controllers allow you to configure two IM or IME volumes per controller. If one volume is already configured, and if there are available disk drives, there are two ways to add a second volume.

Option 1:

1. In the configuration utility, select an adapter from the Adapter List.
2. Select the **RAID Properties** option to display the current volume.
3. Press <C> to create a new volume.
4. Continue with either step 4 in “[Creating an IM Volume](#)“ on page 49 or step 4 in “[Creating an IME Volume](#)“ on page 50 to create a second volume.

Option 2:

1. On the Adapter List screen, use the arrow keys to select an Intel® IT/IR RAID Controller.
2. Press <Enter> to go to the **Adapter Properties** screen, shown in [Figure 10](#).
3. On the Adapter Properties screen, use the arrow keys to select **RAID Properties** and press <Enter>.
4. Continue with either step 4 in “[Creating an IM Volume](#)“ on page 49 or step 4 in “[Creating an IME Volume](#)“ on page 50 to create a second volume.

Managing Hot Spares

You can create one or two global hot-spare disks to protect the IM or IME volumes on an Intel® IT/IR RAID controller. Usually, you create global hot spares at the same time you create the IM/IME volume.

To add global hot-spare disks after an IM/IME volume has been created, follow these steps:

1. On the View Array screen, select **Manage Array**.
2. In the Manage Array screen select **Manage Hot Spares**, as shown in [Figure 12](#).



Figure 12. Manage Array Screen

3. Select a disk from the list by pressing the <+> or <-> key, or the space bar.
4. After you select the global hot-spare disk, press <C>.

An error message appears if the selected disk is not at least as large as the smallest disk used in the IM/IME volume(s). The global hot-spare disk must have 512 byte blocks, it cannot have removable media, and the disk type must be either SATA with extended command set support or SAS with SMART support.

If SATA disks are used for the IM/IME volume(s), the hot-spare disk must also be a SATA disk. If SAS disks are used, the hot-spare disk must also be a SAS disk. An error message appears if the selected disk is not the same type as the disks used in the IM/IME volumes.

5. [Optional] Select a second hot-spare disk.
6. Select **Save changes** then exit this menu to commit the changes.

The configuration utility pauses while the global hot spares are added.

To delete a global hot spare, follow these steps:

1. Select **Manage Hot Spare** on the Manage Array screen.
 2. Select **Delete Hot Spare** and then press <C>.
- If there are two hot spares, select one to delete.
3. Select **Save changes** then exit this menu to commit the changes.

The configuration utility pauses while the global hot spare is removed.

Other Configuration Tasks

This section explains how to perform other configuration and maintenance tasks for IM and IME volumes.

Viewing Volume Properties

To view the properties of volumes, follow these steps:

1. In the LSI MPT* SAS BIOS Configuration Utility, select an adapter from the Adapter List.
2. Select the **RAID Properties** option.

The properties of the current volume are displayed. If global hot spares are defined, they are also listed.

***Note:** If you create one volume using SAS disks, another volume using SATA disks, and another using global hot-spare disks, the hot-spare disks will only appear when you view the volume that has the same type of disks as the hot-spare disks.*

3. If two volumes are configured, press <Alt> + <N> to view the other array.
4. To manage the current array, select **Manage Array** and press <Enter>.

Synchronizing an Array

The Synchronize Array command forces the firmware to resynchronize the data on the mirrored disks in the array. It is seldom necessary to use this command, because the firmware automatically keeps the mirrored data synchronized during normal operation. When you use this command, one disk of the array is placed in a degraded state until the data on the mirrored disks has been resynchronized.

To force the synchronization of a selected array, follow these steps:

1. Select **Synchronize Array** on the Manage Array screen.
2. Press <Y> to start the synchronization, or <N> to cancel it.

Activating an Array

An array can become inactive if it is removed from one controller or computer and moved to another one. The Activate Array option allows you to reactivate an inactive array that has been added to a system. This option is only available when the selected array is currently inactive.

To activate a selected array, follow these steps:

1. Select **Activate Array** on the Manage Array screen.

2. Press <Y> to proceed with the activation, or press <N> to abandon it.

After a pause, the array will be active.

Note: *If there is a global hot-spare disk on the controller to which you have moved the array, the BIOS checks when you activate the array to determine if the hot spare is compatible with the new array. An error message appears if the disks in the activated array are larger than the hot-spare disk or if the disks in the activated array are not the same type as the hot-spare disk (SATA versus SAS).*

Deleting an Array

Caution: *If a volume has been deleted, it cannot be recovered. Before deleting an array, be sure to back up all data on the array that you want to keep.*

To delete a selected array, follow these steps:

1. Select **Delete Array** on the Manage Array screen.
2. Press <Y> to delete the array.

After a pause, the array is deleted. If there is a remaining array and one or two hot-spare disks, the BIOS checks the hot-spare disks to determine if they are compatible with the remaining array. If they are not compatible (i.e., they are too small or the wrong disk type), the firmware deletes them as well.

Note: *When an IM volume is deleted, the data is preserved on the primary disk. When an IME volume is deleted, the master boot records of all disks are deleted.*

Locating a Drive or Multiple Drives in a Volume

You can use the LSI MPT* SAS BIOS Configuration Utility to locate and identify a specific physical disk drive by flashing the drive's LED. You can also use the utility to flash the LEDs of all the disk drives in a RAID volume. There are several ways to do this:

- When you are creating an IM or IME volume, and a disk drive is set to Yes as part of the volume, the LED on the disk drive is flashing. The LED is turned off when you have finished creating the volume.
- You can locate individual disk drives from the SAS Topology screen. To do this, move the cursor to the name of the disk in the Device Identifier column and press <Enter>. The LED on the disk flashes until the next key is pressed.
- You can locate all the disk drives in a volume by selecting the volume on the SAS Topology screen. The LEDs flash on all disk drives in the volume.

Note: *The LEDs on the disk drives will flash as described above if the firmware is correctly configured and the drives or the disk enclosure supports disk location.*

Selecting a Boot Disk

You can select a boot disk in the SAS Topology screen. The selected disk is moved to scan ID 0 on the next boot, and remains at this position. This makes it easier to set the BIOS boot device options and to keep the boot device constant during device additions and removals. There can only be one boot disk.

To select a boot disk, follow these steps:

1. In the LSI MPT* SAS BIOS Configuration Utility, select an adapter from the Adapter List.
2. Select the **SAS Topology** option.

The current topology is displayed. If the selection of a boot device is supported, the bottom of the screen lists the <Alt> + option. This is the key for toggling the boot device. If a device is currently configured as the boot device, the Device Info column on the Topology screen will show the word “Boot.”

3. To select a boot disk, move the cursor to the disk and press <Alt> + .
4. To remove the boot designator, move the cursor down to the current boot disk and press <Alt> + .

This controller will no longer have a disk designated as boot.

5. To change the boot disk, move the cursor to the new boot disk and press <Alt> + .

The boot designator will move to this disk.

Note: *The firmware must be configured correctly in order for the <Alt> + feature to work.*

IS Configuration Overview

You can use the LSI MPT* SAS BIOS Configuration Utility to create one or two IS volumes, with up to a total of 12 drives, on an Intel® IT/IR RAID Controller. Each volume can have from two to ten drives. Disks in an IS volume must be connected to the same Intel® IT/IR RAID Controller, and the controller must be in the BIOS boot order.

Although you can use disks of different size in IS volumes, the smallest disk determines the “logical” size of each disk in the volume. In other words, the excess space of the larger member disk(s) is not used. Usable disk space for each disk in an IS volume is adjusted down to a lower value in order to leave room for metadata. Usable disk space may be further reduced to maximize the ability to interchange disks in the same size classification. The supported stripe size is 64 kbytes.

For more information about Integrated Striping volumes, see [“Features” on page 47](#).

Creating IS Volumes

The LSI MPT* SAS BIOS Configuration Utility is part of the Fusion-MPT* BIOS. When the BIOS loads during boot and you see the message about the utility, press <Ctrl> + <C> to start it. After you do this, the message changes to:

```
Please wait, Invoking Configuration Utility...
```

After a brief pause, the main menu of the LSI MPT* SAS BIOS Configuration Utility appears. On some systems, however, the following message appears next:

```
Configuration Utility will load following initialization!
```

In this case, the utility will load after the system has completed its power-on self test.

Follow the steps below to configure an Integrated Striping (IS) volume with the LSI MPT* SAS BIOS Configuration Utility. The procedure assumes that the required controller(s) and disks are already installed in the computer system. You can configure an IM/IME volume and an IS volume on the same Intel® IT/IR RAID Controller.

1. On the Adapter List screen of the LSI MPT* SAS BIOS Configuration Utility, use the arrow keys to select a RAID adapter.
2. Press <Enter> to go to the Adapter Properties screen, as shown in [Figure 13](#).




Figure 13. Adapter Properties Screen

3. On the Adapter Properties screen, use the arrow keys to select **RAID Properties** and press <Enter>.
4. When you are prompted to select a volume type, select **Create IS Volume**.
The Create New Array screen shows a list of disks that can be added to a volume.
5. Move the cursor to the RAID Disk column. To add a disk to the volume, change the No to Yes by pressing the <+> or <-> key, or the space bar. As disks are added, the Array Size field changes to reflect the size of the new volume.

There are several limitations when creating an IS (RAID 0) volume:

- All disks must be either SATA (with extended command set support) or SAS (with SMART support).
- Disks must have 512 byte blocks and must not have removable media.
- There must be at least two and no more than ten drives in a valid IS volume. Hot-spare drives are not allowed.




Figure 14. Create New Array Screen

6. When you have added the desired number of disks to the array, press <C> and then select **Save changes**. Then exit the menu to commit the changes. The configuration utility pauses while the array is created.

Creating a Second IS Volume

The Intel® IT/IR RAID Controllers allow you to configure two IS volumes, or an IS volume and an IM/IME volume. If one volume is already configured, and if there are available disk drives, there are two ways to add a second volume.

Option 1: Perform the following steps:

1. In the LSI MPT* SAS BIOS Configuration Utility select an adapter from the Adapter List. Select the **RAID Properties** option.
The current volume will be displayed.
2. Press <C> to create a new volume.
3. Continue with Step 4 of “Creating IS Volumes,” to create a second IS volume.

Option 2: Perform the following steps:

1. On the Adapter List screen, use the arrow keys to select an Intel® IT/IR RAID Controller.
2. Press <Enter> to go to the Adapter Properties screen, shown in [Figure 13](#).

3. On the Adapter Properties screen, use the arrow keys to select **RAID Properties** and press <Enter>.
4. Continue with step 4 in “Creating IS Volumes“ on page 56 to create a second IS volume.

Other Configuration Tasks

This section explains how to perform other configuration and maintenance tasks for IS volumes.

Viewing IS Volume Properties

To view the properties of IS volumes, follow these steps:

1. In the configuration utility, select an adapter from the Adapter List.
2. Select the **RAID Properties** option.
The properties of the current volume are displayed.
3. If more than one volume is configured, press <Alt> + <N> to view the next array.
4. To manage the current array, press <Enter> when the Manage Array item is selected.

Activating an Array

An array can become inactive if, for example, it is removed from one controller or computer and moved to another one. The “Activate Array” option allows you to reactivate an inactive array that has been added to a system. This option is only available when the selected array is currently inactive.

To activate an array once it has been selected, follow these steps:

1. Choose **Activate Array** on the Manage Array screen.
2. Press <Y> to proceed with the activation, or press <N> to abandon it. After a pause, the array will be active.

Deleting an Array

Caution: *Once a volume has been deleted, it cannot be recovered. Before deleting an array, be sure to back up all data on the array that you want to keep.*

To delete a selected array, follow these steps:

1. Select **Delete Array** on the Manage Array screen.

2. Press <Y> to delete the array, or press <N> to abandon the deletion. After a pause, the firmware deletes the array.

Note: The master boot records of all disks are deleted and data cannot be recovered.

Locating a Disk Drive, or Multiple Disk Drives in a Volume

You can use the LSI MPT* SAS BIOS Configuration Utility to locate and identify a specific physical disk drive by flashing the drive's LED. You can also use the utility to flash the LEDs of all the disk drives in a RAID volume. There are several ways to do this:

- When you are creating an IS volume, and a disk drive is set to Yes as part of the volume, the LED on the disk drive is flashing. The LED is turned off when you are finished creating the volume.
- You can locate individual disk drives from the SAS Topology screen. To do this, move the cursor to the name of the disk in the Device Identifier column and press <Enter>. The LED on the disk flashes until the next key is pressed.
- You can locate all the disk drives in a volume by selecting the volume on the SAS Topology screen. The LEDs flash on all disk drives in the volume.

Note: The LEDs on the disk drives will flash as described above if the firmware is correctly configured and the drives or the disk enclosure supports disk location.

Selecting a Boot Disk

You can select a boot disk in the SAS Topology screen. This disk is then moved to scan ID 0 on the next boot, and remains at this position. This makes it easier to set BIOS boot device options and to keep the boot device constant during device additions and removals. There can be only one boot disk.

To select a boot disk, follow these steps:

1. In the LSI MPT* SAS BIOS Configuration Utility, select an adapter from the Adapter List.
2. Select the **SAS Topology** option.

The current topology is displayed. If the selection of a boot device is supported, the bottom of the screen lists the <Alt> + option. This is the key for toggling the boot device. If a device is currently configured as the boot device, the Device Info column on the SAS Topology screen will show the word "Boot."

3. To select a boot disk, move the cursor to the disk and press <Alt> + .
4. To remove the boot designator, move the cursor down to the current boot disk and press <Alt> + .

This controller will no longer have a disk designated as boot.

5. To change the boot disk, move the cursor to the new boot disk and press <Alt> + .

The boot designator will move to this disk.

Note: *The firmware must be configured correctly for the <Alt> + feature to work.*

7 Intel® RAID BIOS Console 2 Utility

The Intel® RAID BIOS Console 2 utility provides a GUI utility to configure and manage RAID volumes. The utility configures disk groups and virtual drives. Because the utility resides in the RAID controller firmware, it is independent of the operating system.

The Intel® RAID BIOS Console 2 utility:

- Selects controller
- Displays controller properties
- Scans devices
- Displays the physical properties of devices
- Configures physical drives
- Defines virtual drives
- Displays virtual drive properties
- Initializes virtual drives
- Checks data for consistency

The Intel® RAID BIOS Console 2 utility provides a Configuration Wizard to guide you through the configuration of virtual drives and physical arrays.

Quick Configuration Steps

This section provides the steps to configure arrays and disk groups, and virtual drives using the Intel® RAID BIOS Console 2 utility. The following sections describe how to perform each action using the Intel® RAID BIOS Console 2 utility. The steps are as follows:

1. Power on the system.
2. Press <Ctrl>+<G> to start the Intel® RAID BIOS Console 2 utility.

Note: Some server boards have a BIOS SETUP option called "Port 60/64 Emulation" (or with other similar name). Please ensure this option is enabled in order to use Intel® RAID BIOS Console 2 successfully.

3. Start the Configuration Wizard.
4. Choose a configuration method.
5. Using the available physical drives, create arrays and disk groups.
6. Using the space in the arrays and disk groups, define the virtual drive(s).
7. Initialize the new virtual drives.

Detailed Configuration Steps using the Intel® RAID BIOS Console 2

Start the Intel® RAID BIOS Console 2 Utility

1. When the system boots, hold down the <Ctrl> key and press the <G> key when the following is displayed:
Press <Ctrl><G> to enter the RAID BIOS Console
After you press <Ctrl>+<G>, the Controller Selection screen appears.
2. Select a controller and click **Start** to begin the configuration.

Note: *If there is a configuration mismatch between the disks and the NVRAM, the utility automatically displays the Select Configuration screen. Choose whether the configuration should be read from the RAID array or from NVRAM. For more information, see “Configuration Mismatch Screen” on page 68.*





Screen and Option Descriptions

This section describes the Intel® RAID BIOS Console 2 screens and options.

Toolbar Options

Table 13 describes the Intel® RAID BIOS Console 2 toolbar icons.

Table 13. Intel® RAID BIOS Console 2 Toolbar Icon Descriptions

Icon	Description
	Return to the main screen.
	Return to the page you accessed immediately before the current page.
	Exit the Intel® RAID BIOS Console 2 utility.
	Silence the alarm.

Main Screen

From the main screen, you can scan the devices connected to the controller, select an Intel® RAID controller, and switch between the Physical Drives view and Virtual Drives view. The main screen also provides access to the following screens and tools:

- Controller Selection
- Controller Properties
- Scan Devices
- Virtual Drives
- Drives
- Configuration Wizard
- Physical View
- Events
- Exit




Figure 15. Intel® RAID BIOS Console 2 Menu

Controller Selection

This option allows you to choose an Intel[®] RAID controller installed in the system.

Adapter No.	Bus No.	Device No.	Type	Firmware Version
0.	11	0	Intel Server RAID Controller RS2BL080	2.0-03-0638

Figure 16. Intel[®] RAID BIOS Console 2 - Controller Selection

Controller Properties Screen

When you select the Controller Selection option on the main screen, the Intel[®] RAID BIOS Console 2 utility displays a list of the Intel RAID controllers in the system.

The Controller Properties screen allows you to view and configure the software and hardware of the selected controller.

Intel Server RAID Controller RS2BL080			
Serial Number	SR91700301	FRU	None
SubVendorID	0x8086	Drive Security Capable	Yes
SubDeviceID	0x9261	Drive Security Enabled	No
PortCount	8	Drive Security Method	FDE Only
HostInterface	PCIe	NVRAMSize	32 KB
Firmware Version	2.0-03-0638	Memory Size	512 MB
FW Package Version	12.0-1-0035	Min Stripe Size	8 KB
Firmware Time	Jun 2 2009;00:23:04	Max Stripe Size	1 MB
WebBIOS Version	3.0-17-Rel	Virtual Drive Count	1
Drive Count	10		

Figure 17. Controller Properties

- **Firmware Version:** The firmware version.
- **Host Interface:** The host interface for the installed RAID controller.
- **NVRAM Size:** The NVRAM size on the RAID controller.
- **Firmware Time:** The firmware release date/time.
- **Min Stripe Size:** The minimum stripe size used to read and write data.
- **WebBIOS Version:** The BIOS version for the Intel[®] RAID BIOS Console 2.
- **Sub Device ID:** The sub-device ID (identification) for the RAID controller.
- **Sub Vendor ID:** The sub-vendor ID (identification) for the RAID controller.
- **Port Count:** Number of ports available.

- **Memory Size:** The memory size of the installed DIMM (Dual In-Line Memory Module).
- **Max Stripe Size:** The maximum stripe size.
- **Physical Disk Count:** The number of physical disks connected to the RAID controller.

Additional Controller Properties

To access the screen that displays the additional controller properties, click **Next** on the Controller Properties screen. To change one of the properties displayed in the screen below, select the new value and click **Submit**.

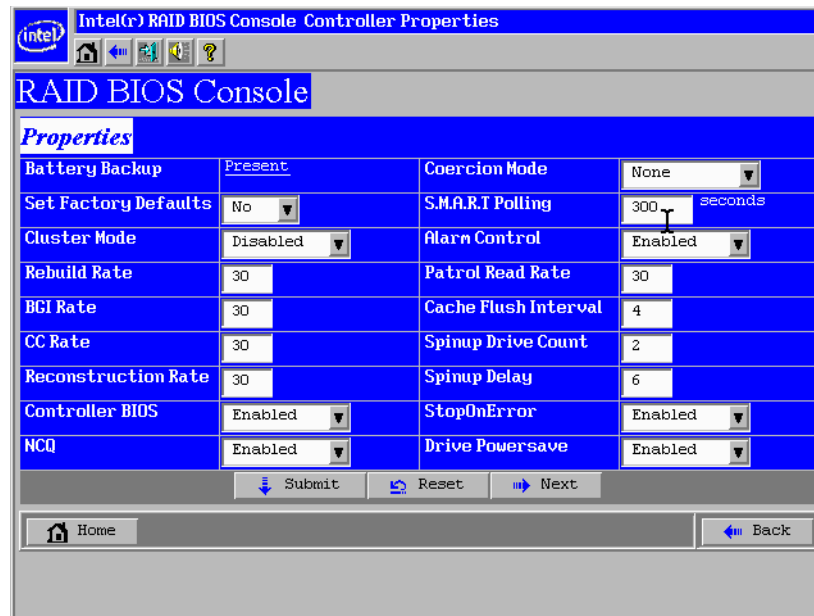


Figure 18. Additional Controller Properties

- **Battery Backup:** Indicates if a battery backup unit is installed.
- **Set Factory Defaults:** Change this field to **Yes** to reset the RAID controller settings to the factory defaults.
- **Cluster Mode:** Enable this field if the RAID controller is used in a cluster.
- **Rebuild Rate:** Enter a number between 0 and 100 to control the rate at which a future rebuild will be performed on a disk group.
- **Patrol Read Rate:** A patrol read is a preventive procedure that monitors physical disks to locate and resolve potential problems that could lead to disk failure. Enter a number between 0 and 100 to control the rate at which patrol reads are performed.
- **BGI Rate (Background Initialization Rate):** Background initialization makes the virtual drive immediately available for use, even while initialization is occurring. Enter a number between 0 and 100 to control the rate at which virtual drives are initialized in the background.
- **CC Rate (Check Consistency Rate):** A consistency check scans the consistency of data on a fault-tolerant disk to determine if the data is corrupted. Enter a number between 0 and 100 to control the rate at which a consistency check is done.
- **Reconstruction Rate:** Enter a number between 0 and 100 to control the rate at which the reconstruction of a virtual drive occurs.

- **Adapter BIOS:** Determines whether the Option ROM is loaded.
- **Coercion Mode:**
 - None: No coercion of size.
 - 128M: The software rounds the drive capacity down to the next 128 MB boundary and then up to the nearest 10 MB until the coerced capacity is larger than the actual drive size. It is then reduced by 10 MB.
 - 1G: The software rounds the drive capacity down to the nearest 1 GB boundary and then down by 1 MB. This corresponds to the terms most drive manufacturers use.
- **PDF Interval:** The PDF interval is the predictive disk failure polling interval. This is the time needed between disk polls to perform SMART polling.
- **Alarm Control:** Disable the alarm to turn off the on-board speaker alarm.
- **Interrupt Throttle Count** and **Interrupt Throttle Time:** Sets the interrupt throttle and count times. This is the number of times that interrupts are coalesced and the amount of time that firmware holds an interrupt before passing it to the host software. Set values lower for better performance—be aware that latency is impacted by these settings.
- **Cache Flush Interval:** This sets the cache flush interval. Valid settings are 2, 4, 6, 8, or 10 seconds.
- **Spinup Drive Count:** This setting controls the number of drives that spin up at one time.
- **Spinup Delay:** After the RAID controller completes its initialization process, the initial delay value defines the number of seconds before the first disk interrogation request is issued to the array or disk group. Do not change this value.
- **Stop On Error:** Stops system POST if any error is detected.
- **NCQ:** Enables NCQ (Native Command Queuing) to optimize physical drive performance and life.
- **Stop CC On Error:** Stops Consistency Check if any error is detected.
- **Schedule CC:** Schedules a Consistency Check.
- **Maintain PD Fail History:** Enables tracking of bad PDs across reboot.

Scan Devices Option

When you select the Scan Devices option on the Main screen, the Intel® RAID BIOS Console 2 checks the physical and virtual drives for any changes of the drive status. The Intel® RAID BIOS Console 2 displays the results of the scan in the physical and virtual drive descriptions.

Virtual Drives Screen

You can access the virtual drives screen by clicking on a virtual drive in the virtual drive list on the main screen. The upper right section of the screen displays the virtual drives that currently exist. The Virtual Drives screen provides options to:

- **Initialize the virtual drives:** The Slow Initialize option initializes the selected virtual drive by writing zeroes to the entire volume. You should initialize each new virtual drive that you configure.

Warning: *Initializing a virtual drive deletes all information on the physical drives that compose the virtual drive.*

- Check consistency (CC): This option verifies the correctness of the redundancy data and is available for arrays and disk groups using RAID 1, 5, 6, 10, 50, or 60. If a difference in the data is found, the Intel® RAID BIOS Console 2 assumes that the data is accurate and automatically corrects the parity value.
- Display the virtual drive properties: Through the Properties option, you can:
 - Display the virtual drive properties (such as RAID level, virtual drive size, and stripe size).
 - Display the read, write, Access, Disk Cache, BGI (Background Initialization), and I/O policies.
 - Change the read, write, Access, Disk Cache, BGI, and I/O policies.
 - Select Write Through, Write Back with BBU, or Always Write Back.
 - Start initialization.
 - Start a consistency check.

After setting any property, click **Go** to perform the selected operation. Click **Change** to apply any policy changes.

Physical Drives Screen

This screen displays the physical drives for each channel or port. From this screen, you can rebuild the physical arrays or disk groups, or view the properties for the physical drive you select.

- Click **Reset** to return to the configuration that existed before you made any changes.
- Select Properties and click **Go** to view the properties. An unconfigured drive can be made into a hot spare from the Properties screen.

Configuration Wizard Option

This option enables you to clear a configuration, create a new configuration, or add a configuration. [“Setting Up a RAID Array Using the Configuration Wizard” on page 68](#) provides detailed steps for using the Configuration Wizard.

Events Screen

This option displays the events generated by physical drives, physical devices, enclosure, the Intel® Smart Battery, and SAS controller. See [Appendix B: “Events and Messages” on page 173](#) for events and message descriptions.

Physical View/Logical View Option

This option toggles between Physical View and Logical View.

Exit

This option allows you to exit and reboot the system.

Configuration Mismatch Screen

A configuration mismatch occurs when the data in the NVRAM and the hard disk drives are different. It automatically displays after POST when a configuration mismatch occurs. The Configuration Mismatch screen allows you to:

- Select **Create New Configuration** to delete the previous configuration and create a new configuration.
- Select **View Disk Configuration** to restore the configuration from the hard disk.
- Select **View NVRAM Configuration** to restore the configuration from the NVRAM.

Setting Up a RAID Array Using the Configuration Wizard

This section provides detailed steps for using the Configuration Wizard to set up a RAID array.

1. Start the Configuration Wizard by selecting the **Configuration Wizard** icon on the Intel® RAID BIOS Console 2 main screen.




Figure 19. Intel® RAID BIOS Console 2 - Configuration Types

2. Select **New Configuration** and click Next.
3. Then select Virtual Drive Configuration and click Next.




Figure 20. Selecting Configuration

4. Choose the configuration method and click Next.

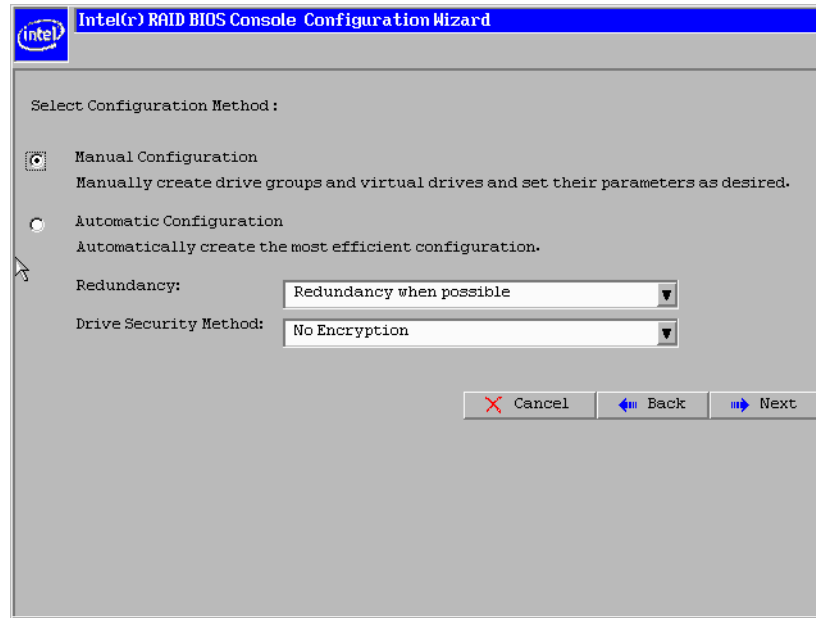


Figure 21. Intel® RAID BIOS Console 2 - Configuration Methods

The following configuration methods options are provided:

- Automatic Configuration

There are two options in Redundancy. **Redundancy When Possible** or **No Redundancy**.

Redundancy When Possible configures RAID 1 for systems with two drives or RAID 5 for systems with three or more drives, or RAID 6 for systems with three or more drives. All available physical drives are included in the virtual drive using all available capacity on the disks.

No Redundancy configures all available drives as a RAID 0 virtual drive.

There is a **Drive Security Method** option which is reserved to be enabled in future.

***Note:** You must designate hot-spare drives before starting auto configuration using all available capacity on the disks.*

- Manual Configuration

Allows you to configure the RAID mode.

***Note:** Automatic Configuration cannot be used for RAID 10, 50, or 60 or with mixed SATA and SAS drives.*

Creating RAID 0, 1, 5, or 6 using Intel® RAID BIOS Console 2 (detailed)

This section describes the process to set up RAID modes using the custom configuration options.

1. When the server boots, hold the <Ctrl> key and press the <G> key when the following is displayed:
Press <Ctrl><G> to enter RAID BIOS Console
The Controller Selection screen appears.
2. Select a controller and click **Start** to begin the configuration.
3. Choose Manual Configuration and click **Next** (see Figure 21).
4. At the Disk Group Definition (DG Definition) screen, hold down the <Ctrl> key and click each drive you want to include in the array or disk group.

See “RAID Levels” on page 9 for the required minimum number of drives that must be added.

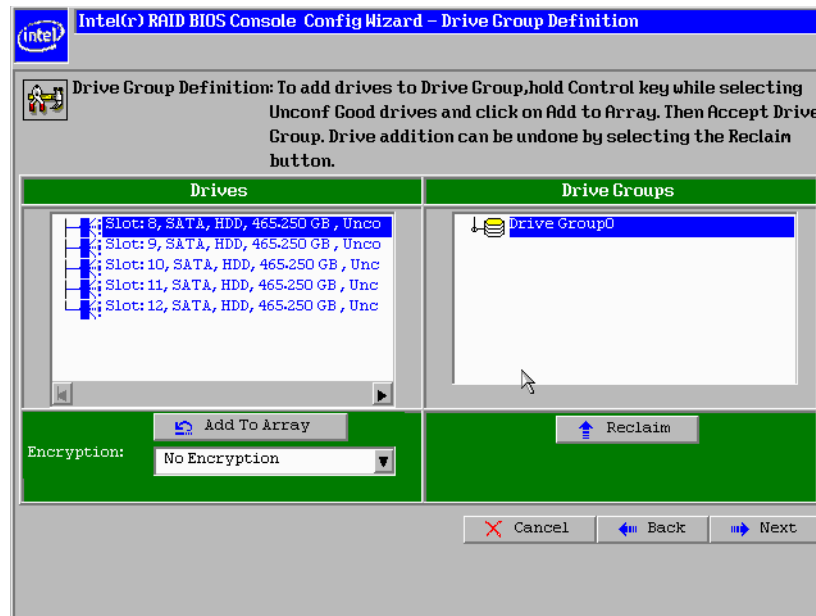


Figure 22. Intel® RAID BIOS Console 2 - Add Physical Drives to Array

5. Click **Add To Array**. If you make a mistake and need to remove drives, click **Reclaim**.
6. Click **Next**.
7. In the next screen, click **Add to Span** and then click **Next**.

8. On the VD Definition window, select RAID 0, 1, 5, or 6 from the first dropdown box.
9. Enter the virtual drive size in the **Select Size** box.
This example shows a specific size. Depending on the RAID level you choose, you may need to manually type in the expected volume size. The possible sizes for some RAID levels are listed on right panel of the screen for reference.
10. If needed, change the Stripe Size, the policies for Access, Read, Write, IO, and Disk Cache and decide whether to use background initialization.

For information about setting these parameters, see [“Setting Drive Parameters” on page 80](#).

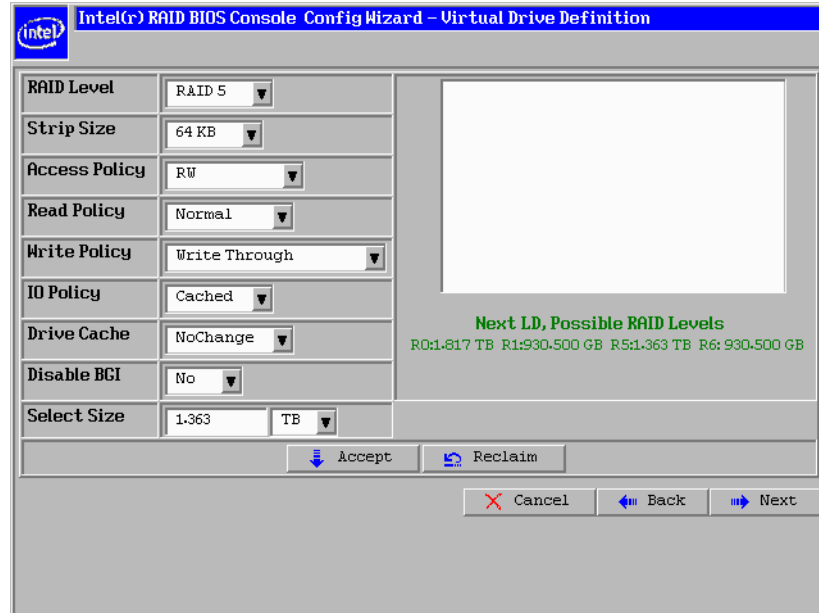


Figure 23. Intel® RAID BIOS Console 2 - Set Array Properties

11. Click **Accept** to accept the changes, or click **Reclaim** to delete the changes and return to the previous settings.
The Intel® RAID BIOS Console 2 configuration utility displays a preview of the configuration.
12. Click **Accept** to save the configuration, or click **Back** to return to the previous screens and change the configuration.




Figure 24. Intel® RAID BIOS Console 2 - Confirm Configuration

13. Click **Accept** as necessary in the screens that follow. You are prompted to save the configuration and then to initialize the virtual drive.

14. Click **Yes** to initialize the new drive.
15. Click Initialize to begin the initialization process.
 - Fast initialization runs a quick preliminary initialization and then runs full initialization in the background after the operating system is booted.
 - Slow initialization may take several hours or even days to complete.




Figure 25. Intel® RAID BIOS Console 2 - Initialization Speed Setting

16. Click **Home** to return to the main configuration screen.
17. Select an additional virtual drive to configure or exit the Intel® RAID BIOS Console 2 configuration utility and reboot the system.

Creating RAID 10, RAID 50, and RAID 60 using Intel® RAID BIOS Console 2

RAID 10, RAID 50, and RAID 60 require setting up multiple RAID arrays/disk groups.

1. When the server boots, hold the <Ctrl> key and press the <G> key when the following is displayed:
Press <Ctrl><G> to enter the RAID BIOS Console
After you press <Ctrl>+<G>, the Controller Selection screen appears.
2. Select a controller and click **Start** to begin the configuration.
3. Select **Custom Configuration** and click **Next** (see [Figure 21](#)).
4. At the Virtual Drive Definition (VD Definition) screen, hold down the <Ctrl> key and click each drive you want included in the first array.

- For RAID 10, use two drives.
- For RAID 50, use at least three drives.
- For RAID 60, use at least three drives.

5. Click **Add To Array**, and then click **Accept DG** in the right pane to confirm.

The first group of drives appears as a disk group in the right pane. These drives are no longer available in the left pane.

6. From the drives that are available in the left pane, choose an additional group of drives and again click **Add To Array**, and click **Accept DG** to confirm.

Each disk group must contain the identical quantity and size of drives.

Multiple drive groups are now displayed in the right pane. You can add up to eight arrays to the right pane for either RAID 10, RAID 50, or RAID 60.

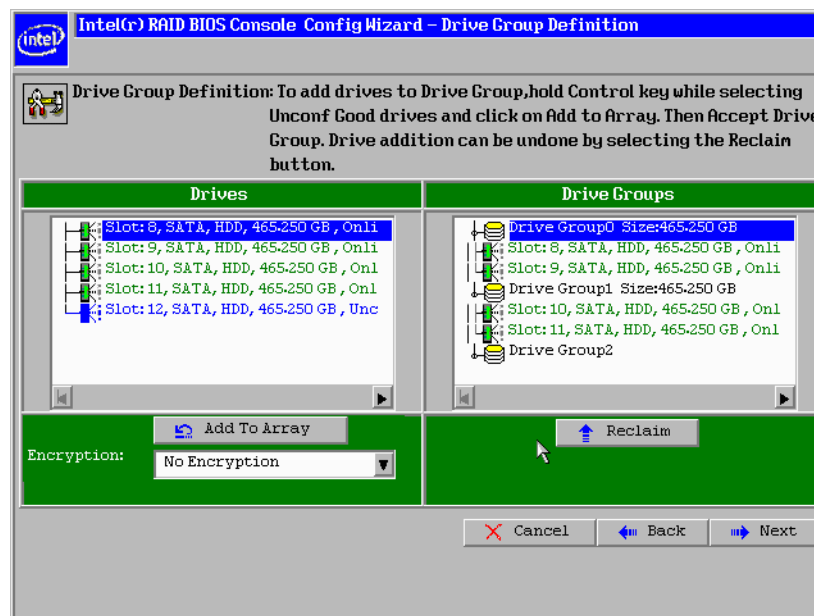


Figure 26. Intel® RAID BIOS Console 2 – Multiple Disk Groups for RAID 10, 50, or 60

7. Select all arrays or disk groups that are to be spanned in the RAID 10, 50, or 60 array by holding down the <Ctrl> key and selecting each array/disk group in the right pane.
8. Click **Next**.
9. In the next screen, click **Add to SPAN** to move all arrays from the left pane to the right pane. Use <Ctrl> to select all SPANs on the right pane.
10. Click **Next**.
11. At the Virtual Drive Definition (VD Definition) screen, select either RAID 10, RAID 50, or RAID 60 from the RAID Level drop-down.
RAID 10 is illustrated below.
12. Select the appropriate Stripe Size, Access Policy, Read Policy, Write Policy, IO Policy, Disk Cache Policy, and Enable/Disable BGI for your application.
For information about setting these parameters, see [“Setting Drive Parameters” on page 80](#).
13. Set the drive size to a number in MB that is a size greater than the size of the RAID 1, RAID 5, or RAID 6 size listed in the disk group.

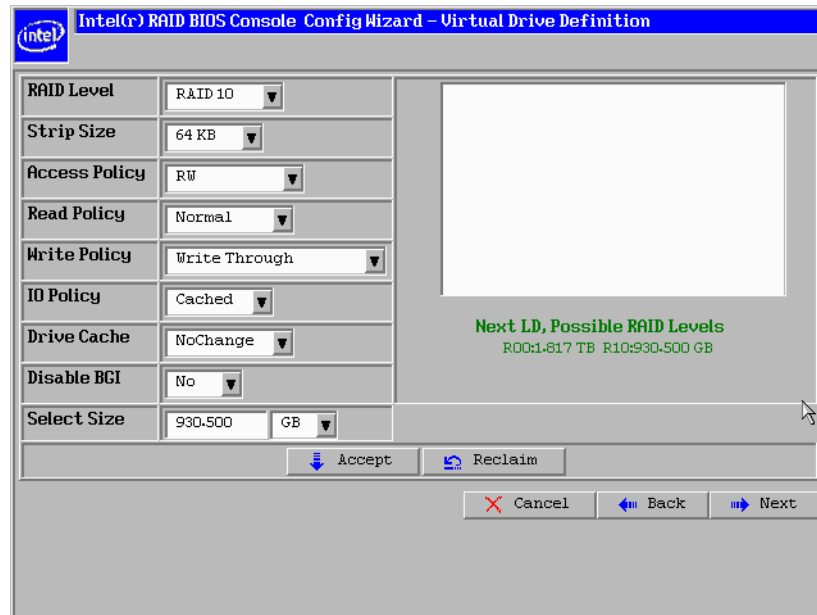


Figure 27. Intel® RAID BIOS Console 2 – Spanning Multiple Arrays

14. Click **Next** if the application does not automatically progress to the next screen.

The configuration preview screen displays the virtual drive as shown below. The configuration preview screen displays the virtual drive (RAID 1 for RAID 10, or RAID 50 or RAID 60).




Figure 28. Intel® RAID BIOS Console 2 - Viewing Completed Settings

15. Click **Accept** to save the configuration.
16. When asked to save the configuration, click **Yes**.
This will store the configuration in the RAID controller.
17. When asked to initialize the drive, click **Yes**.
18. Select **Fast Initialize** and click **Go**.
The drives will initialize based on the RAID settings.

Note: *Slow Initialize initializes the entire drive and may take several hours to complete.*

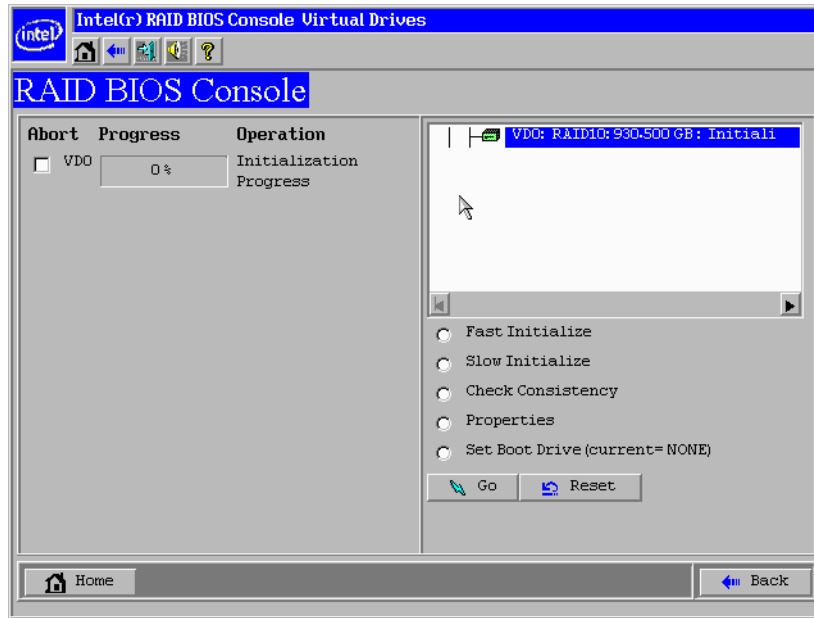


Figure 29. Intel® RAID BIOS Console 2 - Initialization Settings

19. Click **Home** at the Intel® RAID BIOS Console 2 screen to return to the main screen.
The RAID 10, RAID 50, or RAID 60 virtual drives are displayed. The following figure shows the RAID 10 virtual drives.

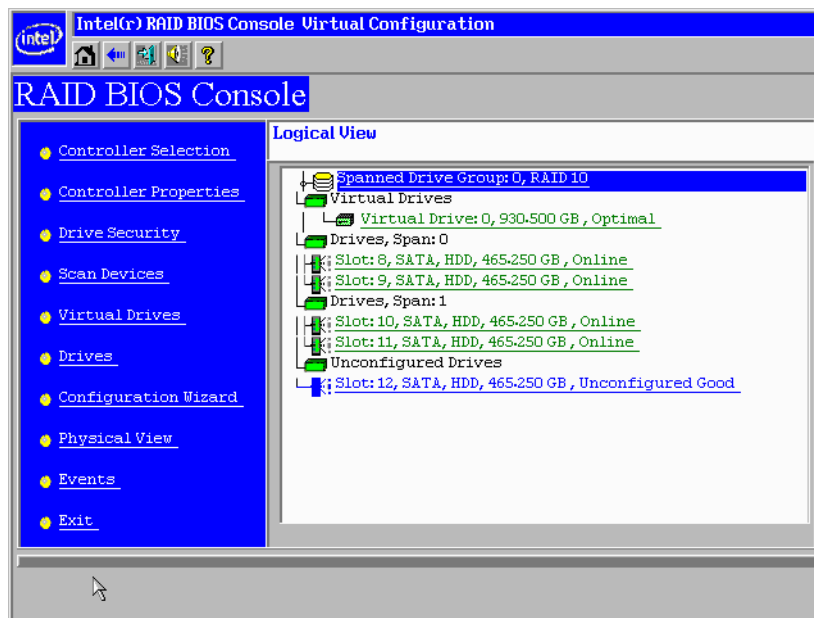


Figure 30. Intel® RAID BIOS Console 2 - RAID 10 Final Screen

20. Under Virtual Drives, select **Virtual Drive 0: RAID 10**, or select **Virtual Drive 0: RAID 50**, or select **Virtual Drive 0: RAID 60** to display the drive properties.




Figure 31. Intel® RAID BIOS Console 2 - RAID 10 Properties Screen




Figure 32. Intel® RAID BIOS Console 2 - RAID 50 Properties Screen

Setting Drive Parameters

The following fields are displayed in the VD Definition screen (see [Figure 23](#) and [Figure 27](#)), which can be used to set the virtual drive parameters:

- **RAID Level:**
 - RAID Level 0: Data striping
 - RAID Level 1: Data mirroring
 - RAID Level 5: Data striping with parity
 - RAID Level 6: Distributed Parity and Disk Striping
 - RAID level 10: Striped mirroring
 - RAID Level 50: Striped RAID 5
 - RAID Level 60: Distributed parity, with two independent parity blocks per stripe
 - **Stripe Size:** Specify the size of the segment written to each disk. Available stripe sizes are 4, 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes.
 - **Access Policy:** Select the type of data access that is allowed for this virtual drive. The choices are Read/Write, Read Only, or Blocked.
 - **Read Policy:** Enables the read-ahead feature for the virtual drive. Read Adaptive is the default setting.
 - Normal: The controller does not use read-ahead for the current virtual drive.
 - Read-ahead: Additional consecutive stripes are read and buffered into cache. This option will improve performance for sequential reads.
 - Adaptive: The controller begins using read-ahead if the two most recent disk accesses occurred in sequential sectors.
 - **Write Policy:** Determines when the transfer complete signal is sent to the host. Write-through caching is the default setting.
 - Write-back caching (Further classified as Write Back with BBU or Always Write Back, which means Write Back is always enabled even if BBU is bad or missing): The controller sends a data transfer completion signal to the host when the controller cache receives all of the data in a transaction. Write-back caching has a performance advantage over write-through caching, but it should only be enabled when the optional battery backup module is installed. The risk of using Always Write Back should be fully recognized.
 - Write-through caching: The controller sends a data transfer completion signal to the host after the disk subsystem receives all the data in a transaction. Write-through caching has a data security advantage over write-back caching.
- Caution:** *Do not use write-back caching for any virtual drive in a Novell NetWare* volume.*
- **IO Policy:** Applies to reads on a specific virtual drive. It does not affect the read-ahead cache.
 - Cached IO: All reads are buffered in cache memory.
 - Direct IO: Reads are not buffered in cache memory. Data is transferred to cache and to the host concurrently. If the same data block is read again, it comes from cache memory.

- **Disk Cache Policy:** The cache policy applies to the cache on physical drives of the current array.
 - Enable: Enable disk cache. Enabling the disk cache in Write-back mode provides little or no performance enhancement, while the risk of data loss due to power failure increases.
 - Disable: Disable disk cache.
 - NoChange: Leave the default disk cache policy unchanged.
- **Disable BGI:** Enable or disable background initialization. Set this option to “Yes” to disable background initialization.
- **Select Size:** Set the size of the virtual drive in megabytes. The right pane of the virtual drive configuration window lists the maximum capacity that can be selected, depending on the RAID level chosen.

Creating a Hot Spare

To create a hot spare, follow these steps:

1. On the main screen, select the drive that should be used as the hot spare.




Figure 33. Intel® RAID BIOS Console 2 – Choosing a Hot Spare Drive

2. Select the disk group.

3. Click one of the following:
 - Click **Make Dedicated HSP** to add the drive as a hot spare dedicated for certain virtual drives.
 - Click **Make Global HSP** if you want to create a global hot spare for all disk groups.




Figure 34. Intel® RAID BIOS Console 2 - Setting a Hot Spare Drive

4. Click **Go** to create the hot spare.

The Drive State changes to **HOTSPARE**, as shown below.




Figure 35. Intel® RAID BIOS Console 2 - Viewing Hot Spare

5. Click **Home** to return to the main screen.

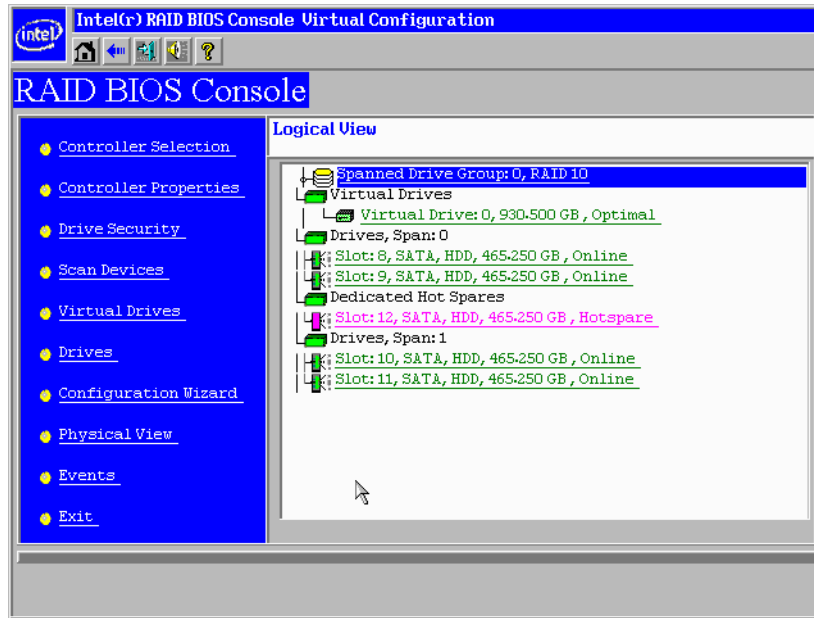


Figure 36. Intel® RAID BIOS Console 2 – Main Screen showing Hot Spare Drive

Viewing Event Details

Events contain information, warnings, and fatal events. Events can be captured on various RAID controller components, such as the battery, physical card, and within the configuration. You can view these using the following steps.

1. On the Main screen, select **Events** from the menu at the left.

The Events screen appears.




Figure 37. Intel® RAID BIOS Console 2 – Event Information Screen

2. Select the component to display from the **Event Locale** list.
3. Select the type of event to display from the **Event Class** drop-down.
4. Type the **Start Sequence#** and the **# of Events** to display.

The following example shows a selection that was made for informational events for the virtual drive, starting at sequence number 120 and displaying 10 events.




Figure 38. Intel RAID BIOS Console 2 - Selecting Events to View

5. Click **Go**.

In the following example, virtual drive events for informational messages were selected, starting with the event at sequence 120. The events screen displays the event information for sequence number 120 in the right pane.

Note: In the following example, the Start Sequence# and the # of Events fields display 0. This is because when you click **Go** to display the events, these fields automatically reset to 0.




Figure 39. Intel RAID BIOS Console 2 - Viewing an Event

6. Click **Next** to view the next message.

8 Intel® RAID Web Console 2

The Intel® RAID Web Console 2 enables you to configure, monitor, and maintain Intel® RAID controllers and the storage-related devices connected to the RAID controller. The Intel® RAID Web Console 2 graphical user interface (GUI) makes it easy to create and manage storage configurations.

Note: *Depending on different version of Intel(R) RAID Web Console 2, the actual screen displayed could be different from this section. This section is for reference only.*

Configuration Functions

The Intel® RAID Web Console 2 enables you configure the controllers, disk drives, battery backup units, and other storage-related devices installed on a system.

- The Configuration Wizard simplifies the process of creating disk groups and virtual drives.
- Auto Configuration mode automatically creates the best possible configuration for the available hardware.

Note: *Auto Configuration cannot be used for RAID 10, 50, or 60 or with mixed SATA and SAS drives.*

- Guided Configuration mode asks you a few brief questions about the configuration, and then creates the array for you.
- Manual Configuration mode gives you complete control over all aspects of the storage configuration.
- The Reconstruction Wizard enables you to increase or reduce the size of a virtual drive and to change the RAID level of an array.

Monitoring Functions

The Intel® RAID Web Console 2 displays information on the status of virtual drives, physical disks, and other storage-related devices on the systems you are monitoring. System errors and events are recorded in an event log file and displayed on the screen. Special device icons appear on the screen to notify you of disk failures and other situations that require immediate attention.

Note: *Intel® RAID Web Console 2 “Client Only” Installation type on the Client system doesn’t support configuring email/alerting settings on local system. Necessary settings for local system need to be configured locally.*

Maintenance Functions

Use the Intel® RAID Web Console 2 to perform system maintenance tasks such as running patrol read operations, updating firmware, and running consistency checks on arrays and disk groups that support redundancy.

Hardware and Software Requirements

The hardware requirements for Intel® RAID Web Console 2 software are as follows.

- PC-compatible computer system with at least one Intel® Xeon® architecture processor or an Intel® Extended Memory 64 Technology (Intel® EM64T) 64-bit processor
- At least 256 Mbytes of system memory
- Hard disk drive with at least 50 MB available free space
- A supported operating system:
 - Microsoft Windows 2000*, Microsoft Windows Server 2003*, Microsoft Windows Server 2008*, or Microsoft Windows XP*
 - Red Hat* Linux Enterprise 3.0, 4.0, or 5.0 with or without corresponding service packs
 - SuSE* Enterprise Linux 9.0, 10.0, or 11.0 with or without corresponding service packs
 - VMWare* ESX 3i, 4i

Installing the Intel® RAID Web Console 2 on a Microsoft Windows* Operating System

To install the Intel® RAID Web Console 2 on a Microsoft Windows 2000*, Microsoft Windows Server 2003*, Microsoft Windows Server 2008*, or Microsoft Windows XP* system, complete the following steps:

1. Insert the *Resource CD* in the CD-ROM drive.
OR
Download the Intel® RAID Web Console 2 installation package from <http://www.intel.com/support/motherboards/server/>.
2. When the Welcome screen is displayed, click **Next**.
3. When the next screen is displayed, read and accept the user license and click **Next**.
The Customer Information screen is displayed, as shown in [Figure 40](#).

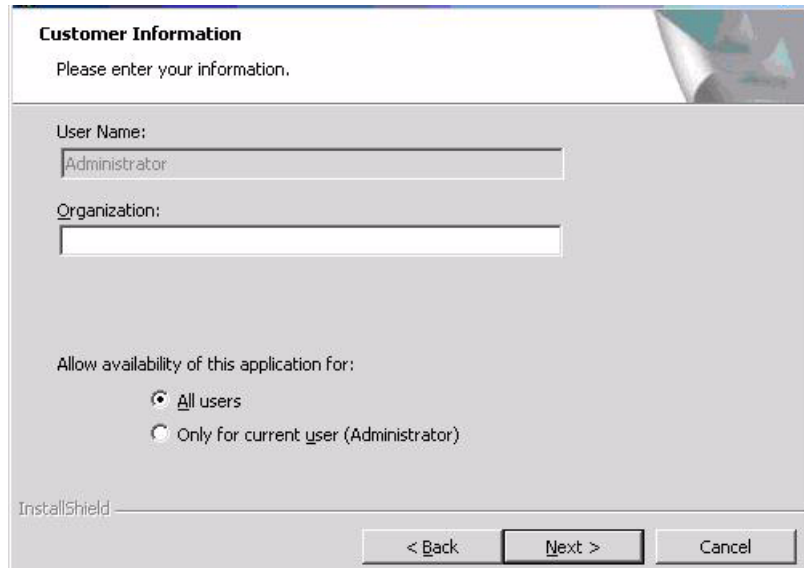


Figure 40. Intel® RAID Web Console 2 – Customer Information Screen

4. Enter your user name and organization name. At the bottom of the screen, select an installation option:
 - If you select **All users**, any user with administrative privileges can view or change the RAID configurations.
 - If you select **Only for current user (Administrator)**, only you can view or change the RAID configurations.
5. Click **Next** to continue.
6. Accept the default Destination Folder, or click **Change** to select a different destination folder. Click **Next** to continue.

The Setup Type screen is displayed, as shown in [Figure 41](#).

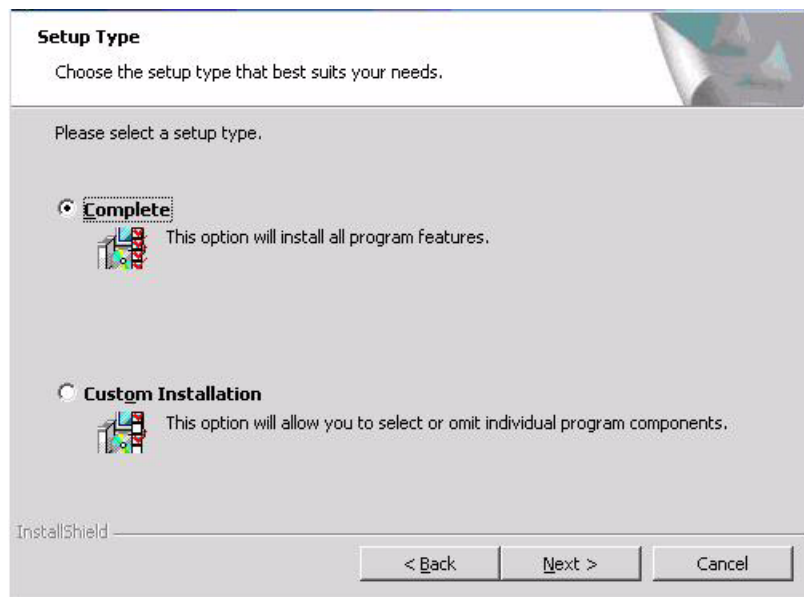


Figure 41. Setup Type Screen

7. Select one of the following setup options.

- Select **Complete** if you are installing the Intel® RAID Web Console 2 on a system.
 - Select **Custom Installation** if you want to select or omit individual program components to install. When selected, more custom options are provided. Please read the instructions on the screen carefully to make sure you select the preferred setup type.
8. Click **Next** to proceed and then click **Finish** to complete the installation process.

Installing the Intel® RAID Web Console 2 on Linux or SuSE* Linux Enterprise Server

To install Intel® RAID Web Console 2 on a system running Red Hat* Linux 3.0, 4.0, or 5.0, or SuSE* Linux Enterprise Server 9,10 or 11, complete the following steps:

1. Unzip the file `ir3_Linux_RWC2_v... .tgz`.
2. In the unzipped files, read the `readme.txt` file and follow the instructions closely, in order to make the installation successful.
3. Run `install.sh` and follow the instructions that display on the screen.

The three setup options are the same as those shown in step 7 of the Microsoft Windows* installation instructions.

Intel® RAID Web Console 2 Support and Installation on VMWare

This section documents the installation of Intel® RAID Web Console 2 on VMWare Classic (with console operating system) and on the VMWare ESX 3i operating system.

Installing Intel® RAID Web Console 2 for VMWare Classic

VMWare does not support any graphics components. In order to install the server component of Intel® RAID Web Console 2 (Intel® RAID Web Console 2 without popup and client), run the script file `ServerInstall.sh` and choose setup type "4". For the noninteractive (silent) mode, run `./ServerInstall.sh -x`.

In order to manage Intel® RAID Web Console 2 on a VMWare system, install Intel® RAID Web Console 2 Client from a remote system.

Uninstalling Intel® RAID Web Console 2 for VMWare

To uninstall the Server Component of Intel® RAID Web Console 2 on VMWare, use the `Uninstall` command in the Program menu or run the script `/uninstaller.sh`.

Note the following points:

1. An Intel® RAID Web Console 2 upgrade is supported in this release. This release can be upgraded by future releases.

2. To shut down the Intel® RAID Web Console 2 Framework service, run the following command:

```
/etc/init.d/vivaldiframeworkd stop
```

It is recommended that you stop the Monitor service before you stop the Intel® RAID Web Console 2 Framework service. To stop the Monitor service run the following command:

```
/etc/init.d/mrmonitor stop
```

Installing Intel® RAID Web Console 2 Support on the VMWare ESX

This section outlines the product requirements needed to support the VMWare ESX operating system. Classic VMWare includes a Service Console that is derived from the Linux 2.4 kernel, but with reduced functionality.

The Intel® RAID Web Console 2 server part cannot be installed directly in VMWare ESX 3i. Management is performed through Intel® RAID Web Console 2 installed on a remote machine (Linux/Windows).

The Linux installer of Intel® RAID Web Console 2 works under console with minimal changes. Hardware RAID is currently supported in ESX 3.x.

Note: *There is a known limitation that virtual drives that are created or deleted will not be reflected to the kernel. The workaround is to reboot the server or to run `esxcfg-rescan <vmhba#>` from COS shell.*

The network communication is a key element for a proper setup. Therefore, we recommend that you install the management on a VM within the ESXi. Follow these steps to install and configure Intel® RAID Web Console 2 support on the VMWare ESX operating system:

1. Network Configuration of the ESXi Host: -
 - Assignment of a ESXi hostname:
 - Even if it is not relevant for your network, you need a FQDN (Fully Qualified Domain Name). Example: local.lsi.com to be entered using the local ESXi console.
 - Configuration of a virtual network environment:
 - i. You can use the already existing Vswitch, which has a VMkernel port already attached for the communication.
 - ii. Alternatively, you can build a new Vswitch without a link to the Host network card.

Which one of the two possibilities to choose depends on your application. It is recommended to choose between both possibilities at a early stage, because the creation of a new Vswitch with VMkernel requires a reboot. For those who want to reach the target as quickly as possible, no change is recommended.

- Configuration of the IP address:
 - Configure the IP address. The address must be accessible by the VM that will be installed next.

2. VM Installation: -

Install the operating system as usual, including the VMWare guest tools. The virtual network card should be linked to a Vswitch that has a VMKernel port attached. For a quick installation, no change is recommended.

3. Intel® RAID Web Console 2 Installation: -

- Install Intel® RAID Web Console 2 with the option "complete".
4. VM Network Configuration: -
- Case 1: Your network contains a DNS server:
Configure a host entry that belongs to your internal zone and make sure that the FQDN of the ESXi server can be resolved. (Example: local.lsi.com and 192.19.221.186)
 - Case 2: Your network does not have a DNS server:
Edit your file C:\windows\system32\drivers\etc\hosts and add another entry:

IP of the ESXi Host	FQDN of the ESXi Host
192.19.221.186	local.lsi.com

5. Final Steps: -
- Reboot the VM and start the Intel® RAID Web Console 2. The ESXi server should now appear in the list of the found hosts. You can now log in with the root account name and password of the ESXi Host.

Limitations

The following are the limitations of this installation and configuration:

- There is no active event notification, for example, by popup or email
- There is no status information for the controller
- There is no user authentication.
- Events are collected as long as Intel® RAID Web Console 2 runs on the Client.
- Intel® RAID Web Console 2 responds more slowly.

Differences in Intel® RAID Web Console 2 for VMware ESXi

The following are some of the differences in the Intel® RAID Web Console 2 utility when you manage a VMWare server.

1. The following limitations apply to the system information exposed through the application:
 - Only the IP address and the Host name display.
 - The operating system type and the operating system architecture do not appear.
 - There is no support for the controller health information.

The following are the Intel® RAID Web Console 2 screens affected:

 - Initial Intel® RAID Web Console 2 framework (hosts) discovery screen: No health information or operating system type display.
 - Server property page: Only the IP address and the Host name display; the operating system type and operating system architecture do not display.
2. Authentication support:

Access control is not supported. There is no support for full view or view only access modes. It is always full view access, and multiple clients can have full view access at the same time on the same server

3. Event Logging:

Full functionality support is available for the VMware ESXi operating system, but it works differently than the normal Intel® RAID Web Console 2 framework mode. The event logging feature for the Intel® RAID Web Console 2 Client connected to a VMware ESXi system behaves as follows:

- There is no support for retrieving initial logs (the events that occurred before a client logs in). Only those events that occur after a client logs in appear in the event logger dialog.
- System log does not display.
- The "Save log" feature is not supported; however, the "Save Log as Text" is still supported.
- The "View Log" option allows you to view the logs saved in a text file on the event logger dialog.
- The event descriptions might not be identical to a normal Intel® RAID Web Console 2 Client because the descriptions come from the firmware through the provider.
- There is no filtering of events, unlike Monitor Service.
- Refreshing of the Intel® RAID Web Console 2 GUI after any updates on the firmware is slower for a client connected to VMWare ESXi hosts, compared to one connected to Windows/Linux/Solaris hosts.

4. Remote discovery and heartbeat mechanism:

For networks that do not have DNS configured, the "hosts" file in the machine on which the Intel® RAID Web Console 2 is installed must be edited as follows:

- Add an entry to map the VMWare host's IP address with the hostname. This is for the discovery to happen correctly. In the absence of this entry, the VMWare host would be discovered as 0.0.0.0.
- Add an entry to map its own IP address (not the loop back address) with the Hostname. This is to ensure that the Alert Event Notifications (AENs) are delivered correctly.

For networks that has DNS configured, the "hosts" file in the machine on which Intel® RAID Web Console 2 is installed must be edited as follows:

- When you do the initial configurations for the VMWare host, provide the correct DNS server IP address.
- In the hosts file of the machine on which Intel® RAID Web Console 2 is installed, add an entry to map its own IP address (not the loop back address) with the Hostname. This is to ensure that the Asynchronous Event Notifications (AENs) are delivered correctly.

5. The VMWare hosts are discovered only when the Framework service starts on the host where Intel® RAID Web Console 2 is installed.

6. The VMWare ESX3i does not support the heartbeat mechanism to let Intel® RAID Web Console 2 know whether VMWare ESX3i is still connected. When the connection to the remote VMWare ESX3i is lost, Intel® RAID Web Console 2 does not indicate this. The only option is to rediscover by restarting the Intel® RAID Web Console 2 framework.

7. This is supported only on a full installation of Intel[®] RAID Web Console 2; standalone, client-only, and server-only modes do not support VMWare ESX3i management.
8. Supported on following guest operating systems:
 - Windows Server 2003* and Windows Server 2008*
 - Linux RHEL 4 and 5
9. The following describes the status of components related to VMWare ESX3i:
 - Intel[®] RAID Web Console 2 client GUI is supported.
 - There is no support for Monitor Configurator; you cannot configure the severity of the AENs.
 - There is no pop-up service support.
 - There is no email and system log support.
 - Monitor service support is not available.
10. For Red Hat Enterprise Linux 5, you must create the following symbolic links:

Note: This step is not required for Intel[®] RAID Web Console 2 version 2.90-02 or later.

- `cd /usr/lib` on RHEL 5
- Search for `libcrypto`, `libssl` and `libsysfs` libraries as follows:


```
ls -lrt libcrypto*, ls -lrt libssl*, ls -lrt libsysfs*
```
- If the files `libcrypto.so.4`, `libssl.so.4`, and `libsysfs.so.1` are missing, manually create sym links as follows:


```
ln -s libcrypto.so libcrypto.so.4
ln -s libssl.so libssl.so.4
ln -s libsysfs.so libsysfs.so.1
```
- If the `.so` files are not present in the `/usr/lib` directory, create a link with the existing version of the library. For example, if `libcrypto.so.6` is present and `libcrypto.so` is not, create the link as follows:


```
ln -s libcrypto.so.6 libcrypto.so.4
```

Running Intel[®] RAID Web Console 2 on VMWare ESX 3.5i U2

If you are using VMWare ESX 3.5i U2, perform the following steps to make Intel[®] RAID Web Console 2 work:

1. Open the maintenance console/shell in ESX3.
 - a. Press ALT+F1.
A shell without any prompt appears.
 - b. Type `unsupported` (all lowercase) and press ENTER.
Typed text is not prompted back.
 - c. Enter your password when prompted.
There is no password by default for the shell. If you have set any password from the "yellow" screen (DCUI), use that password.
You are prompted (`#`) next.
2. Enable `ssh` for remote copy.

- a. Type the following command.

```
vi /etc/inetd.conf
```
 - b. Search for `ssh` in the file.
By default, the line that contains `ssh` has comments.
 - c. Remove the comment by deleting the symbol `#` in front of the line.
 - d. Save the file and exit.
3. Restart the `inetd` daemon for the changes to take effect.
 - a. Type the following command to get the pid for `inetd`:

```
ps | grep inetd
```
 - b. Type the following command to kill the `inetd` process:

```
Kill -9 <inetd pid>
```
 - c. Type the following command to restart the `inetd` daemon:

```
#inetd
```
 4. Type the following command to use `scp` to copy `storelib` from a remote machine to the following path:

```
/lib dir scp <user@ip:path to  
storelib>/libstorelib.so.2.53 /lib/libstorelib.so
```
 5. Restart `SFCB` and check its status.
 - a. Type the following command to restart `SFCB`:

```
/etc/init.d/sfcbd restart
```
 - b. Type the following command to check the status of `SFCB`:

```
/etc/init.d/sfcbd status
```

Note: *The updated Storelib library in the /lib directory does not persist across reboots. Each time you restart the VMWare host, you have to follow this procedure to replace the Storelib library.*

Starting the Intel® RAID Web Console 2

Complete the following steps to start the Intel® RAID Web Console 2 (Different Operating Systems may have slightly different directories):

- **Microsoft Windows***: Select **Start | All Programs | RAID Web Console 2 | StartupUI**, or double-click the Intel® RAID Web Console 2 icon on the desktop.
- **Red Hat* Enterprise Linux 3 U6**: Select **Start | System Tools | RAID Web Console 2 StartupUI**.
- **SuSE* Linux Enterprise Service 9 SP1**: Select **Start | System | More Programs | RAID Web Console 2 StartupUI**.

Intel® RAID Web Console 2 Screens

This section describes the main Intel® RAID Web Console 2 screens. When you start the Intel® RAID Web Console 2, the Select Server window is displayed.

Note: Depending on the revision of Intel® RAID Web Console 2 software installed, the following screen may vary slightly.

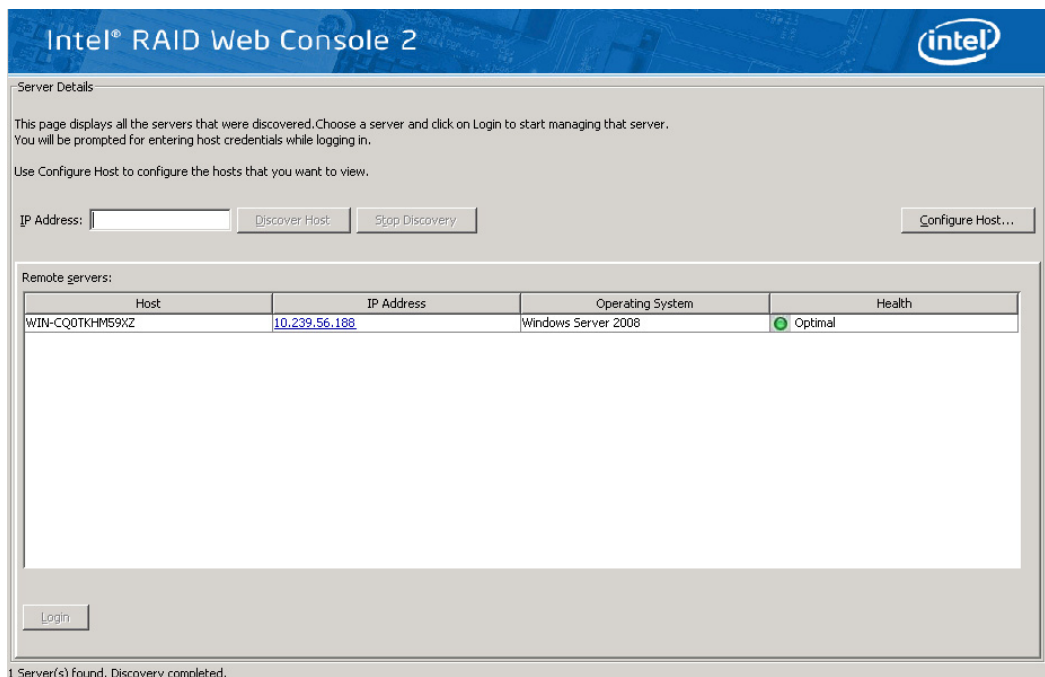


Figure 42. Intel® RAID Web Console 2 - Select Server Screen

Note: To access systems on a different subnet, type in the box in middle-left side of the screen the IP address of a system in the desired subnet where Intel® RAID Web Console 2 is running and click Discover Host. If you check the Connect to remote framework box, you can also access a standalone server running Intel® RAID Web Console 2, if it has a network connection.

Color Coding: If the circle in the server icon is yellow instead of green, it means that the system is running in a degraded state (for example, a disk drive used in a virtual drive has failed). If the circle is red, the storage configuration in the system has failed.

To log in to a system, follow these steps:

1. Double-click the icon of the system from "Remote servers" list that you want to access. The Server Login window appears.



Figure 43. Intel® RAID Web Console 2 - Login Screen

2. Select an access mode from the drop-down menu.
 - Select **Full Access** if you need to view the current system configuration and change the configuration.
 - Select **View Only** if you only need to view the system configuration.
3. Enter your user name and password and click **Login**.

If your user name and password are correct for the login mode you have chosen, the main screen displays.

Note: Intel® RAID Web Console 2 uses the username and password of your server operating system. In Microsoft Windows*, you can use the username Administrator and appropriate password to log on to Intel® RAID Web Console 2. In Linux, you can use the username root and appropriate password to log on to Intel® RAID Web Console 2. If you do not know the password, contact the server system administrator to obtain the password.




Figure 44. Intel® RAID Web console 2 dashboard

The information displayed on this Dashboard screen provides an overview of the RAID ecosystem on the selected server system. Some frequently used functions are listed in the Dashboard tab, such as Create Virtual Drive etc.




Figure 45. Intel® RAID Web Console 2 – Main Screen

The information on the screen in [Figure 45](#) shows the connected devices such as RAID controllers, physical devices, battery and so on.


The following subsections describe this screen in more detail.


Physical/Virtual View Panel

Depending on the tab selected, the left panel displays either the Physical View or the Virtual View of the system and the devices in it.

Note: The word ‘Virtual’ maybe displayed as ‘Logical’ depending on different versions of the utility.

- The Physical View shows the hierarchy of physical devices in the system. At the top of the hierarchy is the system itself. Controllers are installed in the system, and each controller has one or more ports. Disk drives and other physical devices are attached to the ports.
- The Virtual View shows the hierarchy of systems, controllers, virtual drives, and arrays and disk groups that are defined for the system.

Small icons represent the servers, controllers, and other devices. A red circle to the right of an icon indicates that the device has failed. For example, this icon indicates that a disk drive has failed: 

A yellow circle to the right of an icon indicates that a device is running in a degraded state. For example, this icon indicates that a virtual drive is running in a degraded state due to the failure of a disk drive: 

Properties/Operations/Graphical View Panel

Depending on the kind of device selected in the left panel and your login mode (full-access or view-only), the right panel has either two or three tabs.

The Properties tab displays information about the selected device.

Note: Depending on different Intel(R) RAID Web Console 2 versions, Operations and Graphical tabs may not exist, or only exist in old version utility. The screen displays in following sections show all these view panels for user reference, in case different versions of Intel® RAID Web Console 2 is installed on server system.

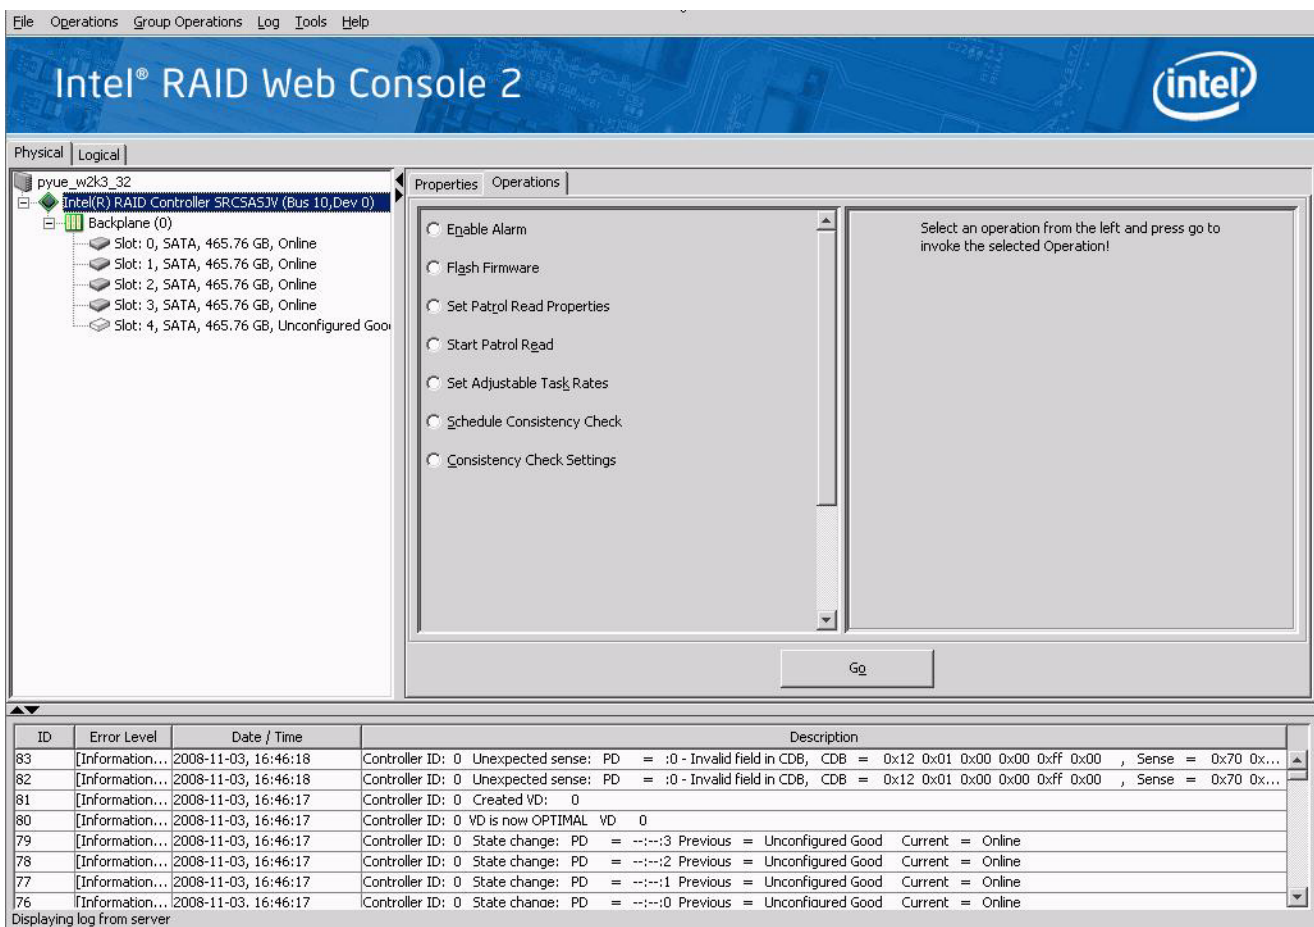


Figure 46. Intel® RAID Web Console 2 - Operations Tab

The Operations tab lists the operations that can be performed on the device selected in the left panel. This tab is available only when you are logged in to Intel® RAID Web Console 2 in Full-access mode. Options available for controllers include: enabling or silencing the alarm, flashing the firmware, and so on. Some types of devices, such as arrays, disk groups, and ports, do not have operations associated with them.

If a physical or virtual drive is selected in the left panel, the Graphical tab is available in the right panel. In the Graphical View, the device's storage is color coded to show used capacity, unused capacity, and so on.

Note: This Graphical View feature may or may not be displayed in the utility depending on the version of the utility.




Figure 47. Intel® RAID Web Console 2 - Graphical Tab (Optional feature)

Event Log Panel

The lower part of the screen displays the event log entries for the system. New event log entries display during the session. Each entry has a time and date stamp, an Error Level that indicates the severity of the event, and a brief description of the event. For information about the event log entries, see [Appendix B: “Events and Messages” on page 173](#)

Menu Bar/Manage Menu

The File menu includes the Exit option to close the Intel® RAID Web Console 2. It also includes a Refresh option to update the screen with the latest RAID configuration information.

Other options may include: Server option to discover and enter other server system's RAID configuration; Check Consistency; Initialize to initialize a RAID; Show progress.

Menu Bar/Go To Menu

The Go To menu has sub-menus like Controller, Drive Group, Physical Drive, Virtual Drive, BBU, etc.

Each of above sub-menu may have some or all items greyed out, unless controller, physical, virtual drive, drive group, BBU or other storage object is selected. For example, the Enable Alarm and Silence Alarm options are available only when a controller is selected. The options also vary depending on the current state of the selected object. For example, if an offline physical drive is selected, the Make Drive Online option is displayed in the menu.

There is a general tip here that, most functions in the Go To menu are equal to right-clicking the storage object selected in Physical or Logical view tab. For example, selecting Go To | Controller | Create Virtual Drive, is equal to selecting the controller in Physical (or Logical) tab, right-clicking it and choosing Create Virtual Drive.

File Menu/Log Menu

The Log menu includes options for saving, clearing, and loading the message log.

File Menu/Tool Menu

The Tools menu includes options for configuring alerts.

File Menu/Help Menu

The Help menu provides access to the online help file and Intel® RAID Web Console 2 version information.

Drive Configuration Tasks

You can use Intel® RAID Web Console 2 to perform the following configuration tasks:

- [“Creating a Virtual Drive Using Simple Configuration” on page 103](#)
- [“Creating a Spanned Disk Group” on page 114](#)
- [“Creating Hot Spares” on page 115](#)
- [“Setting Adjustable Task Rates” on page 117](#)
- [“Adding a Drive to a Virtual Disk” on page 118](#)

- “Removing a Drive or Changing the RAID Level of a Virtual Disk” on page 121
- “A Reconstruct operation begins on the virtual disk. You can monitor the progress of the reconstruction in the Group Show Progress window. Select Manage -> Show Progress.” on page 122
- “Changing Virtual Disk Properties” on page 122
- “Deleting a Virtual Disk” on page 123
- “Managing Configurations” on page 124

Create Virtual Drive

Use the Create Virtual Drive to create disk groups and virtual disks. The Create Virtual Drive can create simple configurations automatically. For more complex configurations, the Create Virtual Drive allows you to customize the configuration parameters according to your needs.

The Modify Drive Group allows you to easily change RAID levels, or to expand or reduce the capacity of existing virtual drives.

Note: *Unless you are logged on to a system with Administrator privileges, you cannot create or modify a storage configuration.*

Creating a Virtual Drive Using Simple Configuration

Simple configuration is the quickest and easiest way to create a new storage configuration. When you select simple configuration mode, the system creates the best configuration possible using the available drives

Note: *You cannot create spanned drives using the simple configuration procedure. To create spanned drives, use the advanced configuration procedure described in “Creating a Virtual Drive Using Advanced Configuration” on page 107*

Follow these steps to create a new storage configuration in simple configuration mode.

1. Perform either of the following steps:
 - Right click on the controller node in the device tree in the left frame of the Intel® RAID Web Console 2 window and select Create Virtual Drive
 - Select the controller node and select Go To -> Controller -> Create Virtual Drive in the menu bar, as shown in [Figure 48](#)

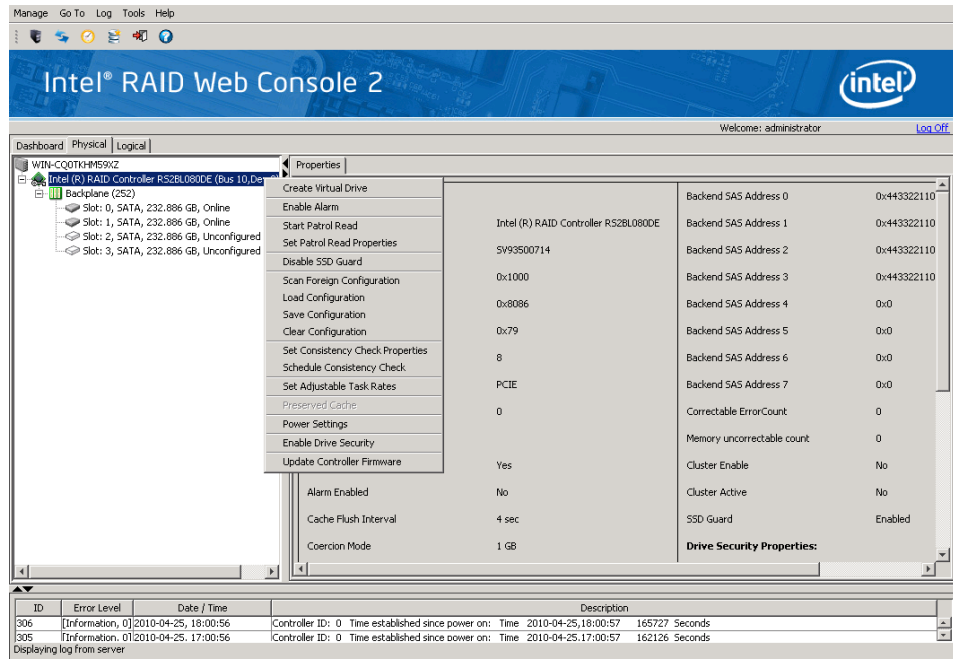


Figure 48. Virtual Drive Creation Menu

The dialog box for the configuration mode (simple or advanced) appears, as shown in Figure 49

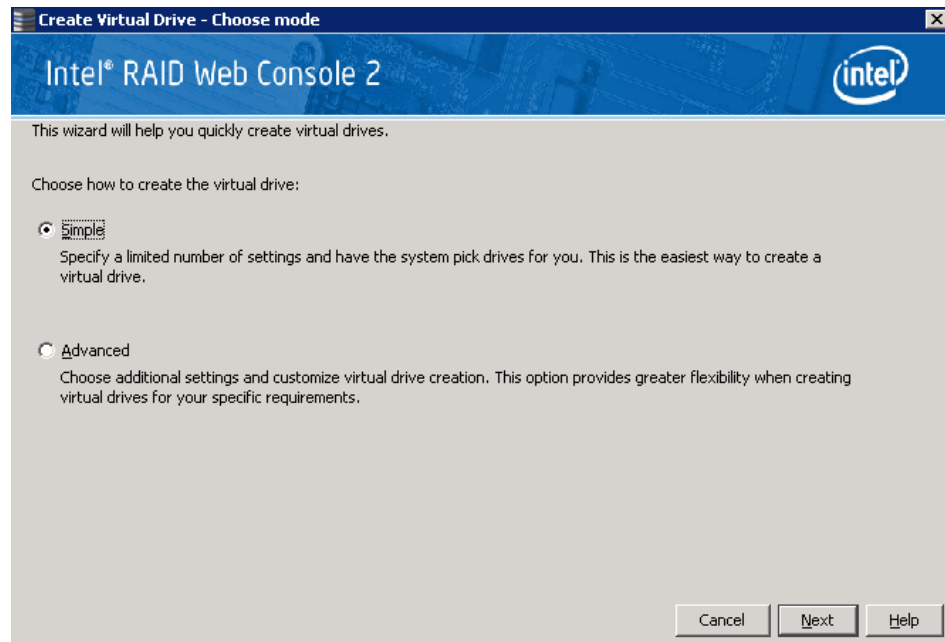


Figure 49. Virtual Drive Creation Mode

2. Click **Simple** and press **Next**.

The Create Virtual Drive screen appears, as shown in [Figure 50](#).




Figure 50. Create Virtual Drive Screen

3. Select the RAID level desired for the virtual drive.

When you use simple configuration, the RAID controller supports RAID levels 1, 5, and 6. In addition, it supports independent drives (configured as RAID 0). The screen text gives a brief description of the RAID level you select. The RAID levels you can choose depend on the number of drives available. To learn more about RAID levels, see Chapter-2 “[RAID Levels](#)” on page 9

4. Click the box next to **Assign a hot spare** if you want to assign a hot spare drive to the virtual drive. Hot spares are drives that are available to replace failed drives automatically in a redundant virtual drive (RAID 1, RAID 5, or RAID 6).

Note: *In the simple configuration procedure, you can assign dedicated hot spares to a maximum of 16 arrays at one time. This is because the dedicated hot spares can support only up to 16 drive groups. If you try to create more than 16 drive groups at one time, dedicated hot spares will not be assigned to drive groups beyond the first 16.*

To create more than 16 drive groups with hot spares, you need at least 35 drives of the same capacity.

5. Select the number of virtual drives that you want to create.
6. Select the capacity for the virtual drives.
Each virtual drive has the same capacity.
7. Click **Next**.

The **Create Virtual Drive - Summary** window appears, as shown in [Figure 51](#). This window shows the selections you made for simple configuration.




Figure 51. Create Virtual Drive - Summary Window

8. Click **Back** to return to the previous screen to change any selections or click **Finish** to accept and complete the configuration.

The new storage configuration will be created and initialized.

Note: *If you create a large configuration using drives that are in powersave mode, it could take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a box appears to identify the drive or drives.*

After the configuration is completed, a dialog box may or may not notify (depending on different utility version) you that the virtual drives were created successfully, as shown in [Figure 52](#). If more drive capacity exists, the dialog box asks whether you want to create more virtual drives. If no more drive capacity exists, you are prompted to close the configuration session.



Figure 52. Option to Create Additional Virtual Drives

9. Select **Yes** or **No** to indicate whether you want to create additional virtual drives.

If you select **Yes**, the system takes you to the Create Virtual Drive screen, as shown in [Figure 50](#). If you select **No**, the utility asks whether you want to close the wizard, as shown in [Figure 53](#).

10. If you selected **No** in step 9, select **Yes** or **No** to indicate whether to close the wizard.



Figure 53. Option to Close the Configuration Wizard

If you select **Yes**, the configuration procedure closes. If you select **No**, the dialog box closes and you remain on the same page.

Creating a Virtual Drive Using Advanced Configuration

The advanced configuration procedure provides an easy way to create a new storage configuration. Advanced configuration gives you greater flexibility than simple configuration because you can select the drives and the virtual drive parameters when you create a virtual drive. In addition, you can use the advanced configuration procedure to create spanned drive groups.

Follow these steps to create a new storage configuration in the advanced configuration mode. In this example, we will create a spanned drive group.

1. Perform either of the following steps:
 - Right click on the controller node in the device tree in the left frame of the Intel[®] RAID Web Console 2 window and select Create Virtual Drive
 - Select the controller node and select Go To -> Controller -> Create Virtual Drive in the menu bar, as shown in [Figure 54](#)

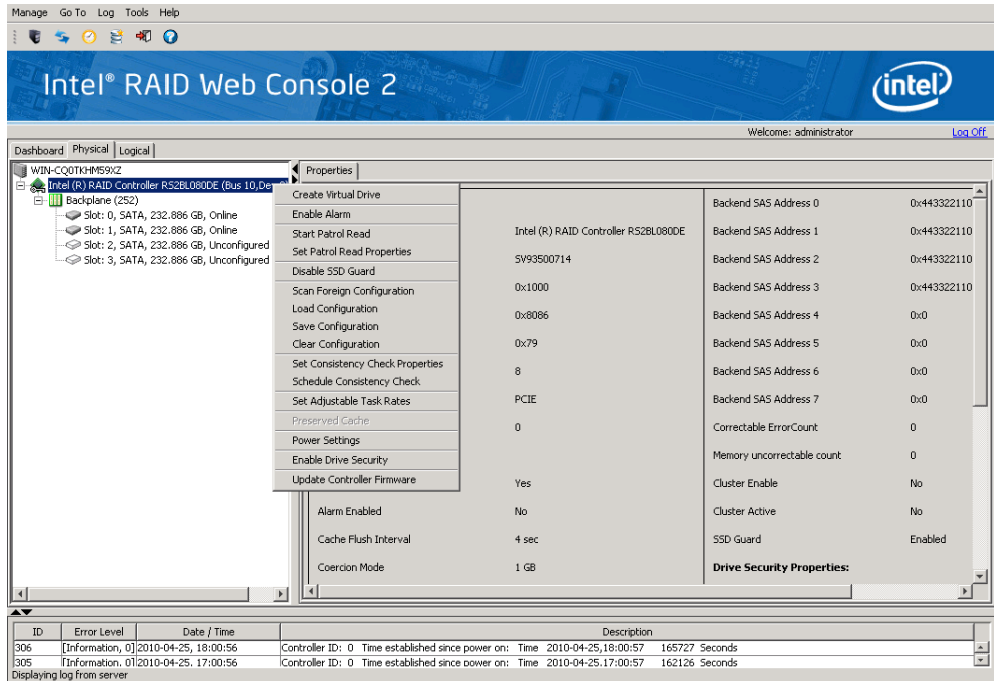


Figure 54. Virtual Drive Creation Menu

The dialog box shown in Figure 55 appears.




Figure 55. Virtual Drive Creation Mode

2. Click **Advanced** and press **Next**.

The Create Drive Group Settings screen appears, as shown in [Figure 56](#).




Figure 56. Create Drive Group Settings Screen

3. Select the following items on the Create Drive Group Settings screen:
 - a. Select the RAID level desired for the drive group from the drop-down menu. To make a spanned drive, select **RAID 10**, **RAID 50**, or **RAID 60** in the **RAID level** field.

Drive Group 0 and **Span 0** appear in the **Drive groups** field when you select RAID 10, 50, or 60.

The RAID controller supports RAID levels 1, 5, 6, 10, 50, and 60. In addition, it supports independent drives (configured as RAID 0 and RAID 00). The screen text gives a brief description of the RAID level you select. RAID levels you can choose depend on the number of drives available. To learn more about RAID levels, see Chapter 2, “[RAID Levels](#)” on page 9

- b. Select unconfigured drives from the list of drives and click **Add>** to add them to the drive group.

The selected drives appear under **Span 0** below **Drive Group 0**, as shown in [Figure 57](#)




Figure 57. Span 0 of Drive Group 0

- c. Click **Create Span** to create a second span in the drive group.
- d. Select *unconfigured* drives from the list of drives and click **Add >** to add them to the drive group.
- e. The selected drives appear under **Span 1** below **Drive Group 0**, as shown in Figure 58




Figure 58. Span 0 and Span 1 of Drive Group 0

- f. Click **Create Drive Group** to make a drive group with the spans.
- g. Click **Next** to complete this step.

The Virtual drive settings window appears, as shown in [Figure 59](#). The drive group and the default virtual drive settings appear. The options to update the virtual drive or remove the virtual drive are grayed out until you create the virtual drive.




Figure 59. Virtual Drive Settings Window

4. Change any virtual drive settings, if desired.
5. Click **Create Virtual Drive**.

The new virtual drive appears under the drive group, as shown in [Figure 60](#). The options **Update Virtual Drive** and **Remove Virtual Drive** are now available. **Update Virtual Drive** allows you to change the virtual drive settings and **Remove Virtual Drive** allows you to delete the virtual drive.

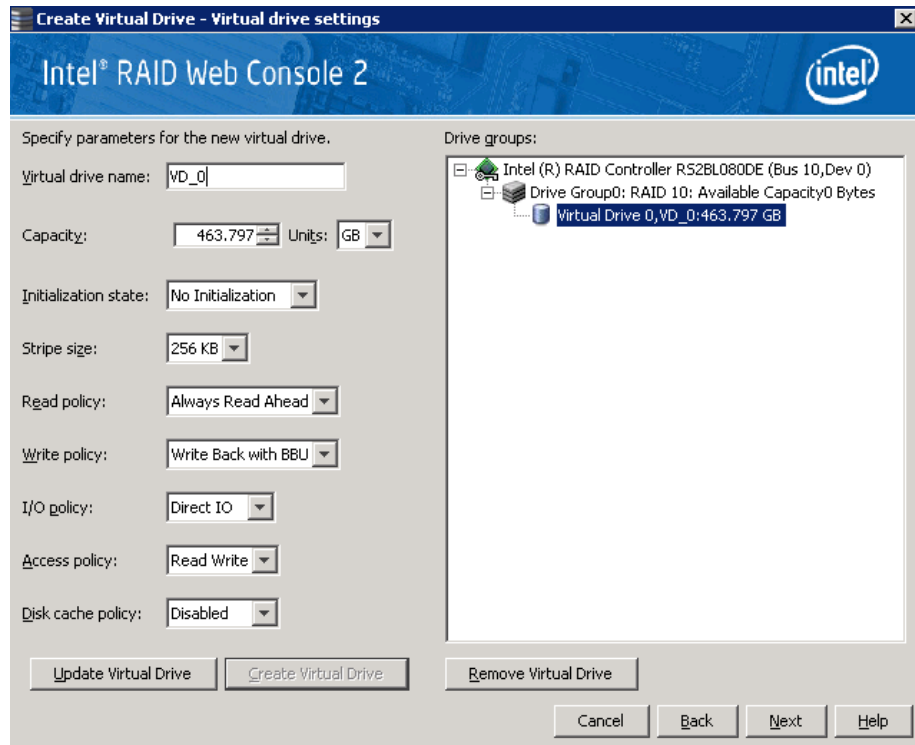


Figure 60. New Virtual Drive 0

6. Click Next.

The **Create Virtual Drive - Summary** window appears, as shown in Figure 61. This window shows the selections you made for advanced configuration.




Figure 61. Create Virtual Drive Summary Window

7. Click **Back** to return to the previous screen to change any selections or click **Finish** to accept and complete the configuration.

The new storage configuration will be created and initialized.

***Note:** If you create a large configuration using drives that are in powersave mode, it could take several minutes to spin up the drives. A progress bar appears as the*

drives spin up. If any of the selected unconfigured drives fail to spin up, a box appears to identify the drive or drives.

8. After the configuration is completed, a dialog box may or may not notify (depending on different utility version) you that the virtual drives were created successfully, as shown in [Figure 62](#). If more drive capacity exists, the dialog box asks whether you want to create more virtual drives. If no more drive capacity exists, you are prompted to close the configuration session.



Figure 62. Option to Create Additional Virtual Drives

9. Select **Yes** or **No** to indicate whether you want to create additional virtual drives.
If you select **Yes**, the system takes you to the **Create Virtual Drive** screen, as shown in [Figure 50](#). If you select **No**, the utility asks whether you want to close the wizard, as [Figure 63](#).
10. If you selected **No** in <hyperactive>step 8, select Yes or No to indicate whether you want to close the wizard.
If you select **Yes**, the configuration procedure closes. If you select **No**, the dialog box closes and you remain on the same page.



Figure 63. Option to Close the Configuration Wizard

Creating a Spanned Disk Group

Spanning allows you to configure multiple arrays as a single virtual disk. Spanned arrays provide additional levels of data redundancy and storage capacity. The Intel® RAID Web Console 2 supports three types of spanned arrays or disk groups:

- RAID 00 (multiple RAID 0 arrays or disk groups)
- RAID 10 (multiple RAID 1 arrays or disk groups)
- RAID 50 (multiple RAID 5 arrays or disk groups)
- RAID 60 (multiple RAID 6 arrays or disk groups)

For spanned disk group creation steps, refer to the section [Creating a Virtual Drive Using Advanced Configuration](#).

Creating Hot Spares

Hot spares are disk drives that are available to automatically replace failed drives in a virtual disk. There are two kinds of hot spares: 1) dedicated hot spares, which are available to one or more specified arrays or disk groups, and 2) global hot spares, which are available to any array or disk group defined on the controller.

To create a global hot spare, follow these steps:

1. In the left panel of the Intel[®] RAID Web Console 2 window, right-click the icon of disk drive that is not assigned to a storage configuration. If it is assigned to a storage configuration, a check mark is displayed on the disk drive icon.
2. In the Menu, click the **Go To -> Physical Drive** tab and select **Assign Global Hot Spare**, as shown in [Figure 64](#).

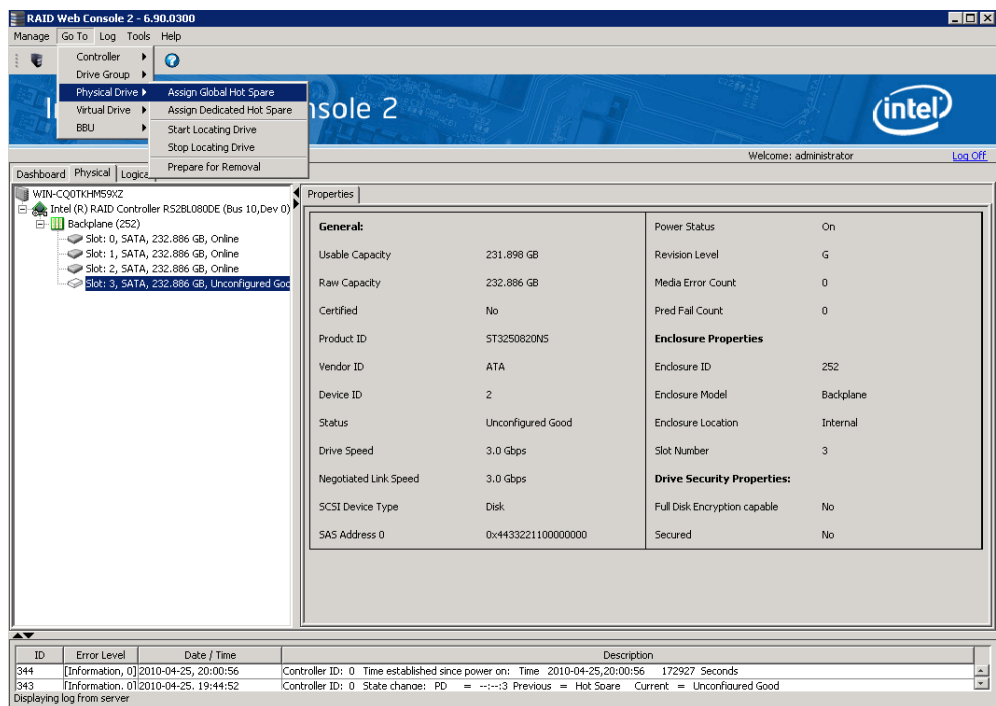


Figure 64. Assign Global Hotspare

You normally create a dedicated hot spare when you create a new configuration with the Manual Configuration option. To add a dedicated hot spare to an existing array or disk group, follow these steps:

1. In the left panel of the Intel[®] RAID Web Console 2 window, click the icon of a disk drive that is not assigned to a storage configuration. If it is assigned to a storage configuration, a check mark is displayed on the disk drive icon.
2. In the Menu, click the **Go To -> Physical Drive** tab and select **Assign Dedicated Hot Spare**, as shown in [Figure 65](#).




Figure 65. Assign Dedicated Hotspare

3. Select the array or disk group to which the hotspare will be dedicated from the list in the pop-up window.

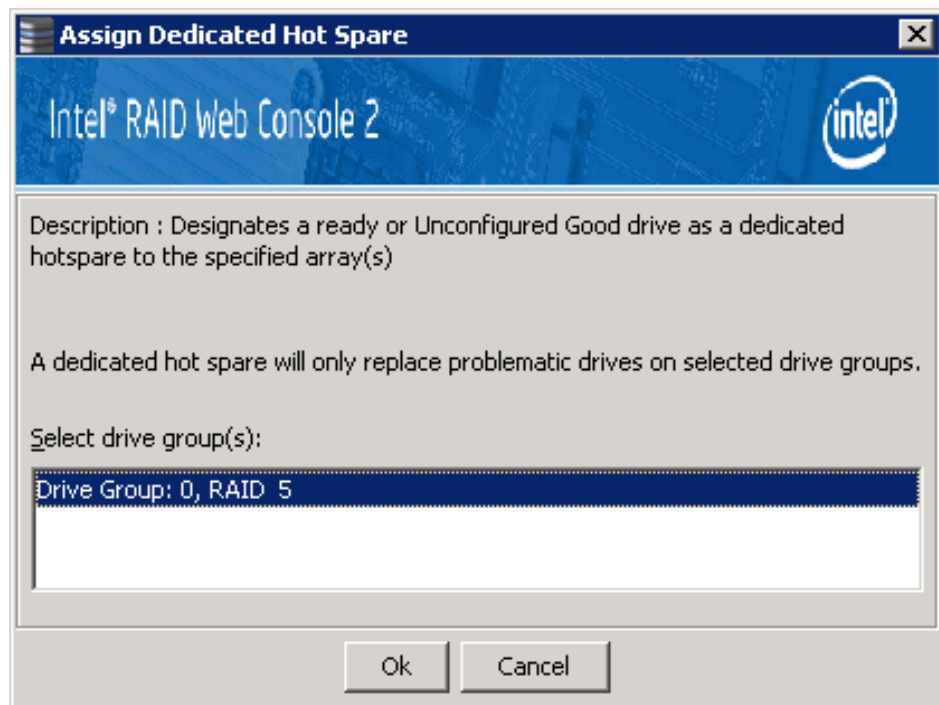


Figure 66. Select Hotspare Drive

4. Click **OK** to create the dedicated hot spare.

Setting Adjustable Task Rates

To change the Rebuild rate and other task rates for a controller, you must first log onto the system in All Access mode (versus View-Only mode). Complete the following steps to set the adjustable task rates:

1. Select a controller icon in the left panel, and select the **Physical View** tab.
2. In the Menu, select the **Go To -> Controller** tab and select **Set Adjustable Task Rates**.

The task rates appear in a new pop-up window, as shown in [Figure 67](#).




Figure 67. Set Adjustable Task Rates

3. Enter changes as needed to any of the task rates:
 - **Rebuild Rate.** Enter a number from 0 to 100 to control the rate at which a rebuild is performed on a disk drive when one is necessary. The higher the number, the faster the rebuild will occur. However, if a higher number is selected, the system I/O rate may be slower.
 - **Patrol Rate.** Enter a number from 0 to 100 to control the rate at which patrol reads are performed. Patrol read is a preventive procedure that monitors physical disks to find and resolve potential problem that might cause a disk failure. The higher the number, the faster the patrol read occurs. However, if a higher number is selected, the system I/O rate may be slower.

- **Background Initialization (BGI) Rate** Enter a number from 0 to 100 to control the rate at which virtual disks are initialized in the background. Background initialization makes the virtual disk immediately available for use, even while the initialization is occurring. The higher the number, the faster the initialization occurs. However, if a higher number is selected, the system I/O rate may be slower.
 - **Check Consistency Rate**. Enter a number from 0 to 100 to control the rate at which a consistency check is done. A consistency check scans the consistency data on a fault tolerant virtual disk to determine if the data is corrupted. The higher the number, the faster the consistency check is done. However, if a higher number is selected, the system I/O rate may be slower.
 - **Reconstruction Rate**. Enter a number from 0 to 100 to control the rate at which reconstruction of a virtual disk occurs. The higher the number, the faster the reconstruction will occur. However, if a higher number is selected, the system I/O rate may be slower.
4. Click **OK** to accept the new task rates.
 5. When the warning message is displayed, click **OK** to confirm that you want to change the task rates.

***Note:** The Controller Go To tab also has options for disabling or silencing the alarm on the controller. Ordinarily, you should leave the alarm enabled so it can warn you of abnormal conditions on the controller. If the alarm is malfunctioning, you may need to silence it.*

Adding a Drive to a Virtual Disk

You can use Intel® RAID Web Console 2 to increase the capacity of an existing virtual disk by adding physical disk drives to it. In order to do this, you must be logged on to the system in All Access mode.

***Warning:** Before you change a virtual disk's capacity, back up the data on the virtual disk.*

To increase the capacity of a virtual disk, follow these steps:

1. User Logical View and select a Drive Group icon in the left panel and then select **Go To -> Drive Group -> Modify Drive Group**.




Figure 68. Starting Modify Drive Group

The Modify Drive Group window is displayed, as shown in Figure 69.




Figure 69. Select RAID level to migrate

2. Click **Add Drive**.

Note: Depending on current RAID level, the targeted migrated RAID level and the available Unconfigured Good drive numbers, pop-up window may appear to allow selection among Add drives, Remove drives, and Migrate RAID level.

The following screen appears:




Figure 70. Selecting Drives to Add

3. In the panel, check the disk drives that you want to add to the virtual disk.
4. When you are finished adding disk drives, click **Next**. The next screen displays, which allows you to review the summary and go back if you need to make corrections as shown in Figure 71.

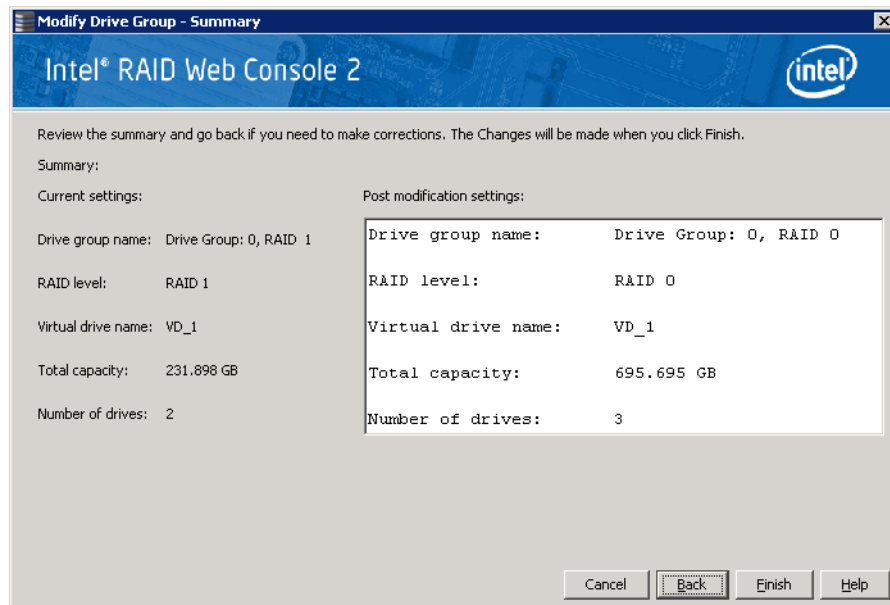


Figure 71. Changing RAID Level

5. When everything is acceptable, click **Finish** to accept the new configuration.
A Reconstruct operation begins on the virtual disk. You can monitor the progress of the reconstruction in the Group Show Progress window. Select **Manage** -> **Show Progress**.

Removing a Drive or Changing the RAID Level of a Virtual Disk

You can use Intel® RAID Web Console 2 to remove a physical disk drive from a virtual disk. In order to do this, you must be logged on to the system in All Access mode.

Warning: Be sure to back up the data on the virtual disk before you change its capacity.

1. Select a Drive Group icon and start the Modify Drive Group, as described in the previous section.
2. Click **Remove Drive**.

Note: Depending on current RAID level, the targeted migrated RAID level and the available Unconfigured Good drive numbers, pop-up window may appear to allow selection among Add drives, Remove drives, and Migrate RAID level.

The following screen appears:

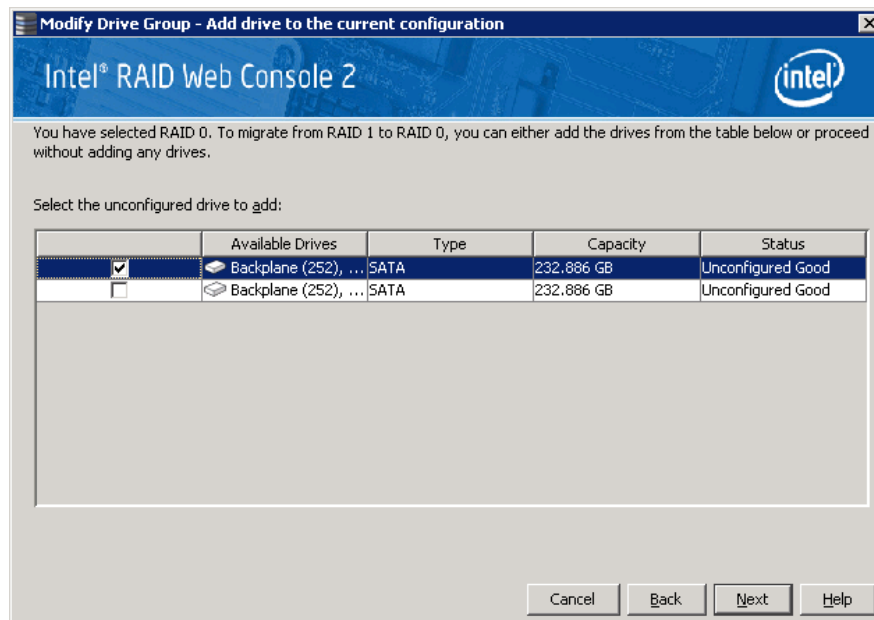


Figure 72. Selecting Drives to Remove

3. In the panel, check the disk drives that you want to remove from the virtual disk.
4. When you are finished removing disk drives, click **Next**. The next screen displays, which allows you to review the summary and go back if you need to make corrections as shown in [Figure 73](#).

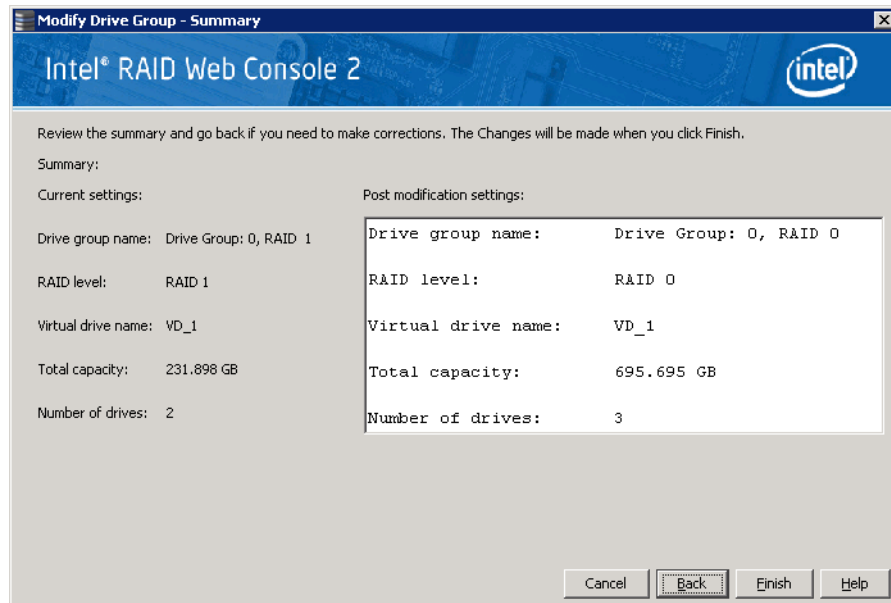


Figure 73. Changing RAID Level

5. When everything is acceptable, click **Finish** to accept the new configuration.

A Reconstruct operation begins on the virtual disk. You can monitor the progress of the reconstruction in the Group Show Progress window. Select **Manage -> Show Progress**.

Changing Virtual Disk Properties

You can change a virtual disk's Read Policy, Write Policy, and other properties after the disk is created. To do this, follow these steps:

1. Select a virtual disk (or drive) icon in the left panel and then select **Go To -> Virtual Drive -> Set Virtual Drive Properties**. Alternatively, click the **Properties** panel and then click **Set Virtual Disk Properties**. [Figure 74](#) shows the Set Virtual Disk Properties screen.

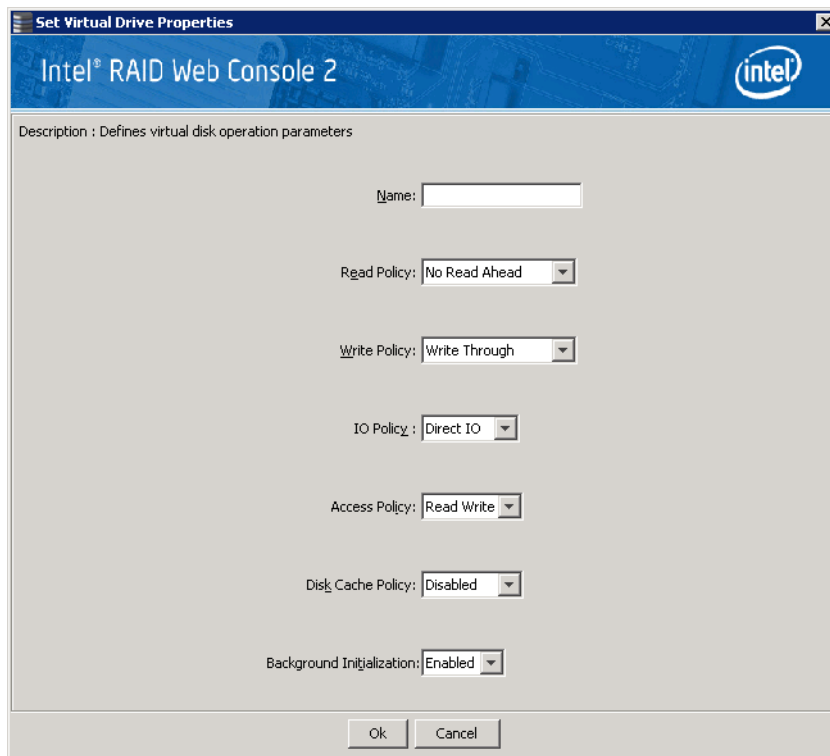


Figure 74. Set Virtual Disk Properties

2. As needed, change the properties in the pop-up window.
3. Click **OK** to accept the changes.

Deleting a Virtual Disk

Warning: Before you delete a virtual disk, back up the data on it.

You can delete one or more virtual disks. Before you can do this, you must be logged on to the system in All Access mode. To delete a virtual disk, follow these steps:

1. Back up all user data on the virtual disk you intend to delete.
2. In the left panel, select the **Virtual** (or **Logical**) tab and click the icon of the virtual disk (or drive) you want to delete.
3. In the menu, select the **Go To -> Virtual Drive** and click **Delete Virtual Drive**.
4. When the warning message displays, click **Yes** to confirm that you want to delete the virtual disk.

Managing Configurations

You can use Intel[®] RAID Web Console 2 to manage the configurations that you create. For example, you can save a storage configuration that you have defined on a controller and load this configuration from the disk to another controller, after first clearing the existing configuration from that controller. In order to do this, you must be logged on to the system in All Access mode. This section explains how to do these tasks.

Saving a Configuration to Disk

You can save an existing controller configuration to a file so you can apply it to another controller. To save a configuration file, follow these steps:

1. Select a controller icon in the left panel of the Intel[®] RAID Web Console 2 window.
2. Select **Go To -> Controller -> Save Configuration**.




Figure 75. Save Configuration to File

The Save dialog box displays as shown in [Figure 76](#).

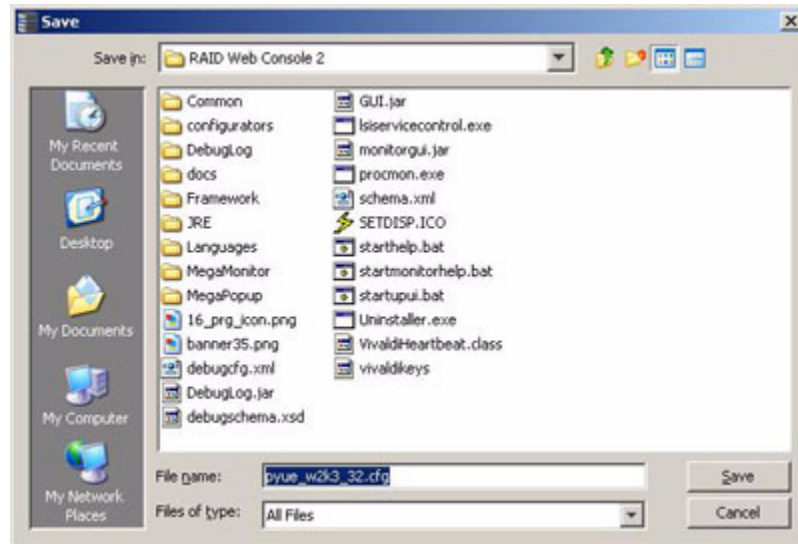


Figure 76. Save Configuration Dialog Box

3. In the dialog box, type a name for the configuration file.
4. Click **Save** to save the configuration file, or accept the default name: hostname.cfg

Clearing a Configuration from a Controller

If you want to manually create a new storage configuration on a controller, or load a configuration file on a controller, you must first clear the existing configuration.

Warning: Before you clear a configuration, make sure to save any data you want to keep! Clearing a configuration deletes all data from the disks in the existing configuration.

To clear a configuration from a controller, follow these steps:

1. Select a controller icon in the left panel of the Intel® RAID Web Console 2 window.
2. Select **Go To -> Controller -> Clear Configuration**.

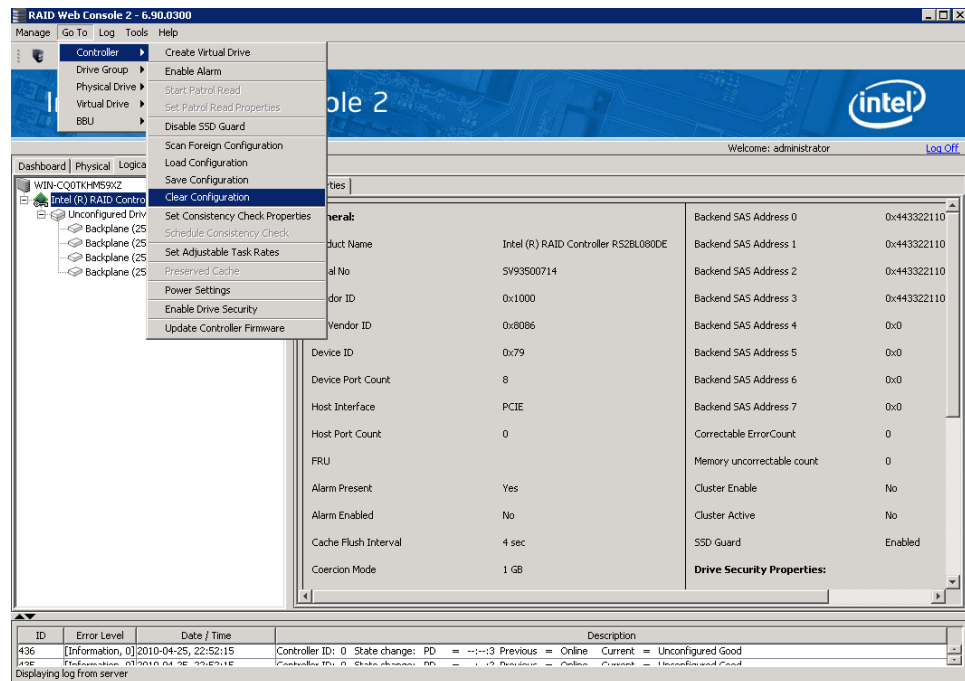


Figure 77. Clear Configuration

A Warning message displays.

3. Click **Yes** to clear the configuration or **No** to cancel the operation.

Adding a Configuration from a File

When you replace a controller or when you want to duplicate an existing storage configuration on a new controller, you can add a saved configuration to the controller.

Warning: When you add a saved configuration to a replacement controller, make sure that the number and size of the physical disks connected to the controller is exactly the same as it was when the configuration was saved.

To add a saved configuration, follow these steps:

1. Select a controller icon in the left panel of the Intel[®] RAID Web Console 2 window.
2. Select **Go To -> Controller -> Load Configuration**.




Figure 78. Add Saved Configuration

A Warning message displays.

3. Click **Yes**.
4. When the Open dialog box is displayed, select the configuration file and click **Open**.
5. View the configuration detail, then select **Apply**.
6. When prompted, confirm the new configuration.

Monitoring System Events and Devices

The Intel® RAID Web Console 2 enables you to monitor the status of disk drives, virtual disks, enclosures, and other devices. You can monitor the following:

- Monitoring System Events
- Monitoring Controllers
- Monitoring Disk Drives and Other Physical Devices
- Monitoring Virtual Disks
- Monitoring Enclosures
- Another method to use the BBU operations is to right-click the BBU icon to open the operations menu and select Start Learn Cycle.
- Monitoring Rebuilds and Other Processes

Monitoring System Events

Intel® RAID Web Console 2 monitors the activity and performance of all controllers in the system and the devices attached to them. When an “event” occurs—such as the completion of a consistency check or the removal of a physical drive—an event message displays in the log at the bottom of the Intel® RAID Web Console 2 screen, as shown in Figure 79. These event messages also display in the Microsoft Windows* application log (the Event Viewer).

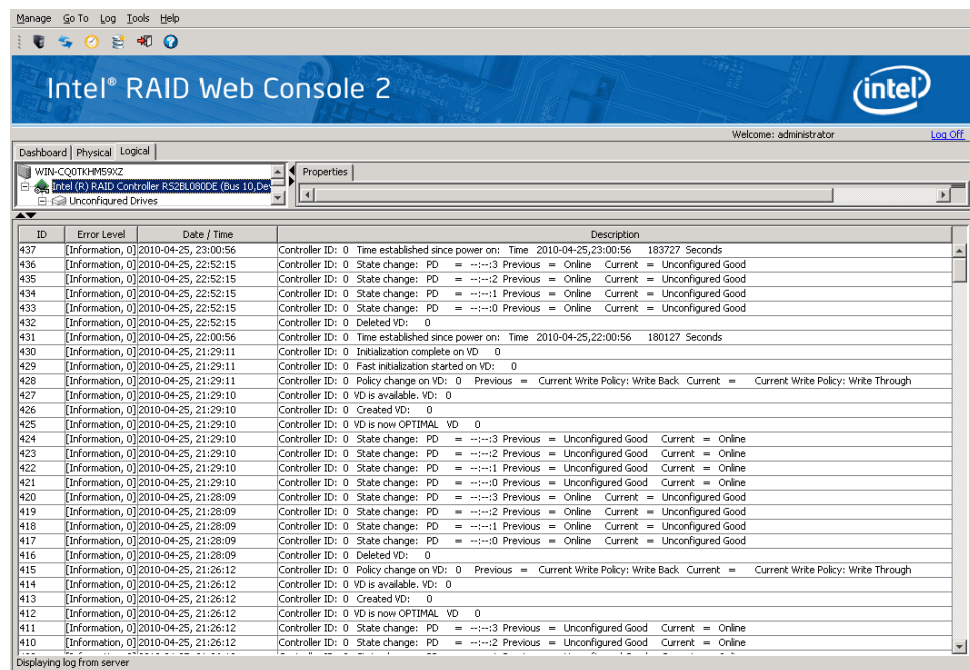


Figure 79. Event Information Window


Each event in the log includes an error level—Information, Warning, Critical, Fatal, or Dead—a date/time stamp, and a brief description. (For a list of all events, see [Appendix B: “Events and Messages” on page 173](#)) The status bar at the bottom of the screen indicates whether the log is a system log (as in [Figure 79](#)) or a log from a locally stored file. (This file could have been generated by a previous “Save Log” action.)

When a system log displays, the Log menu has four options:

- **Save Log:** Saves the current log to a .log file.
- **Save as Text:** Saves the current text to a .txt file.
- **Clear:** Clears the current log information, if you have full access (versus View-only access).
- **Load:** Enables you to load a local .log file.

When a local log is displayed, the status bar lists the name of the local log file, and the Log menu has an additional option, **Rollback to Current Log**, which enables you to retrieve the system log.

Monitoring Controllers

When Intel® RAID Web Console 2 is running, you can see the status of all controllers in the left panel. If the controller is operating normally, the controller icon looks like this: . If the controller has failed, a small red circle is displayed to the right of the icon.

To display complete controller information, click on a controller icon in the left panel. [Figure 80](#) shows the Controller Information window on the right panel.





Figure 80. Controller Information

Note the following:

- The *Rebuild rate*, *Patrol read rate*, *Reconstruction rate*, *Consistency check rate*, and *BGI rate* (background initialization) are all user-selectable. For more information, see “[Setting Adjustable Task Rates](#)” on page 117.
- The *BBU Present* field indicates whether a battery backup unit is installed.
- The *Alarm Present* and *Alarm Enabled* fields indicate whether the controller has an alarm to alert the user with an audible tone when there is an error or problem on the controller. There are options on the controller Properties tab for silencing or disabling the alarm. For more information, see “[Setting Adjustable Task Rates](#)” on page 117.

Monitoring Disk Drives and Other Physical Devices

When the Intel® RAID Web Console 2 is running, you can see the status of all physical disk drives and other physical devices in the left panel. If the physical drive is operating normally, the controller icon looks like this: . If the physical drive has failed, a small red circle is displayed to the right of the icon.

To display complete physical drive Information, click on a physical drive icon in the left panel. [Figure 81](#) shows the Properties panel for a physical drive.

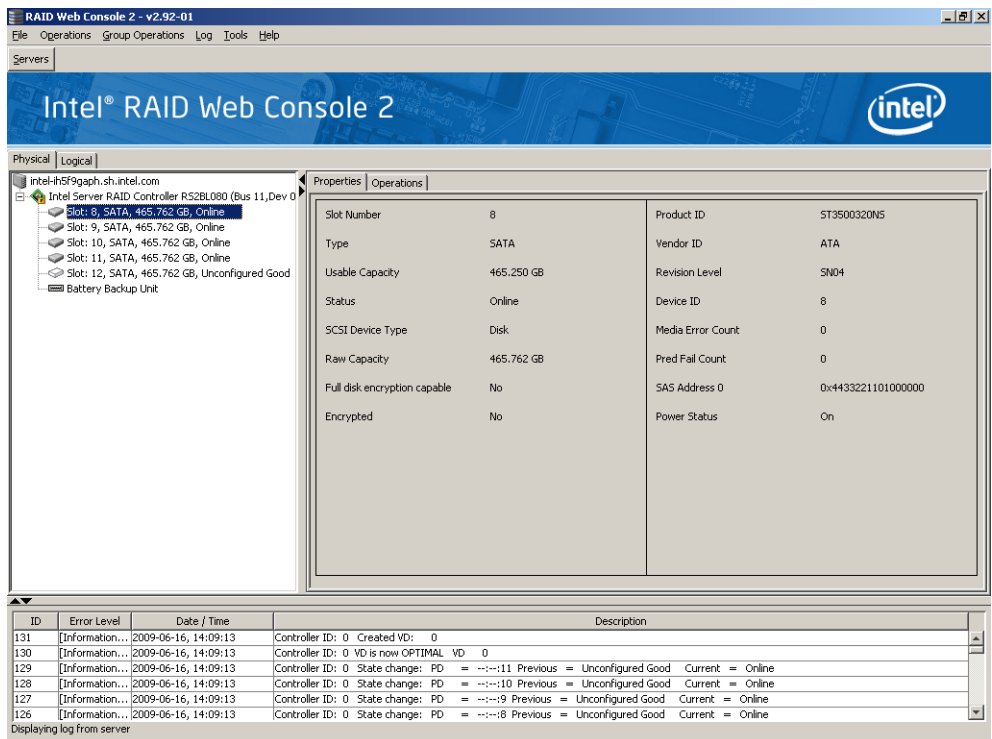


Figure 81. Physical Drive Information

There are no user-selectable properties for physical devices. Icons for other physical devices such as CD-ROM drives and DAT drives may also display in the left panel.

If the physical drives are in a disk enclosure, you can identify which physical drive is represented by a disk icon on the left. To do this, follow these steps:

1. Click the physical disk icon in the left panel.

2. Click the **Go To -> Physical Drive** tab.
3. Click **Start Locating Drive** (see [Figure 82](#)).




Figure 82. Locating a Physical Drive

The LED on the physical disk drive in the enclosure starts blinking to show its location.

Note: LEDs on global or dedicated hot-spare drives do not blink.

4. To stop the disk drive light from blinking, click **Stop Locating Drive**.

Running a Patrol Read

A Patrol Read periodically verifies all sectors of physical disks that are connected to a controller, including the system reserved area in the RAID configured drives. Patrol Read works for all RAID levels and for all hot-spare drives. A patrol read is initiated only when the controller is idle for a defined period and has no other background activities.

To enable and configure Patrol Read, follow these steps:

1. Click a controller icon in the left panel.
2. Select the **Go To -> Controller** tab and click **Set Patrol Read Properties** to display a pop-up window, as shown in [Figure 83](#).






Figure 83. Patrol Read Configuration

3. Select an Operation Mode for patrol read. The options are:
 - **Automatic:** Patrol Read runs automatically at the time interval you specify on this screen.
 - **Manual:** Patrol Read runs only when you manually start it by selecting **Start Patrol Read** from the controller options screen (see [Figure 83](#)).
 - **Disabled:** Patrol Read does not run at all.
4. (Optional) Specify a maximum count of physical drives to include in the patrol read. The default number is 255; if you want, you can specify a lower number.
5. (Optional) Select virtual disks on this controller to exclude from the Patrol Read. The existing virtual disks are listed in the gray box. To exclude a virtual disk, select the check box next to it.
6. Enter the Interval at which the Patrol Read will run.
7. If **Run Patrol Read non-stop** check box is selected, Patrol Read will run continuously.
8. Click **OK** to enable these Patrol Read options.

Monitoring Virtual Disks

When the Intel® RAID Web Console 2 is running, you can see the status of all virtual disks. If a virtual disk is operating normally, the icon looks like this: . If the virtual disk is running in degraded mode—for example, if a physical disk has failed—a small yellow circle is displayed to the right of the icon, which looks like this: .

When the Virtual tab is selected, the panel on the left shows which physical disks are used by each virtual disk. In [Figure 84](#) you can see that the virtual disk uses physical disks 1, 2, and 3. Multiple virtual disks can use the same physical disk.

To display complete virtual disk information, click the **Virtual** tab in the left panel, click on a virtual disk icon in the left panel. [Figure 84](#) shows the Properties tab for a virtual disk in the right panel.

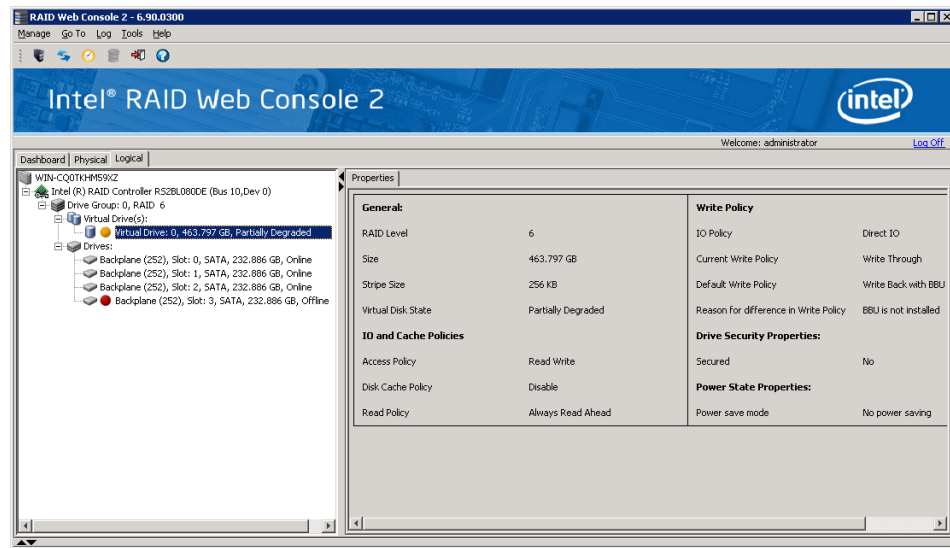


Figure 84. Virtual Drive Properties

The RAID level, stripe size, and access policy of the virtual disk are set when it is configured.

Note: You can change the Read Policy, Write Policy, and other virtual disk properties by selecting **Go To -> Virtual Drive -> Set Virtual Drive Properties**.

Monitoring Enclosures

When the Intel® RAID Web Console 2 is running, you can see the status of all enclosures that are operating normally. Information about the enclosure is displayed in the right panel. [Figure 85](#) shows it.

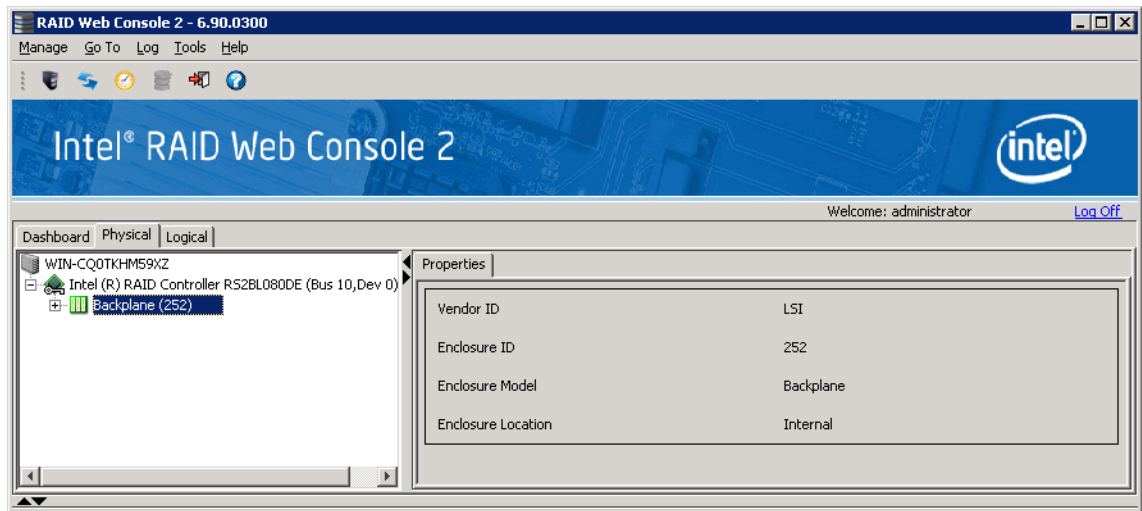


Figure 85. Enclosure Information

The display in the **Dashboard** tab also includes enclosure information in the **Properties** part.

Monitoring Battery Backup Units

When Intel® RAID Web Console 2 software is running, you can monitor the status of all of the BBUs connected to controllers in the server. If a BBU is operating normally, the icon looks like this: . If it has failed, a red dot appears next to the icon.

To show the properties for a BBU, perform the following steps:

1. Click the **Physical** tab on the main menu to open the physical view.
2. Select the BBU icon in the left panel.

The BBU properties, such as the battery type, temperature, and voltage, appear in the right panel, as shown in [Figure 86](#).




Figure 86. Battery Backup Unit Information

The BBU properties include the following:

- The number of times the BBU has been recharged (Cycle Count)
- The full capacity of the BBU, plus the percentage of its current state of charge, and the estimated time until it will be depleted
- The current BBU temperature, voltage, current, and remaining capacity
- If the battery is charging, the estimated time until it is fully charged

Battery Learn Cycle

Learn Cycle is a battery calibration operation performed by the controller periodically to determine the condition of the battery. You can start battery learn cycles manually or automatically. To choose automatic battery learn cycles, enable automatic learn cycles. To choose manual battery learn cycles, disable automatic learn cycles.

If you enable automatic learn cycles, you can delay the start of the learn cycles for up to 168 hours (7 days). If you disable automatic learn cycles, you can start the learn cycles manually, and you can choose to receive a reminder to start a manual learn cycle.

Setting Learn Cycle Properties

To set the learn cycle properties, perform the following steps:

1. Click the **Physical** tab to open the physical view.
2. Select the BBU icon in the left panel.

3. Click the **Go To -> BBU** tab.

The BBU operations appear, as shown in [Figure 87](#) (Depending on different utility version, the BBU operation also may appear in a new pop-up window).

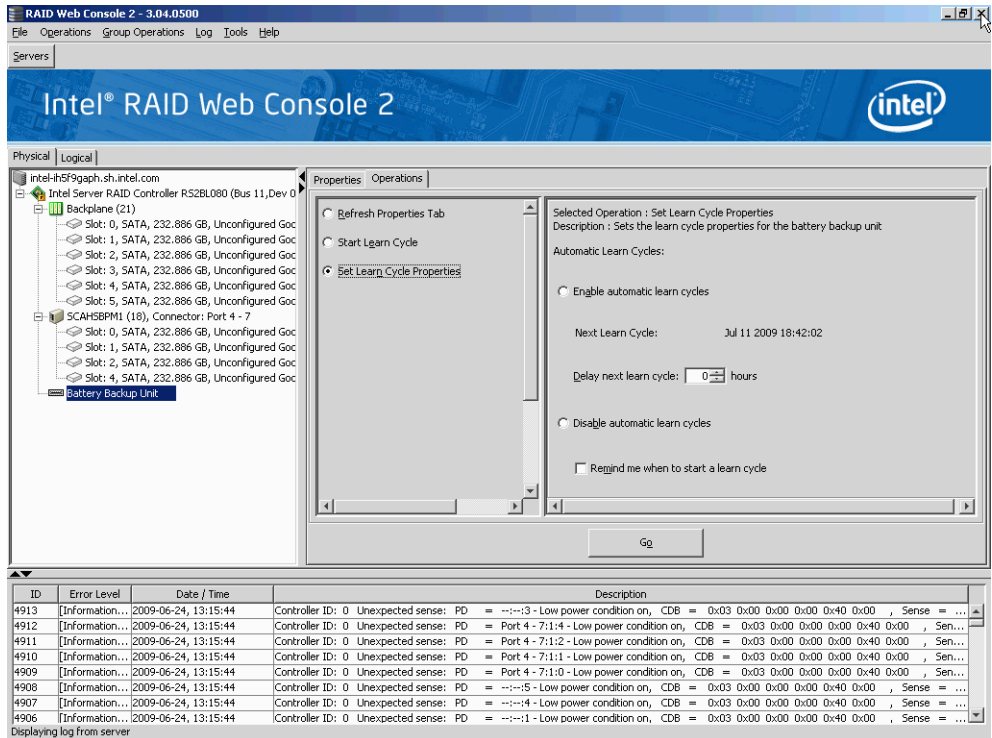


Figure 87. Battery Backup Unit Operations

4. Select **Set Learn Cycle Properties**.
The options appear in the right frame.
5. To enable automatic learn cycles, click **Enable automatic learn cycles** and click **Go**.
You can delay the start of the next learn cycle by up to 7 days (168 hours) using the **Delay next learn cycle** field.
6. To disable automatic learn cycles, click **Disable automatic learn cycles** and click **Go**.
You can start the learn cycles manually. In addition, you can check the box next to the field **Remind me when to start a learn cycle** to receive a reminder to start a manual learn cycle.

Starting a Learn Cycle Manually

To start the learn cycle properties manually, perform the following steps:

1. Click the **Physical** tab to open the physical view.
2. Select the BBU icon in the left panel.
3. Click the **Go To -> BBU** tab.

The BBU operations appear, as shown in [Figure 87](#) (Depending on different utility version, the BBU operation also may appear in a new pop-up window).

4. Click **Start Learn Cycle** and click **OK**.

Another method to use the BBU operations is to right-click the BBU icon to open the operations menu and select **Start Learn Cycle**.

Monitoring Rebuilds and Other Processes

The Intel® RAID Web Console 2 allows you to monitor the progress of rebuilds and other lengthy processes in the Group Show Progress window, shown in [Figure 88](#). You open this window by selecting **Manage -> Show Progress**.




Figure 88. Group Show Progress Window

Operations on virtual disks appear in the left panel of the window, and operations on physical drives appear in the right panel. The types of operations that appear in this window are as follows:

- Initialization of a virtual disk (see [“Initializing a Virtual Disk”](#) on page 138).
- Rebuild (see [“Rebuilding a Drive”](#) on page 140).
- Reconstruction (see [“Adding a Drive to a Virtual Disk”](#) on page 118).
- Consistency check (see [“Running a Consistency Check”](#) on page 139).

Note: A Reconstruction process cannot be aborted. To abort any other ongoing process, click **Abort** next to the status indicator. Click **Abort All** to abort all ongoing processes. Click **Close** to close the window.

Maintaining and Managing Storage Configurations

This section explains how to use Intel® RAID Web Console 2 to maintain and manage storage configurations. You must log on to the system in Full Access mode to complete these maintenance and management tasks. The following maintenance and management functions can be done:

- Initializing a Virtual Disk
- Running a Consistency Check
- Rebuilding a Drive
- Removing a Drive
- Flashing the Firmware

Initializing a Virtual Disk

When you create a new virtual disk with the Configuration Wizard, you can choose to initialize the disk immediately. To initialize a virtual disk after the configuration process, follow these steps:

1. Select the **Virtual** tab in the left panel, and click the icon of the virtual disk to initialize.
2. Select **Manage** -> **Initialize**.




Figure 89. Selecting Initialize

The Group Initialize dialog box is displayed.

3. Select the virtual disk(s) to initialize. If you want to use this option, select **Fast Initialization**.

Fast Initialization quickly formats the virtual disk by writing zeros to the first few sectors of the physical disks in the virtual disk. Depending on the number and size of the physical disks in the virtual disk, regular initialization takes longer.

4. Click **Start** to begin the initialization.

If needed, you can monitor the progress of the initialization. See [“Monitoring Rebuilds and Other Processes” on page 137](#) for more information.

Running a Consistency Check

You should periodically run a consistency check on fault-tolerant virtual disks. A consistency check scans the virtual disk to determine whether consistency data is corrupted and needs to be restored. It is especially important to do this if you suspect that the virtual disk consistency data is corrupted.

To run a consistency check, follow these steps:

1. Select **Manage -> Check Consistency**.

The Group Consistency Check window is displayed, as shown in [Figure 90](#).

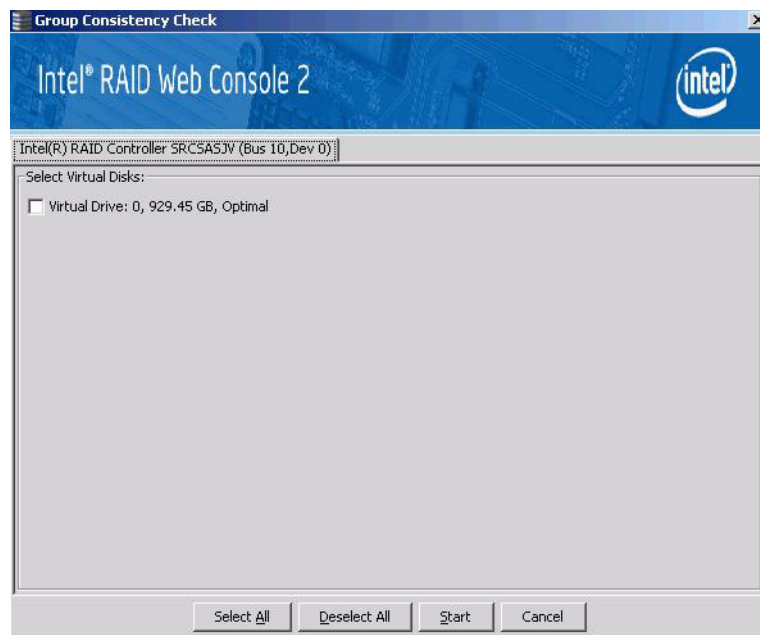


Figure 90. Group Consistency Check Window

2. Select the virtual disks you want to check, or click **Select All** to select all disks.
3. Click **Start** to begin.

You can monitor the progress of the consistency check. For more information, see [“Monitoring Rebuilds and Other Processes” on page 137](#).

Note: You can also run a consistency check by selecting the virtual disk icon in the left panel and right-click the option **Start Consistency Check**.

Scanning for New Drives

The Intel® RAID Web Console 2 normally detects newly installed disk drives and other storage devices and displays icons for them on the main screen. If the Intel® RAID Web Console 2 does not detect a new drive, you can use the **Scan for Foreign Configuration** option to find it. To do this, follow these steps:

1. Select a controller icon in the left panel.
2. Select **Go To -> Controller -> Scan Foreign Configuration**.

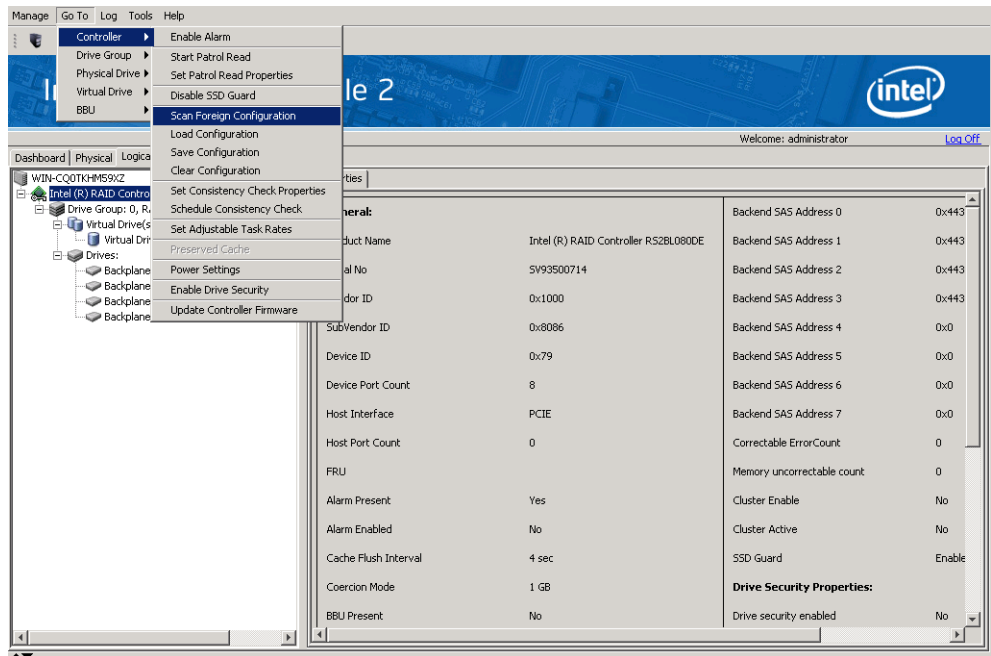


Figure 91. Scan for Foreign Configuration

If the Intel® RAID Web Console 2 detects any new disk drives, it displays a list of them on the screen.

Rebuilding a Drive

If a single drive in a fault tolerant system fails, the system is protected from data loss by the parity data (in RAID 5, RAID 50, or RAID 60) or by data redundancy (RAID 1, RAID 10). You must replace the failed drive and rebuild the drive's data on a new drive to restore the system to fault tolerance. (Or, if the drive is still operational, you can rebuild the data on the failed drive.) If dedicated or global hot-spare disks are available, the failed drive is rebuilt automatically without any user intervention.

If a drive has failed, a red circle is displayed to the right of the disk drive icon: 🟡🔴 . A small yellow circle is displayed to the right of the icon of the virtual disk that uses this physical disk: 🟡🟡 . This indicates that the virtual disk is in a degraded state, but the data is still intact.

If you need to rebuild a physical drive, follow these steps:

1. Right click the icon of the failed drive and select **Rebuild**.

- When the warning message displays, click **Yes**. If the drive is still good, a rebuild starts. You can monitor the progress of the rebuild in the Group Show Progress window by selecting **Go To -> Show Progress**. If the drive cannot be rebuilt, an error message is displayed, and you must replace the drive before a rebuild can occur. Continue with the next step.
- Click the icon of the failed drive in the left panel.
- Right-click and select **Prepare for Removal**.

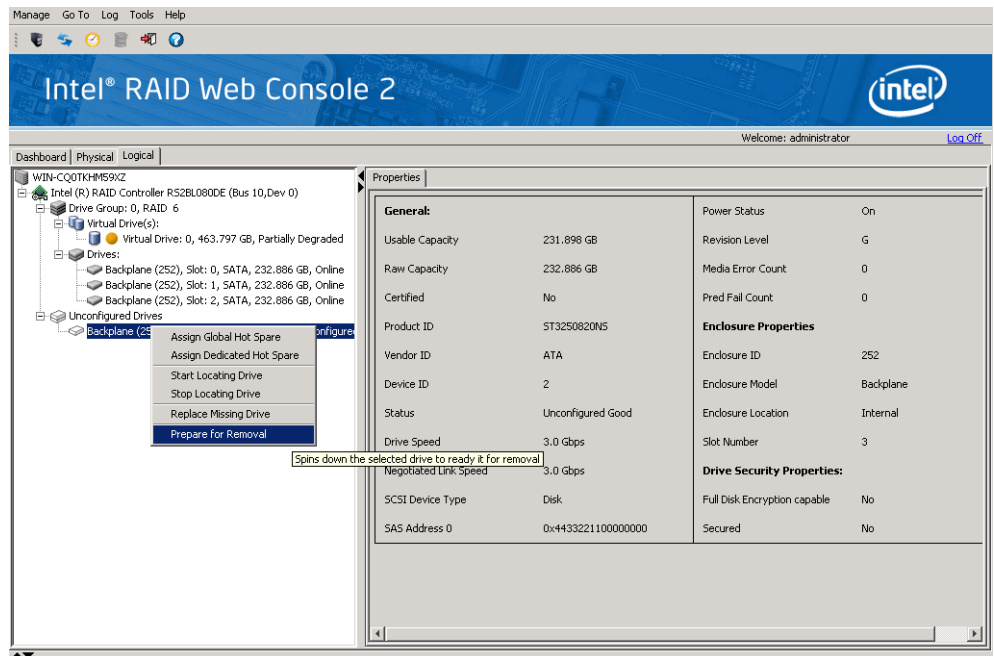


Figure 92. Preparing Drive for Removal

- Physically remove the failed drive and replace it with a new drive of equal or greater capacity. When the new drive spins up, the drive icon changes to green, and the rebuild begins automatically. You can monitor the progress of the rebuild in the Group Show Progress window by selecting **Manage -> Show Progress**.

Note: If you want to force a disk drive into Fail status to trigger a rebuild, right-click the drive icon and select **Make Drive Offline**.

Removing a Drive

You may sometimes need to remove a non-failed drive that is connected to the controller. For example, you may need to replace the drive with a larger drive. Follow these steps to remove a drive safely:

- Click the icon of the drive in the left panel and right-click it.
- Select **Make Drive Offline**.
- Right-click again and select **Mark Drive as Missing**.

4. Right-click again and select **Prepare for Removal**. This will change drive status to **Ready to Remove**
5. If you change your mind, right-click it and select **Undo Prepare for Removal**. Otherwise, wait until the drive spins down and then remove it.

Warning: *While the system is powered off, never replace a drive that has not failed (and is not marked by the controller as failed). A drive must always be failed before it is replaced in an array.*

Flashing the Firmware

The Intel® RAID Web Console 2 enables you to easily upgrade the controller firmware. To flash the controller firmware, follow these steps:

1. In the left panel, click on the icon of the controller you need to upgrade.
2. Right-click it and select **Update Controller Firmware**.
3. Browse for the .rom flash update file and click **OK**.

The Intel® RAID Web Console 2 displays the version of the existing firmware and the version of the new firmware file.

4. When you are prompted to ask if you want to upgrade the firmware, click **Yes**.

The controller is updated with the new firmware code contained in the .ROM file.

Enabling RAID Premium Features

Enabling Full Disk Encryption feature

This feature is supported by Intel® RAID controller RS2BL080DE, RS2PI008DE, and by other controllers with Intel® RAID Premium Feature Key AXXRPFKDE installed. This feature also requires that Self-Encrypting Drives (SEDs) are connected to the RAID controller. Refer to Intel® RAID Premium Feature Key AXXRPFKSSD, AXXRPFKDE, and AXXRPFKSNSH Installation Guide (E88588-00x) for physical installation of the Premium Feature Key.

This feature enables authentication key management for data security protection for Self Encrypting Drives (SED), auto-lock, and instant erase of self-encrypting drives.

- It transparently auto-Locks Volumes; Volumes are automatically locks when powered down.
- It provides secure data when a drive is misplaced, failed, experiencing drive theft, experiencing Server theft “Smash & Grab”
- It provides Instant Secure Erase: Before returning leased Drives; Before redeploying with the datacenter.
- It helps reduce risk of data being compromised when drives eventually leave the data center.
- It significantly reduces the cost and time of repurposing or retiring drives.

- It is 100% transparent to storage systems and end users; appears the same as traditional non-encrypting drives; operates at full drive speed - No impact to performance; automatically encrypts or decrypts all data.

The following steps describe the steps to use Full Disk Encryption feature:

- Enable this feature within the Intel® RAID Web Console 2.
- Build a logical volume containing Self-Encrypting Disks.
- Specify that the volume be secured.

Detailed steps are as below:

1. Select the RAID controller and choose **Physical** tab in the left panel of Intel® RAID Web Console 2. Check **Drive Security Properties** in the right panel to see whether the RAID controller is **Drive Security capable**.

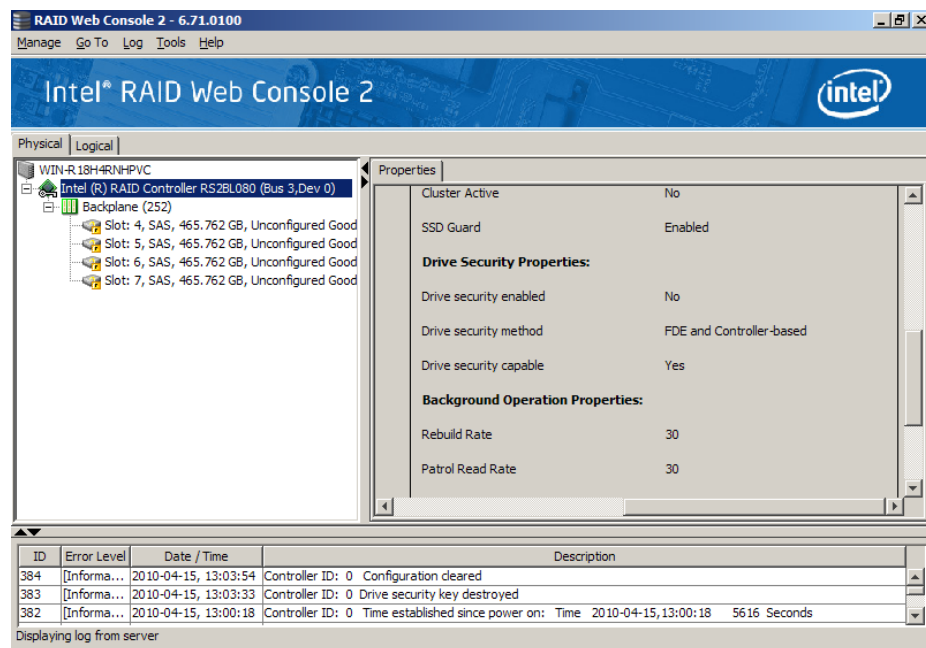


Figure 93. Check Controller Security status

2. Select the physical drive and choose **Physical** tab in the left panel of Intel® RAID Web Console 2. Check **Drive Security Properties** in the right panel to see whether the RAID controller is **Full Disk Encryption capable**.

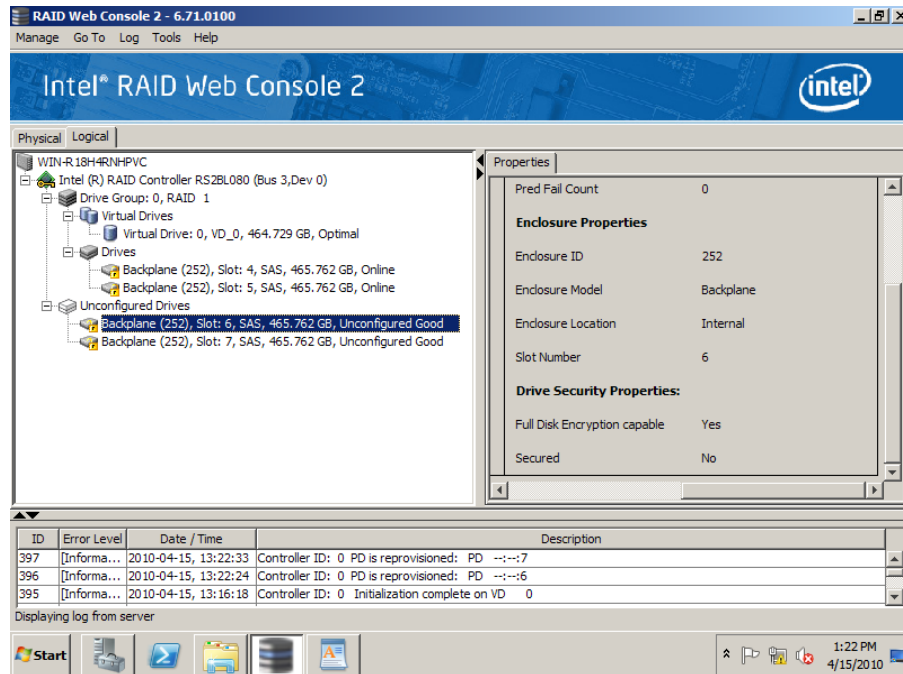


Figure 94. Check Drive Security status

3. Right-click on the controller and select **Enable Drive Security**.




Figure 95. Enable Drive Security

4. In the Security Wizard, click **Next**.

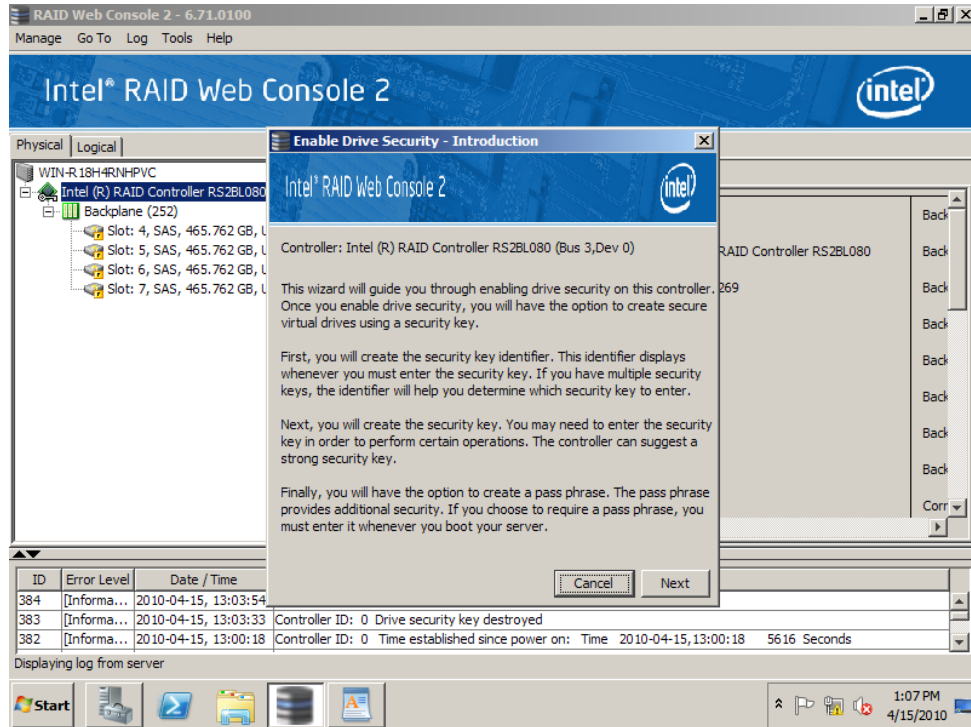


Figure 96. Start Security Wizard

5. Create the required Security key Identifier. Use the default suggestion, or create your own, then click **Next**.

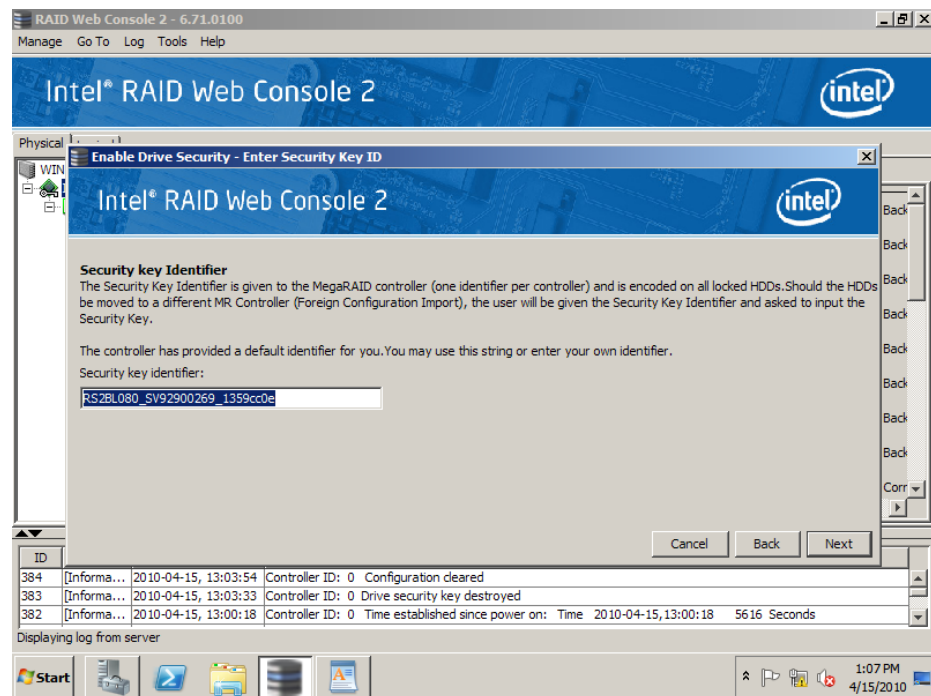


Figure 97. Enter Security Key Identifier

6. Create the required Security key. Use the default suggestion, or create your own, then click **Next**.




Figure 98. Enter Security Key

7. Create an optional Pass Phrase for additional security, then click **Next**.




Figure 99. Enter Pass Phrase

- Accept the confirmation to complete the wizard and enable drive security, then click **Finish**.




Figure 100. Confirm Enable Drive Security

9. Controller properties now indicate **Drive security enabled** is **Yes**.




Figure 101. Check Drive Security Enabled status

10. Select Unconfigured Good drives and select Drive security method, here FDE to create the virtual drive.




Figure 102. Select Full Disk Encryption

11. Follow same steps for creating a RAID virtual drive. After a virtual drive is created, Selecting this virtual drive in left panel shows its **secured** status and **Drive security method** status in the **Drive Security Properties** part in the right panel. Log information at the bottom of the window also indicates the virtual drive security status. Full Disk Encryption feature has now been enabled.




Figure 103. Create RAID Virtual Drive with FDE enabled

12. (Optional) For Self-Encrypting Drives (SEDs) that need to have their data permanently erased for security reason, follow normal steps to delete the secured RAID virtual drive, and mark all SEDs as Unconfigured Good physical drives. Then select one SED, right-click it and select **Instant Erase**.




Figure 104. Instant Secure Erase

13. Accept confirmation to permanently erase disk drive.




Figure 105. Confirm Secure Erase

Enabling Snapshot feature

This feature is supported by Intel® RAID Controllers with Intel® RAID Premium Feature Key AXXRPFKSNSH installed. Refer to Intel® RAID Premium Feature Key AXXRPFKSSD, AXXRPFKDE, and AXXRPFKSNSH Installation Guide (E88588-00x) for physical installation of the Premium Feature Key.

The MegaRAID Recovery feature, also known as Snapshot, offers a simplified way to recover data and provides automatic protection for the boot volume. You can use MegaRAID Recovery to take a snapshot of a volume and to restore a volume or file.

Snapshot functionality allows users to capture data changes to the volume. If data is deleted accidentally or maliciously, you can restore the data from view or roll back to a snapshot at a previous Point in Time (PiT).

You can make up to 8 snapshots of Point in Times for each volume.

You can either restore files or roll back applications to a previous point in time.

Recovers lost data from disk in minutes, minimizing the downtime experienced by users and eliminating the lengthy process of restoring data from tape.

Allows isolation of a server's boot volume to a separate virtual drive; by enabling the Auto Snapshot feature, users are guaranteed that bootable snapshots exist, in case corruption of operating system occurs.

***Note:** MegaRAID Recovery Snapshots are not a substitute for a tested backup strategy. When Recovery Snapshot is enabled, changes to the selected volume are tracked. When a roll back to a snapshot at a previous Point in Time (PiT) is done, you are "undoing" changes to the volume that caused corruption or a file deletion. The first snapshot does not create a full volume backup. Recovery Snapshot is equivalent to Microsoft* System Restore with additional features and capabilities.*

Detailed Configuration steps are as below:

1. Select the **Logical** view. Select and highlight a **virtual drive** from the list of virtual drives (If required, initialize, allocate and format the virtual drive).
2. Select **Go To -> Virtual Drive -> Enable MegaRAID Recovery** on the menu bar.

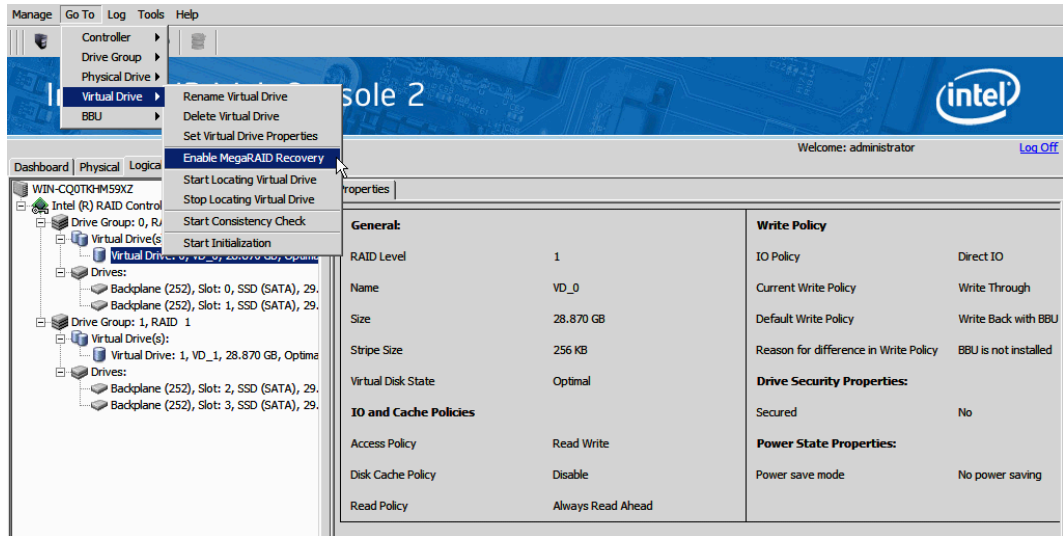


Figure 106. Enable MegaRAID Recovery

3. On the Enable MegaRAID Recovery Wizard screen, select the virtual drive to use as the Snapshot Repository in the **Snapshot Repository** field. (You can leave the suggested virtual drive as the Snapshot Repository.)
4. Enter the capacity to use in the Snapshot Repository for changes to the base virtual drive. The available capacity is the largest free block of capacity on the snapshot repository virtual drive.




Figure 107. Enter the Capacity for Snapshot Repository

5. Choose whether to have a snapshot taken on reboot. To enable this option, check the box next to the **Take snapshot on reboot** field. If you select this option, a snapshot is taken on boot after every successful shutdown. You can use this snapshot of the boot virtual drive to restore the operating system on the virtual drive in case the virtual drive becomes corrupted.

Caution: *Before you select this option, make sure that you backup all data to the virtual drive. If there is any existing data on this virtual drive, it will be lost after converting to the Snapshot repository.*

6. Click **OK**.
7. A confirmation dialog box appears. Check the box next to the Confirm field to enable snapshots on the selected virtual drive and then click **Yes**.




Figure 108. Confirm Enable Snapshot

8. This virtual drive becomes a snapshot repository. Use it only for storing snapshot-related data.

Caution: *After you enable snapshots on this virtual drive, you cannot change the allocated percentage of capacity or the snapshot repository without first disabling snapshots and losing any snapshot data.*

Note: *With the current generation of the MegaRAID Recovery product there is no guidance on the size of the repository other than saying that it is completely dependent on the amount of data being written and the frequency of snapshots being taken. This is different in every application. Currently a customer will have to use empirical data on the behavior of their application to estimate this. If the repository runs out of space they will have to delete the oldest snapshot. Then at some time when it is appropriate they will have to disable the feature and re-enable it, re-sizing the repository.*

The existing snapshots are lost when the feature is disabled. Repository virtual drives will be lost if you disable snapshots.

In future development, users will be able to grow the size of the repository dynamically, without losing snapshots.




Figure 109. Snapshot Base is shown

- Following steps create Snapshots and Views, based on above configurations for Snapshot Base and Snapshot Repository.

Note: You can create up to eight snapshots of a volume. snapshots are shown in chronological order from the oldest to the newest. Each snapshot is a PiT snapshot of the virtual drive that is the Snapshot Base

The snapshots appear on the timeline from the oldest on the left to the newest on the right. If you create the maximum number of snapshots allowed, eight, the **Create Snapshot** button is disabled. You cannot create an additional snapshot unless you delete a snapshot.

After you create the snapshots, you can create views of the PiT snapshots. You can search the views to find a snapshot that does not contain corrupt data or a snapshot that contains deleted data, depending on the situation. You can reboot and roll back to that snapshot to restore the data.

- Click the **Logical** view on the main menu screen. Click the **Snapshot Base** virtual drive in the left frame. Click the **Snapshots** tab in the right frame.




Figure 110. Enter Snapshot Name

- Enter the snapshot name in the **Enter snapshot name** field. Click **Create Snapshot**.

- Repeat above step so that additional snapshots can be created. The snapshot you create appears in the snapshot timeline. The oldest snapshot is on the left end of the timeline.




Figure 111. Create Snapshot

- (Assuming you experience a file deletion or corruption, following steps create MegaRAID Recovery Views based on snapshots created previously, and can roll back snapshots.) Click the **Logical** view on the main menu screen. Click the **Snapshot Base** virtual drive in the left frame. Click the **Snapshots** tab in the right frame (This screen shows the Snapshot Base details and any existing snapshots). Click **Create View**.

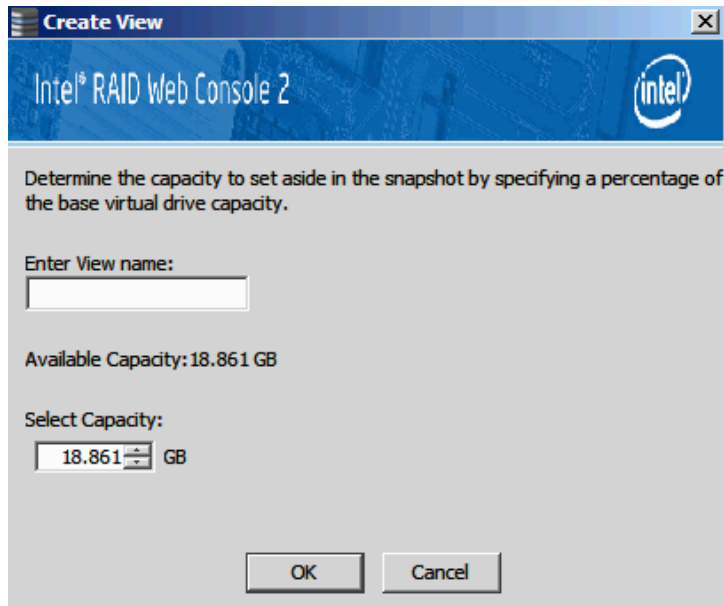


Figure 112. Create View

- Enter the name of the view in the **Enter View name** field and the capacity to set aside in the snapshot. Click **OK**.

Note: The “View name” can be found in the snapshots screen view in the View Details information box after a snapshot has been created

15. Enter into “**My computer**”, and then you can see a new drive.
16. Copy the files in the new drive to the original drive and cover them. Now the system rolls back to the selected PiT.
17. Following steps delete and clear snapshots.

Note: When deleting single snapshots, you can only delete the oldest snapshot.

18. Follow these steps to delete a single snapshot:

Click the **Logical** view on the main menu screen; Click the **Snapshot Base** virtual drive in the left frame; Click the **Snapshots** tab in the right frame; Click the oldest snapshot in the timeline; Click the **Delete Snapshot** button.

19. Follow these steps to clear (delete) all snapshots:

Click the **Logical** view on the main menu screen; Click to select the **Snapshot Base** virtual drive in the left frame; Click **Go To -> Virtual Drive -> Clear Snapshots** on the menu bar (Or right click the Snapshot Base Virtual Drive and click Clear Snapshots); A confirmation dialog box appears. It warns that any snapshot-related data that is on the associated Snapshot Repository virtual drives will be lost if you clear the snapshots; Check the box next to the **Confirm** field and click **Yes** and the snapshots will be cleared.

20. Follow these steps to Set MegaRAID Recovery Properties:

Click the **Logical** view on the main menu screen; Click the **Snapshot Base** virtual drive in the left frame; Click **Go To -> Virtual Drive -> Set MegaRAID Recovery Properties** on the menu bar; Check the box to the properties selections you desire and click **OK**.

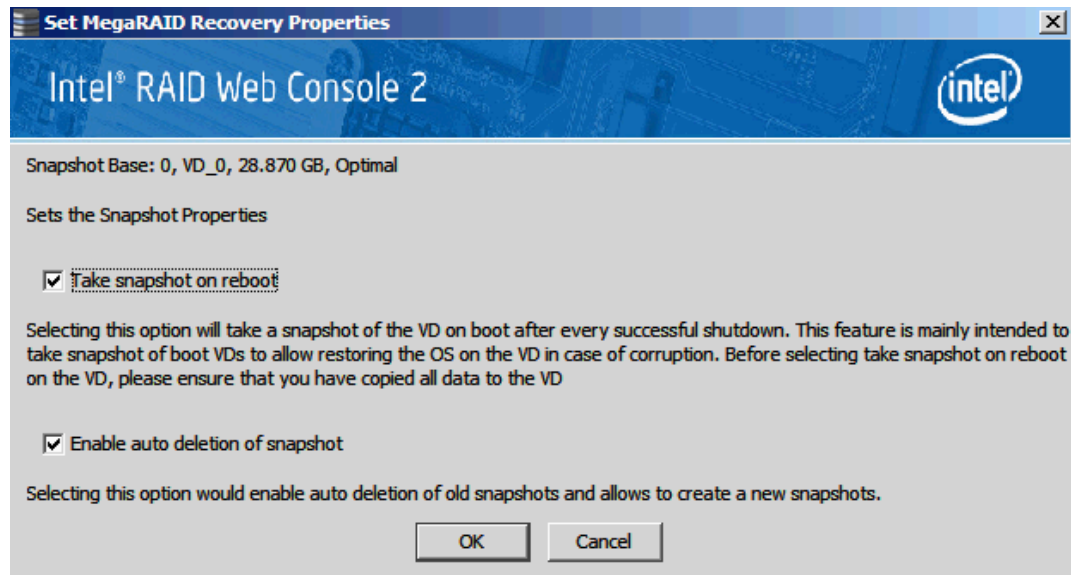


Figure 113. Set MegaRAID Recovery Properties

21. Follow these steps to Disable MegaRAID Recovery:

Click the **Logical** view on the main menu screen; Click the **Snapshot Base** virtual drive in the left frame; Click **Go To -> Virtual Drive -> Disable MegaRAID Recovery** on the menu bar.

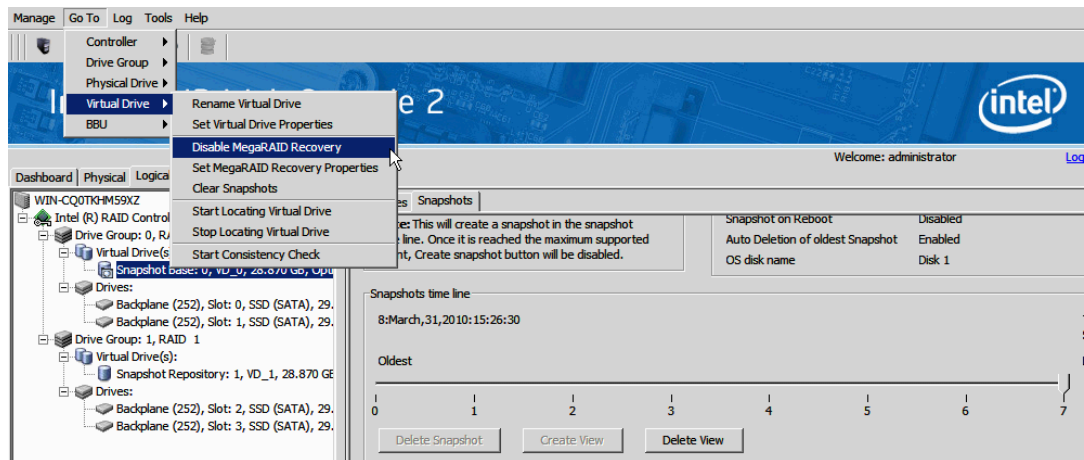


Figure 114. Disable MegaRAID Recovery

22. A confirmation dialog box appears. It warns that any snapshot-related data that is on the associated Snapshot Repository virtual drives will be lost if you disable snapshots; Check the box next to the **Confirm** field and click **Yes** . The snapshots are disabled.




Figure 115. Confirm Disable Snapshots

23. After snapshots are successfully created (in step 12, Figure 111), following steps roll back to a snapshot at a previous Point in Time (PiT).

Note: If the operating system is installed on a Snapshot Base, and snapshots are created for it and stored in Snapshot Repository, the rolling back steps can help recover the operating system into a previous status.

24. Reboot the server system and press <Ctrl> + <G> to start the RAID BIOS Console 2 Utility. When you see the screen below, press **Start**.




Figure 116. Adapter Selection

- In the home page, from the **Logical View** in the right panel, click the **Snapshot Base** which has the operating system installed.




Figure 117. Selecting Snapshot Base

- In the new screen, select **Adv Opers** (which means Advanced Operations), and click **Go**.




Figure 118. Selecting Advanced Operations

27. In the Advanced Operations screen, select **Rollback**.

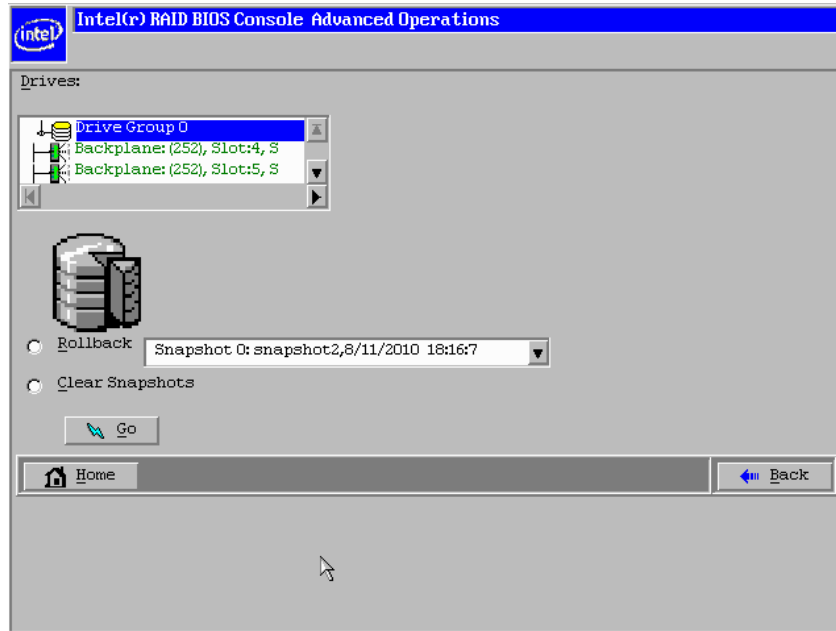


Figure 119. Selecting Rollback

28. From the drop-down menu, select a previous created snapshot, and then click **Go**.




Figure 120. Selecting a snapshot

29. A Confirm Page should appear. Please read carefully the warning message. *“You have chosen to restore a previous snapshot time line. Doing so will rollback to the snapshot to its state at the date and time you selected. You may lose any changes you have made subsequent to that date and time. Are you sure you want to restore snapshot <snapshot name> <date> <time>?”*. When confirmed, click **Yes**.




Figure 121. Confirm Page

30. The screen should go back to the display below. Click Home to go back to home page. Then exit the RAID BIOS Console 2 Utility and reboot the server system. You should be able to enter the operating system into a selected previous status.

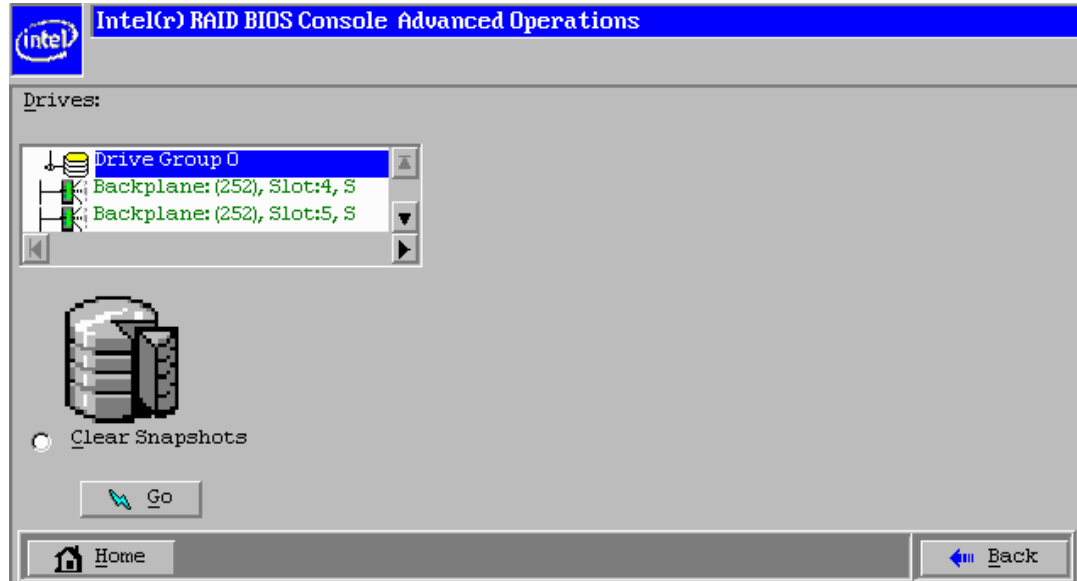


Figure 122. Rollback operation is done

Enabling Super Sized Cache

This feature is supported by Intel® RAID Controllers with Intel® RAID Premium Feature Key AXXRPFKSSD installed. Refer to Intel® RAID Premium Feature Key AXXRPFKSSD, AXXRPFKDE, and AXXRPFKSNSH Installation Guide (E88588-00x) for physical installation of the Premium Feature Key.

Super Size Cache (SSC) using Solid State Devices (SSD) as controller cache allows for very large data sets to be present in that cache to improve overall application performance. SSC anticipate up to a five times performance improvement in READ-intensive applications and provide maximum IOPs.

This solution is ideal for cost sensitive server environments, by offering a high-performance upgrade that requires only a small investment in SSD technology. This solution also reduces the need to invest expensive IT cycles to hand-tune applications running on hard drive-only disk arrays.

Detailed Configuration steps are as below:

1. Select the **Dashboard** view. Click **Create super sized cache (SSC)** in the **Actions** tab; Or right click the **RAID Controller** and select **Create SSCD** (Super Sized Cache Drive).




Figure 123. Create SSC from Dashboard

2. The **Create Super Sized Cache Drive group** screen appears. Select the **unconfigured SSD drive**. Click **Add >** -> **Create SSC Drive Group** -> **Next**.

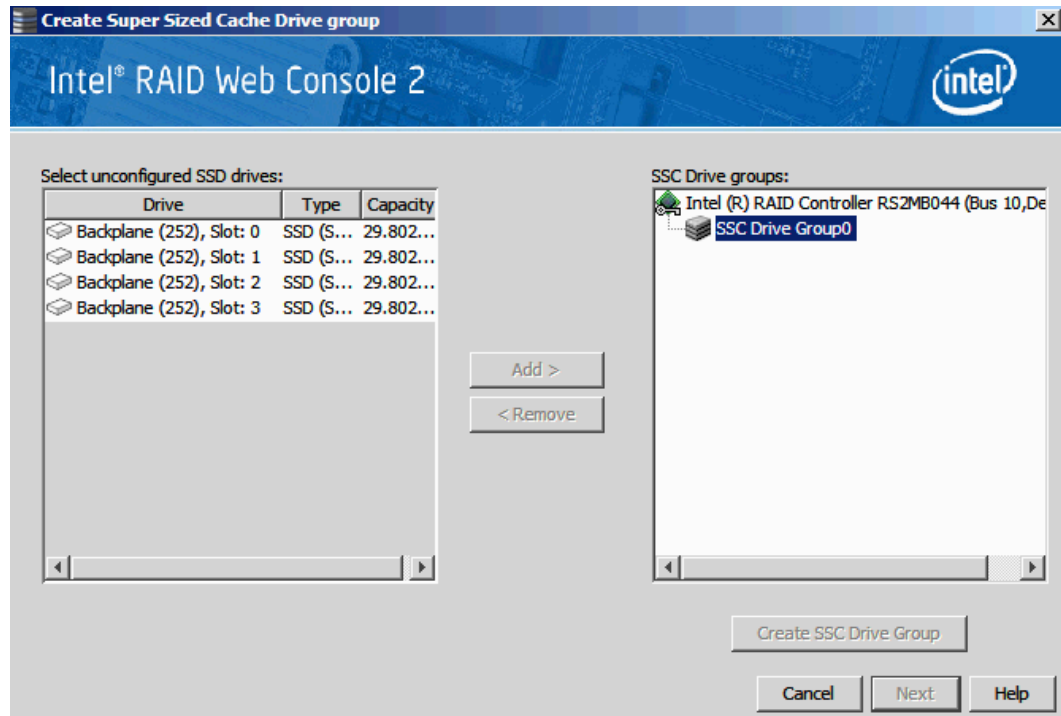


Figure 124. Create SSC Drive Group

3. The **Create Super Sized Cache Drive** screen appears. Click **Create** to create a SSC drive group. Click **Next**.




Figure 125. Create SSCD name

4. Create Super Sized Cache screen appears. Click **Finish** to confirm the information and create the Super Sized Cache Disk (SSCD).




Figure 126. SSC Summary

5. You can see **Super Sized Cache Drive Group** in the logical view.

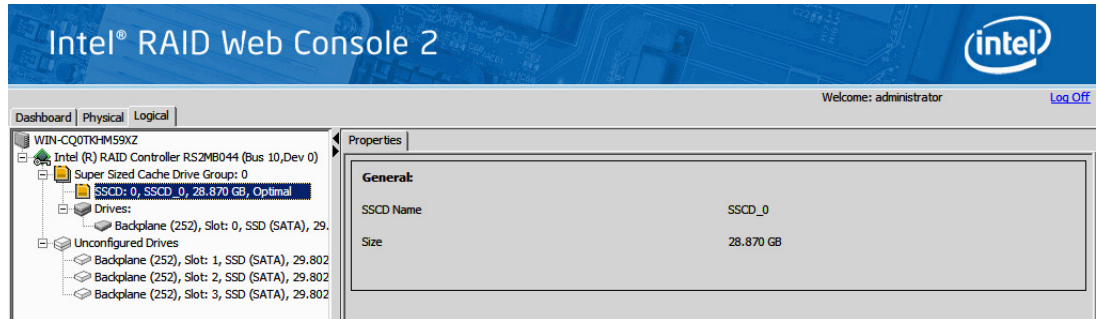


Figure 127. SSCD status shown

- Following steps delete a Super Sized Cache disk:

Click the **Logical** view tab on the main menu screen; Click the **SSCD virtual drive** in the left frame; Click **Go To -> Virtual Drive -> Delete Virtual Drive** on the menu bar.

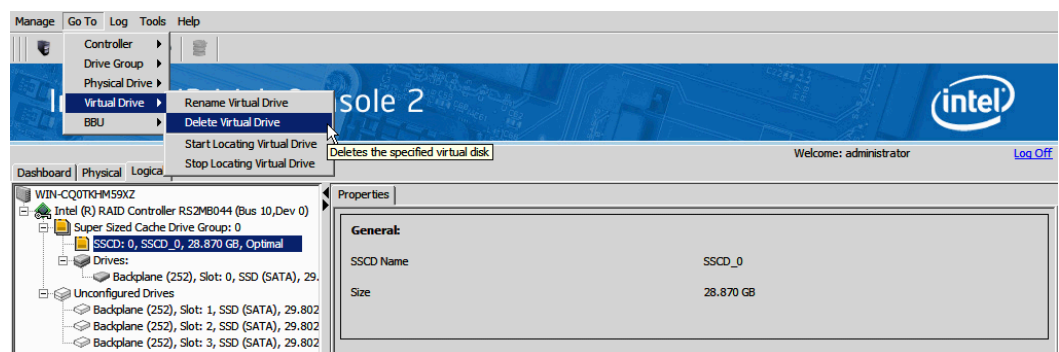


Figure 128. Delete SSCD

- A confirmation dialog box appears. Check the box next to the **Confirm** field and click **Yes**. The SSCD is deleted.

Appendix A: Creating a Virtual Drive Using Advanced Configuration

The advanced configuration procedure provides an easy way to create a new storage configuration. Advanced configuration gives you greater flexibility than simple configuration because you can select the drives and the virtual drive parameters when you create a virtual drive. In addition, you can use the advanced configuration procedure to create spanned drive groups.

Note: Depending on different versions of utility, the steps in screenshots may change, and below sections are for reference purpose.

Follow these steps to create a new storage configuration in the advanced configuration mode. In this example, we will create a spanned drive group.

1. Perform either of the following steps:

- Right click on the controller node in the device tree in the left frame of the Intel[®] RAID Web Console 2 window and select Create Virtual Drive
- Select the controller node and select Operations->Create Virtual Drive in the menu bar, as shown in [Figure 129](#)




Figure 129. Virtual Drive Creation Menu

The dialog box shown in [Figure 130](#) appears.




Figure 130. Virtual Drive Creation Mode

2. Click **Advanced** and press **Next**.

The Create Drive Group Settings screen appears, as shown in [Figure 131](#).

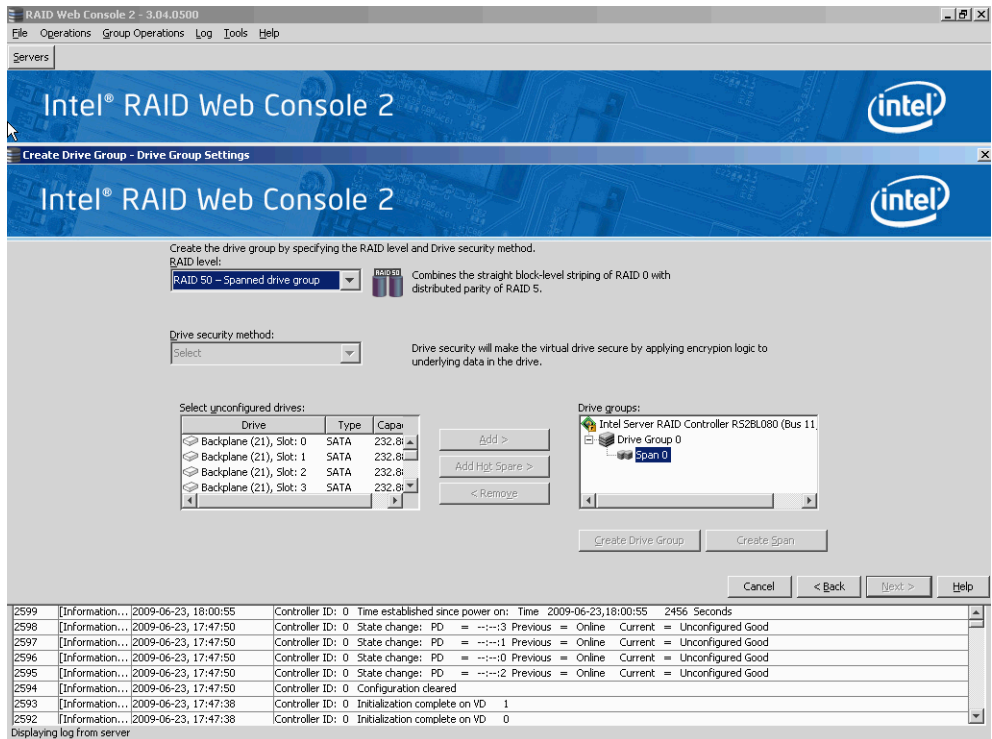


Figure 131. Create Drive Group Settings Screen

3. Select the following items on the Create Drive Group Settings screen:
 - a. Select the RAID level desired for the drive group from the drop-down menu. To make a spanned drive, select **RAID 10**, **RAID 50**, or **RAID 60** in the **RAID level** field.

Drive Group 0 and **Span 0** appear in the **Drive groups** field when you select RAID 10, 50, or 60.

The RAID controller supports RAID levels 1, 5, 6, 10, 50, and 60. In addition, it supports independent drives (configured as RAID 0 and RAID 00). The screen text gives a brief description of the RAID level you select. RAID levels you can choose depend on the number of drives available. To learn more about RAID levels, see Chapter 2, “RAID Levels” on page 9.

- b. Select unconfigured drives from the list of drives and click **Add>** to add them to the drive group.

The selected drives appear under **Span 0** below **Drive Group 0**, as shown in [Figure 132](#)




Figure 132. Span 0 of Drive Group 0

- c. Click **Create Span** to create a second span in the drive group.
- d. Select *unconfigured* drives from the list of drives and click **Add>** to add them to the drive group.
- e. The selected drives appear under **Span 1** below **Drive Group 0**, as shown in [Figure 133](#)




Figure 133. Span 0 and Span 1 of Drive Group 0

- f. Click **Create Drive Group** to make a drive group with the spans.
- g. Click **Next** to complete this step.

The Virtual drive settings window appears, as shown in Figure 134. The drive group and the default virtual drive settings appear. The options to update the virtual drive or remove the virtual drive are grayed out until you create the virtual drive.




Figure 134. Virtual Drive Settings Window

4. Change any virtual drive settings, if desired.
5. Click **Create Virtual Drive**.

The new virtual drive appears under the drive group, as shown in [Figure 135](#). The options **Update Virtual Drive** and **Remove Virtual Drive** are now available. **Update Virtual Drive** allows you to change the virtual drive settings and **Remove Virtual Drive** allows you to delete the virtual drive.




Figure 135. New Virtual Drive 0

- Click Next.

The **Create Virtual Drive - Summary** window appears, as shown in Figure 136. This window shows the selections you made for advanced configuration.




Figure 136. Create Virtual Drive Summary Window

- Click **Back** to return to the previous screen to change any selections or click **Finish** to accept and complete the configuration.

The new storage configuration will be created and initialized.

Note: If you create a large configuration using drives that are in powersave mode, it could take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a box appears to identify the drive or drives.

8. After the configuration is completed, a dialog box notifies you that the virtual drives were created successfully, as shown in [Figure 137](#). If more drive capacity exists, the dialog box asks whether you want to create more virtual drives. If no more drive capacity exists, you are prompted to close the configuration session.




Figure 137. Option to Create Additional Virtual Drives

9. Select **Yes** or **No** to indicate whether you want to create additional virtual drives.
If you select **Yes**, the system takes you to the **Create Virtual Drive** screen, as shown in [Figure 50](#). If you select **No**, the utility asks whether you want to close the wizard, as [Figure 138](#).
10. If you selected **No** in <hyperactive>step 8, select Yes or No to indicate whether you want to close the wizard.
If you select **Yes**, the configuration procedure closes. If you select **No**, the dialog box closes and you remain on the same page.




Figure 138. Option to Close the Configuration Wizard

Appendix B: Events and Messages

This appendix lists the Intel® RAID Web Console 2 events that may display in the event log.

The Intel® RAID Web Console 2 monitors the activity and performance of all controllers in the server and the devices attached to them. When an “event” such as the completion of a consistency check or the removal of a physical drive occurs, an event message is displayed in the log displayed at the bottom of the Intel® RAID Web Console 2 screen. The messages are also logged in the Microsoft Windows* Application log (Event Viewer). Error event levels are:

- **Progress:** This is a progress posting event. Progress events are not saved in NVRAM.
- **Info:** Informational message. No user action is necessary.
- **Warning:** Some component may be close to a failure point.
- **Critical:** A component has failed, but the system has not lost data.
- **Fatal:** A component has failed, and data loss has occurred or will occur.
- **Dead:** A catastrophic error has occurred and the controller has died. This is seen only after the controller has been restarted.

The following table lists all of the Intel® RAID Web Console 2 event messages.

Table 14. MFI Messages

Number	Type	Description
0	Info	Firmware initialization started (PCI ID %04x/%04x/%04x/%04x)
1	Info	Firmware version %s
2	Fatal	Unable to recover cache data from TBBU
3	Info	Cache data recovered from TBBU successfully
4	Info	Configuration cleared
5	Warning	Cluster down; communication with peer lost
6	Info	Virtual drive %s ownership changed from %02x to %02x
7	Info	Alarm disabled by user
8	Info	Alarm enabled by user
9	Info	Background initialization rate changed to %d%%
10	Fatal	Controller cache discarded due to memory/battery problems
11	Fatal	Unable to recover cache data due to configuration mismatch
12	Info	Cache data recovered successfully
13	Fatal	Controller cache discarded due to firmware version incompatibility
14	Info	Consistency Check rate changed to %d%%
15	Dead	Fatal firmware error: %s
16	Info	Factory defaults restored
17	Info	Flash downloaded image corrupt
18	Critical	Flash erase error
19	Critical	Flash timeout during erase
20	Critical	Flash error
21	Info	Flashing image: %s
22	Info	Flash of new firmware image(s) complete
23	Critical	Flash programming error
24	Critical	Flash timeout during programming
25	Critical	Flash chip type unknown
26	Critical	Flash command set unknown
27	Critical	Flash verify failure
28	Info	Flush rate changed to %d seconds
29	Info	Hibernate command received from host
30	Info	Event log cleared
31	Info	Event log wrapped
32	Dead	Multi-bit ECC error: ECAR=%x
33	Warning	Single-bit ECC error: ECAR=%x
34	Dead	Not enough controller memory
35	Info	Patrol Read complete
36	Info	Patrol Read paused
37	Info	Patrol Read Rate changed to %d%%
38	Info	Patrol Read resumed
39	Info	Patrol Read started
40	Info	Rebuild rate changed to %d%%

Number	Type	Description
41	Info	Reconstruction rate changed to %d%%
42	Info	Shutdown command received from host
43	Info	Test event: '%s'
44	Info	Time established as %s; (%d seconds since power on)
45	Info	User entered firmware debugger
46	Warning	Background Initialization aborted on %s
47	Warning	Background Initialization corrected medium error (%s at %lx
48	Info	Background Initialization completed on %s
49	Fatal	Background Initialization completed with uncorrectable errors on %s
50	Fatal	Background Initialization detected uncorrectable double medium errors (%s at %lx on %s)
51	Critical	Background Initialization failed on %s
52	Progress	Background Initialization progress on %s is %s
53	Info	Background Initialization started on %s
54	Info	Policy change due to BBU on %s from %s to %s
55	Info	Policy change due to user on %s from %s to %s
56	Warning	Consistency Check aborted on %s
57	Warning	Consistency Check corrected medium error (%s at %lx
58	Info	Consistency Check done on %s
59	Info	Consistency Check done with corrections on %s
60	Fatal	Consistency Check detected uncorrectable double medium errors (%s at %lx on %s)
61	Critical	Consistency Check failed on %s
62	Fatal	Consistency Check failed with uncorrectable data on %s
63	Warning	Consistency Check found inconsistent parity on %s at strip %lx
64	Warning	Consistency Check inconsistency logging disabled on %s (too many inconsistencies)
65	Progress	Consistency Check progress on %s is %s
66	Info	Consistency Check started on %s
67	Warning	Initialization aborted on %s
68	Critical	Initialization failed on %s
69	Progress	Initialization progress on %s is %s
70	Info	Fast initialization started on %s
71	Info	Full initialization started on %s
72	Info	Initialization complete on %s
73	Info	LD Properties updated to %s (form %s)
74	Info	Reconstruction complete on %s
75	Fatal	Reconstruction of %s stopped due to unrecoverable errors
76	Fatal	Reconstruct detected uncorrectable double medium errors (%s at %lx on %s at %lx)
77	Progress	Reconstruction progress on %s is %s
78	Info	Reconstruction resumed on %s
79	Fatal	Reconstruction resume of %s failed due to configuration mismatch
80	Info	Reconstructing started on %s

Number	Type	Description
81	Info	State change on %s from %s to %s
82	Info	PD Clear aborted on %s
83	Critical	PD Clear failed on %s (Error %02x)
84	Progress	PD Clear progress on %s is %s
85	Info	PD Clear started on %s
86	Info	PD Clear completed on %s
87	Warning	Error on %s (Error %02x)
88	Info	Format complete on %s
89	Info	Format started on %s
90	Critical	Hot Spare SMART polling failed on %s (Error %02x)
91	Info	PD inserted: %s
92	Warning	PD %s is not supported
93	Warning	Patrol Read corrected medium error on %s at %lx
94	Progress	Patrol Read progress on %s is %s
95	Fatal	Patrol Read found an uncorrectable medium error on %s at %lx
96	Critical	Predictive failure: CDB: %s
97	Fatal	Patrol Read puncturing bad block on %s at %lx
98	Info	Rebuild aborted by user on %s
99	Info	Rebuild complete on %s
100	Info	Rebuild complete on %s
101	Critical	Rebuild failed on %s due to source drive error
102	Critical	Rebuild failed on %s due to target drive error
103	Progress	Rebuild progress on %s is %s
104	Info	Rebuild resumed on %s
105	Info	Rebuild started on %s
106	Info	Rebuild automatically started on %s
107	Critical	Rebuild stopped on %s due to loss of cluster ownership
108	Fatal	Reassign write operation failed on %s at %lx
109	Fatal	Unrecoverable medium error during rebuild on %s at %lx
110	Info	Corrected medium error during recovery on %s at %lx
111	Fatal	Unrecoverable medium error during recovery on %s at %lx
112	Info	PD removed: %s
113	Warning	CDB: %s
114	Info	State change on %s from %s to %s
115	Info	State change by user on %s from %s to %s
116	Warning	Redundant path to %s broken
117	Info	Redundant path to %s restored
118	Info	Dedicated Hot Spare PD %s no longer useful due to deleted array
119	Critical	SAS topology error: Loop detected
120	Critical	SAS topology error: Unaddressable device
121	Critical	SAS topology error: Multiple ports to the same SAS address
122	Critical	SAS topology error: Expander error
123	Critical	SAS topology error: SMP timeout

Number	Type	Description
124	Critical	SAS topology error: Out of route entries
125	Critical	SAS topology error: Index not found
126	Critical	SAS topology error: SMP function failed
127	Critical	SAS topology error: SMP CRC error
128	Critical	SAS topology error: Multiple subtractive
129	Critical	SAS topology error: Table to table
130	Critical	SAS topology error: Multiple paths
131	Fatal	Unable to access device %s
132	Info	Dedicated Hot Spare created on %s (%s)
133	Info	Dedicated Hot Spare %s disabled
134	Critical	Dedicated Hot Spare %s no longer useful for all arrays
135	Info	Global Hot Spare created on %s (%s)
136	Info	Global Hot Spare %s disabled
137	Critical	Global Hot Spare does not cover all arrays
138	Info	Created %s}
139	Info	Deleted %s}
140	Info	Marking LD %s inconsistent due to active writes at shutdown
141	Info	Battery Present
142	Warning	Battery Not Present
143	Info	New Battery Detected
144	Info	Battery has been replaced
145	Critical	Battery temperature is high
146	Warning	Battery voltage low
147	Info	Battery is charging
148	Info	Battery is discharging
149	Info	Battery voltage is normal
150	Fatal	Battery needs to be replacement
151	Info	Battery relearn started
152	Info	Battery relearn in progress
153	Info	Battery relearn completed
154	Critical	Battery relearn timed out
155	Info	Battery relearn pending: Battery is under charge
156	Info	Battery relearn postponed
157	Info	Battery relearn will start in 4 days
158	Info	Battery relearn will start in 2 day
159	Info	Battery relearn will start in 1 day
160	Info	Battery relearn will start in 5 hours
161	Info	Battery removed
162	Info	Current capacity of the battery is below threshold
163	Info	Current capacity of the battery is above threshold
164	Info	Enclosure (SES) discovered on %s
165	Info	Enclosure (SAF-TE) discovered on %s
166	Critical	Enclosure %s communication lost

Number	Type	Description
167	Info	Enclosure %s communication restored
168	Critical	Enclosure %s fan %d failed
169	Info	Enclosure %s fan %d inserted
170	Critical	Enclosure %s fan %d removed
171	Critical	Enclosure %s power supply %d failed
172	Info	Enclosure %s power supply %d inserted
173	Critical	Enclosure %s power supply %d removed
174	Critical	Enclosure %s SIM %d failed
175	Info	Enclosure %s SIM %d inserted
176	Critical	Enclosure %s SIM %d removed
177	Warning	Enclosure %s temperature sensor %d below warning threshold
178	Critical	Enclosure %s temperature sensor %d below error threshold
179	Warning	Enclosure %s temperature sensor %d above warning threshold
180	Critical	Enclosure %s temperature sensor %d above error threshold
181	Critical	Enclosure %s shutdown
182	Warning	Enclosure %s not supported; too many enclosures connected to port
183	Critical	Enclosure %s firmware mismatch
184	Warning	Enclosure %s sensor %d bad
185	Critical	Enclosure %s phy %d bad
186	Critical	Enclosure %s is unstable
187	Critical	Enclosure %s hardware error
188	Critical	Enclosure %s not responding
189	Info	SAS/SATA mixing not supported in enclosure; PD %s disabled
190	Info	Enclosure (SES) hotplug on %s was detected, but is not supported
191	Info	Clustering enabled
192	Info	Clustering disabled
193	Info	PD too small to be used for auto-rebuild on %s
194	Info	BBU enabled; changing WT virtual disks to WB
195	Warning	BBU disabled; changing WB virtual disks to WT
196	Warning	Bad block table on PD %s is 80% full
197	Fatal	Bad block table on PD %s is full; unable to log block %lx
198	Info	Consistency Check Aborted Due to Ownership Loss on %s
199	Info	Background Initialization (BGI) Aborted Due to Ownership Loss on %s
200	Critical	Battery/charger problems detected; SOH Bad
201	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); warning threshold exceeded
202	Critical	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); critical threshold exceeded
203	Critical	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); further reporting disabled
204	Critical	Enclosure %s Power supply %d switched off
205	Info	Enclosure %s Power supply %d switched on
206	Critical	Enclosure %s Power supply %d cable removed
207	Info	Enclosure %s Power supply %d cable inserted

Number	Type	Description
208	Info	Enclosure %s Fan %d returned to normal
209	Info	BBU Retention test was initiated on previous boot
210	Info	BBU Retention test passed
211	Critical	BBU Retention test failed!
212	Info	NVRAM Retention test was initiated on previous boot
213	Info	NVRAM Retention test passed
214	Critical	NVRAM Retention test failed!
215	Info	%s test completed %d passes successfully
216	Critical	%s test FAILED on %d pass. Fail data: errorOffset=%x goodData=%x badData=%x
217	Info	Self check diagnostics completed
218	Info	Foreign Configuration Detected
219	Info	Foreign Configuration Imported
220	Info	Foreign Configuration Cleared

